PCBNEW Wishlists, Blueprints, Design and Implementation Notes

Brian Bidulock*

September 6, 2010

Abstract

This documents contains some of the design and implementation notes that I have made along the way while making modifications on pcbnew.

Preface

License

Notice

This document is based on KiCad, the open source CAD system.

Optranex Incorporated is making this documentation available as a reference point for the KiCad development project. While *Optranex* believes that these specifications are well-defined in this release of the document, minor changes may be made prior to products conforming to the specifications being made available.

Foreword

Overview

This document provides the feature enhancements, blue prints and wish lists for the pcbnew component of KiCad, the free[1] EDA tool. These enhancements, blueprints and wish lists derive from my desire to use KiCad for a series of high-speed optical network add-in cards (frequencies ranging from 2.5 Gbps to 40 Gbps). These are medium density PCI Express add-in cards. They contain upwards of 20,000 board features. My experiences with manually routing these cards, and attempting a plot-and-go set of fabrication outputs, is what has driven this development.

Purpose

My purpose is making these modifications has been twofold:

- To make KiCad usable for designs up to 40 Gbps. That is, to make KiCad superior for designs up to 40 Gbps.
- To make KiCad usable for modern techniques (such as back-drilling). That is, to bring KiCad up to date.
- To make KiCad generate plot-and-go manufacturing outputs. That is, to make KiCad outputs directly usable for manufacturing processes.
- To make KiCad generate a wider range of manufacturing and intermediate outputs (Gerber, DPF, Excellon-2, IDF, GenCAD, GenCAM, GenX, 258X).

Most of these items affect only pcbnew (and to some degree the module editor and gerbview) in the KiCad suite.

Scope

This document is not meant to be a textbook on manufacturing printed wiring assemblies. Where specific fabrication or manufacturing processes are described, read these passages with the understanding that, although I am an Electrical Engineer, that I am not an expert on fabrication or manufacturing processes. The fabrication and manufacturing requirements listed are taken from fabricator and assembler customer DFM sheets, industry standards (where publicly available), research papers, and university level courses on PCB design and manufacturing. From what I can see on google books, there are many good books written on the topic, by subject matter experts, and I refer you to their guidance. The descriptions in this document were for the purposes of documenting the reasoning behind a specific pcbnew feature, attribute, or parameter, rather than illuminating the process. For further information and detail, see *References* on *page 155*.

Organization

*bidulock@openss7.org

1. As in speech.

Contents

List of Figures

List of Tables

1 Introduction

KiCad has been selected for as the preferred open-source EDA program for several reasons:

- KiCad has an intuitive user interface. This permits a faster learning curve for using the package. This is particularly true for pcbnew. For example, I could never figure out how to use the pcb program from the gEDA suite.
- KiCad uses a simple and straightforward set of ASCII text formatted files to store its data. Programs always have bugs it seems. However, a well documented and straight-forward file format, save in ASCII text, always allows data to be recovered, no matter how bad the problem. Text files can easily be archived using widely available source code control systems. It is actually considered a benefit that these file formats are *not* in XML format. XML format is fine for exchanging data between different implementations of a program; but, for simply storing data, a user-alterable ASCII text format is better.
- KiCad has a 3D viewer that views the final result of the board. The gEDA suite has no equivalent. I cannot stress to much the importance of the 3D viewer. This is fundamental for ensuring that a complex design of mechanical and electrical components are properly placed and spaced on the board. Also, the 3D viewer helps the designer better visualize the resulting board in 3 dimensions: something that is not always possible when looking at only flattened copper layers from the top of the board.
- KiCad always appears to offer ways to get around difficulties when it comes to exceptional design requirements, using the tools available, rather than needing to modify and specialize the code. This indicates that some correct design decisions were taken during its development.

1.1 Fundamentals

1.1.1 Manufacturing Principles

Full scale mass production runs of a product can absorb significant fixed costs, including massive outlays such as pilot plants. Full scale mass production should never be of concern to the KiCad project for a number of reasons:

1. If a production engineering department launching mass production of their product wants some feature from KiCad they can easily build it.
2. Such departments can afford all kinds of tools and methods.
3. Sometimes things are done backward for mass production (such as designing the enclosure *before* the board).

On the other end of the spectrum, prototype and short production runs are sensitive to fixed costs, particularly where the fixed costs recur with every supply order, or recur depending upon the supplier that is chosen, or transaction costs.

Some fixed costs are under control of the supplier and others are under the control of the designer. There are many instances in fabrication and assembly where the fabricator or assembler will presume to manipulate or re-engineer the design data while tooling. This is almost always a manual process. Many board fabricators actually suggest that they perform these manipulations, because they know so much better than the designer, or because the designer cannot possibly comprehend their processes, ... The real reason is that a board fabricator wants to raise the transaction cost of leaving them for another fabricator as high as possible, to lock-in their customers at the prototype stage, and then hold their data captive for the remaining life of the product.

This is so much so that customers looking for short turnaround, short production runs for JITS (Just In Time Supply), have encourage the birth of plot-and-go or build to spec. services.[2]

For the in-house (or really in-house, as in: in-my-basement) prototyping shop, fixed costs associated with manual processes and setup should be avoided. Your time and my time are valuable, and anybody that we are paying a salary to is valuable too. Where a modification to design data would be made by a fabricator, the in-house shop must perform the modifications themselves. They might not be able to afford the CAM systems or other software necessary to perform those modifications.

Hopefully you get the point: to avoid lock-in by one manufacturing sub-contractor, to reduce the transaction costs associated with switching sub-contractors, to avoid tedious or error-prone manual setup, yet to provide the best manufacturing result with minimal capabilities, KiCad should have the ability to generate manufacturing outputs that are as close as possible to those used in the manufacture of the card. KiCad should also be capable of generating design intent outputs (a subset of build to spec) for mass production.

1.1.2 Fabrication Principles

There are a number of instances where board fabricators attempt lock-in or elevated transaction costs. Some examples follow:

- Fabricators recommend that on high-density designs that the solder mask be plotted 1:1, and that the fabricator adjust the solder mask, expanding apertures, performing gang-relief on high-pitch lands, placing solder dams between SMT contacts where necessary and possible, and performing encroachment on vias, an any number of design principles that need to be applied to the primary or secondary mask layers.
- Fabricators will clip legend (silk screen), adjust line widths upward to minimums, even move character fields and adjust character heights. Fabricators will even recommend no silk screen whatsoever (but that is another matter).
- Fabricators adjust stack-ups, choose dielectrics, number and type of prepreg plies, perform resin venting and plating thieving, etch compensation, and a host of other copper and dielectric related variables.
- Fabricators perform panelization, select route paths, perform bevels, choose break-away tab locations, and determine whether holes are routed or drilled, whether slots are drilled or routed, and a host of other mechanical features.
- Fabricators reverse engineer net lists for Gerber verification, DRC checks, and bare-board testing.

Fabricators providing prototype or plot-and-go services will not perform many of these functions. In-house fabrication may not be able to perform these services. The CAD system should enable the designer to perform all of these functions instead. There is no parameter or process that a board fabricator can use successfully that cannot be modelled in the CAD system and specified as *build-to-spec* instead of *design-intent*.

That is, the CAD system should perform full DFM (Design for Manufacturing) analysis.

1.1.3 Assembly Principles

Turn-key, kitted, or short production run assembly shops are much better at avoiding fixed costs, both for themselves and their customers. They really want your assembly work to keep their lines running. They avoid custom tooling at every step. These shops are also better at specifying their data requirements as close to build as possible.

2. I is my wife's opinion that the US Department of Defense has helped in the latter regard: build-to-spec.

Some turn-key and short production run assembly shops will even demand that board be supplied by their approved board fabricator just to avoid throw-it-over-the-wall issues with board fabricators. It just goes to show how unenlightened many board fabricators are to the needs of assembly.

The CAD system should perform full DFA (Design for Assembly) analysis.

1.1.4 Test Principles

The CAD system should perform full DFT (Design for Test) analysis.

2 Internal Changes

KiCad has evolved from a simple set of programs through incremental change. There come a point in the life-cycle of any large software project where it cannot grow organically. This first set of requirements that are orthogonal to the original design require changes at the level of the rewrite for further growth. As any large project including user interface, the investment of users into the current or historical user interface in terms of learning and understanding is a heavy burden. For the most part, drastic changes should be internalized and hidden from the user interface where possible. Proper MCV (Model Controller View) architectures permit the model to be changed without affecting the view or the control. Unfortunately, KiCad was not designed closely to the MCV principles.

To meet the objectives laid out in this document, a number of changes were made to pcbnew to circumvent initial design shortsightedness and to release the program from legacy implementation. This section describes the rationale behind some of those changes.

Performance. Non-scalable performance is the major reason that I started making changes to pcbnew. I have a board design with several large BGAs. Wishbones to ground and supply planes number in the thousands. There are over 500 1uF ceramic chip decoupling capacitors, of which all of them connect one terminal to the ground plane. There are approximately 22,000 features on the board. When upgrading to the latest stable release, editing the board ground to a halt on very fast (3GHz) dual core processor, making pcbnew largely unusable for editing the board. When investigating the performance difficulties, I discovered that the entire list of board items was being sorted and linearly searched again and again for the simple task of adding a dogbone to a ground via. So I ripped pcbnew apart and added some constant time scalable design patterns to circumvent the scalability and performance issues.

See *Sec. 2.1(10)* for details of software design patterns and changes made to improve performance.

Fabrication Output. Having to cram several hundred termination resistors between two sets of SSTL buses between FPGA and SDRAM chips in a fine pitch BGA, having to length-match all of the stubs and place the $200 \times 51\Omega$ chip resistors, I found myself running out of room and reaching for blind or buried vias to handle the routing. The board is restricted to 0.062" thickness and I was already at 14 layers. To my chagrin, I discovered that the BBV outputs generated by pcbnew are unusable by a fabricator. Outputting drill files for each layer pair is simply not the way that it is done. Also, some of the techniques for high-speed design (back-drilling) and some BBV fabrication methods (depth-control-drilled blind vias) were simply not supported. Really, only through vias are supported in pcbnew. Looking into fabrication outputs, I was reaching for CAM outputs such as [GenCAM], [GenX] and [258X] as well as addressing problems with the generation of NC drill files and Gerbers. To address some of the modelling, I added support for these as well as [DPFv7] and IDF [Kehmeier and Makowski, 1998]. Better support for [GenCAD] and support for IPC-D-356.

See *Sec. 2.2(12)* for an overview of the software changes necessary to support new and improvded fabrication outputs.

High-Speed Design. The board design I was working on required laying out quite a few 666Mhz SSTL logic traces, 622MHz PECL clocks, 622MHz LVDS differential pairs, and 2.488GHz CML differential pairs, and 2.5GHz PCIe lanes. I found myself needing better tools for figuring out the track impedance, propagation delay differences between the various stubs of the SSTL

traces, and needing to control both microstrip and stripline configurations. Because I manually laid out the board[3], I was looking to apply approximate calculations. Things such as copper roughness, etch factor, dielectric constant and thickness, copper weights, press factors, plating factors, via tuning, back-drilling to remove stubs, ground plane stitching, isolation traces, swept traces, minimum arc, maximum decoupled segments, segment by segment maximum mismatching: none of these did pcbnew help with. The new netclasses stuff that was added looked like a promising start for real characterization of busses, signal logic families, differential pair rules, and the like. I needed to extend the netclasses and net information to support these high-speed design features. Additional high-speed design features having to do with removing parasitic capacitance by removing non-functional pads, ground-plane cutouts under high-frequency pads, back-drilling, and the like required better approaches to fabrication outputs.

See *Sec. 2.3(12)* for details of feature additions to support high-speed design.

DRC. When investigating the design guidelines of a number of board fabricators and assembly houses, I discovered that pcbnew was performing only rudimentary trace and space calculations. For example, placing a trace over a board edge did not raise an error: not even a warning. Almost none of the DFM guidelines proclaimed by fabricators and DFA guidelines proclaimed by assembly shop were included in the checks. This, coupled with the poor fabrication outputs, would easily result in a board design that I spent months on hand-routing traces that simply could not be manufactured, nor assembled. As most of these DFM and DFA rules are mathematical calculations, it seemed like a good job for a computer program to me.

Early on when starting to use pcbnew, I relied heavily on it ability to layout differential pairs by specifying the clearance and then using the DRC while moving features to slap the pair up against its mate. I was using 5 mil traces with a 7 mil gap (12 mil spacing) which is pretty common for PCIe lanes, OC-48 CML traces, high frequency LVDS pairs and the like. To my chagrin, after laying these out, a DRC check insisted that each one of these traces was too close to the other. I had to set the clearance to 6.9 mils to keep the DRC from triggering on each corner. Problem: precision of calculation of the square root of 2 with an integer precision of 2 digits (decimil). Integer math (and trig functions) in one place, and full double floating point calculations in the other. Result: change everything to floating point, factor it all down to basic geometric shapes. The track layout was rounding up (placement of cursor), and the DRC was rounding down, or visa versa. Part of this is units. Nevertheless, incorporating DFM rules and full floating point DRC calculations was called for.

See *Sec. 2.4(12)* for deficiencies and new software approaches to improve DRC handling.

Ratsnests. What a pig. It works find when there are only a handful of pads in a net. But when there are 2000 BGA balls strewn across the board, everything grids to a halt. I rotated a module (just a JTAG connector), and, because it had a ground pin, it took 3 minutes to respond. Problem: ratsnest calculations were recalculating the ratnest relation for each net for each pad in the rotated module. But just about every module has a ground pin. When a JTAG connector is rotated, *only the pads that are moving need ratnests recalculated,* ***not every pad in the net***. So, I reworked ratsnest calculations to use a real Lee algorithm for pads that move using the constant-time proximity maps. The three ratslines on edited tracks too.

See *Sec. 2.5(13)* for details on ratsnest calculation changes.

Layers. I had some long trials with technical layers:

- Because I have some pretty high-density stuff underneath fine pitch BGAs, I need to reach for something like dry-film, via capping, or hole plugging. Also, I was considering using Via-Through-Pad. However, pcbnew had none of the technical layers necessary to support selective via capping or hole filling.
- Also, I have some of those thermal pad QFN packages that really need conductive filled thermal vias. Again, hole-filling technical layer.
- I considered get rid of the $200 \times 51\Omega$ chip resistors used for the SSTL termination by using embedded resistors. Yet again another missing technical layer: resistor secondary etch mask.
- I was jamming 0.1uF ceramic chip decoupling capacitors on the back of the board under the BGA fields and discovered that silk was not a good courtyard. Also, because they are all terminated on one terminal to V_{CC} and the other to V_{SS}, they could be placed right on top of each other and DRC didn't care! Enter the lack of an IPC CAPC1005L footprint, and you get the point: Courtyard layer.
- Trying to avoid the impedance variability of solder mask over 2.5GHz microstrip, I considered a common approach: take the solder mask away and gold plate them instead. Then you need peel-off masks and a gold mask for the area. Two more layers: peel-mask and finish (gold mask).
- I have those QFNs (synchronous bucks and programmable regulators) that dissipate all their heat into the board. Well, solder mask is a good thermal insulator too. Where is the heat to go? Stencil printed board heat sinks is the cure! Another layer.
- After trying to cram in all of the silk for those capacitors and resistors and drawing boxes of reference designators, I though it would be a good idea to print the component outline on assembly layers for documentation (particularly after the board fabricator clipped away all of the silk anyway). No can do. Another layer: component outlines.

Looking at [GenCAM], all of these layers are supported in the CAM file format. But pcbnew doesn't know about them. Problem: 32-bit bit mask with no room left. So, I built a completely new mask class for layers and provided for a largely unlimited number of copper layers. This was not easy.

See *Sec. 2.6(13)* for details on the need for new layers and the software design patterns used.

Internal Units. Most of the round-off error problems and lack of accuracy for metric gridded things (like, gee a mix of BGAs with imperial and metric pitches). This causes a problem when working on the metric grid. Everything is off. Type in a coordinate and it jumps up or down on you. Having big BGA devices (with over 1000 balls) on a metric grid, the dogbone results in these freaking little 1 mil segments jumping from the pad to the 45° dogbones. Several thousand of the on the card. No track operation to get rid of them. Not only that, these freakin' little things cause the trace to shift over slightly from the correct 45° position, resulting in (you guessed it) DRC errors on every dogbone under every BGA. Not to mention that they almost double the number of copper features on the board, and mostly on the ground and power nets where the scalability problems were. So, I tried just unlocking the 45° degree tracks and dragging that little 1 mil segment to 0 mil and then cleaning them away. Guess

3. Real men lay out their own boards: wimps use autorouters.

what? The dogbone moved! DRC violations on every dogbone. Solution: use a proper high precision internal unit. Changed from decimils to nanometers. Now everything on a metric or imperial grid down to 10μ" or 0.1 microns locks to a hard grid with no round-off. This was not easy either.

What do you thing is next? Angles of course: 1/10th of a degree just doesn't cut it.

See *Sec. 2.7(15)* for details on the changes to internal units.

Modules. The first time I looking in a .brd file and saw my 500 chip capacitors with their entire module definition repeated over and over again, I though it was a little ridiculous, but didn't think much more of it until I attempted to output [GenCAM], [GenX], [258X] and IDF [Kehmeier and Makowski, 1998] file formats. Those formats always separate package (footprint, decal) from placement. The package is output once and then referred to by each placement. Now how was I going to do that? Refactor modules every time one of these files was read or written? Don't think so. So, the internal module format was redesigned to split package from placement.

See *Sec. 2.8(16)* for details on fundamental changes to the representation of modules.

2.1 Performance

The internal structure of pcbnew grew over time and evolved from a "C" language program. Unfortunately, cruft has resulted in linear (or worse) searches being performed on what can be rather large lists of items. Design patterns to avoid this difficulty in the future and cure the performance problems resulting from the current scheme have been installed. This section describes some of these design patterns.

Proximity Maps. Historically, when pcbnew was performing DRC and other proximity tests (right-click, ratsnest), it was iterating through entire lists of items performing hit tests on points and complex geometric calculations. This was particularly wasteful of resources, quite time consuming, and did not scale at all. It produced staggering and long delays while editing a complex board. The proximity map pattern is a solution to this problem.

See *Sec. 2.1.1(10)* for details on proximity maps.

Connect Maps. Historically, when pcbnew was testing for connectivity at a connection point between two connectible items, it was iterating through entire lists of items performing match test on all of the connection points for the item. Also, of particular difficulty was checking for connectivity to copper zones. This was extremely wasteful of resources, was time-consuming and did not scale. It produced staggering and long delays while editing a complex board. The connect map pattern is a solution to this problem.

See *Sec. 2.1.2(10)* for details on the connect map design pattern.

Dirty Classes. Historically, pcbnew has maintained a set of dirty flags for the board and performed sorting and reordering of all items on the board when changes were made. Full sorts were performed, regardless of the number of items changed. In fact, in many instances full sorts were invoked by single changes. A lack of understanding of the requirements for updating when a change was made to a single item often resulted in simply marking the entire board dirty and performing several full sorts. This has turned out to be extremely wasteful of resources, does not scale, and produces staggering and long delays while editing a complex board. The dirty class pattern is a solution to this problem.

See *Sec. 2.1.3(11)* for details on the dirty class design pattern.

Connection Graphs. Historically, when pcbnew was checking for connectivity between copper layer features, it was iterating through net-code sorted lists of items looking for connections. In particular cases (ground nets in particular), it was quite wasteful of resources, time consuming, and did not scale. It produced staggering and long delays or hesitation while editing a complex board. The connection graph pattern is a solution to the problem.

See *Sec. 2.1.4(11)* for details on the connection graph design pattern.

2.1.1 Proximity Maps

The file include/class_proximity.h contains definitions and implementation for classes PROXIMITY_GRID and PROXIMITY_ITEM.

Previously pcbnew used to iterate through entire lists of items making distance calculations to determine the location of an item or feature. This was extremely wasteful of resources and time consuming. The proximity map design pattern was used to avoid this waste.

The proximity map design pattern breaks the board area (both inside and outside board edges) into a grid of fixed-sized cells. The grid is represented as an unordered multi-map that uses a key that contains the grid location of a cell in the proximity grid and a layer specification indicating the layers on which the item may be found. The value contains a pointer to the proximity item occupying the cell.

A proximity item is an object that can occupy cells in a proximity grid. It contains a pointer to the proximity grid to which it belongs, and a set of references to the cells (entries in the proximity map) that the proximity item occupies. Whenever a proximity item moves or is added to or removed from a board, the proximity grid is updated with the new locations (or, when being removed, the locations are removed). A proximity item is deemed to "occupy" a cell whenever a point exists on the proximity item that also exists within the proximity cell.

To check whether a given area is occupied by a proximity item is a simply matter of looking up the proximity grid cells that area occupies. Once these items that are "near" to the area are iterated, the precise relationship between the area and the items can be determined (calculated). The lookup of the proximity cells uses a constant-time hash function that is not dependent upon the number of items on the board (it is only dependent on the density of items on the board).

This pattern is very useful for DRC/DFM checks, tests for copper connectivity or clearances, as well as locating functions such as "right-click" and bounding box redraws.

It took a significant coding effort to replace all of the location functions in pcbnew with the optimizations provided by the proximity map pattern.

2.1.2 Connect Maps

The file include/class_connect.h contains definitions and implementation of classes CONNECTION_MAP and CONNECTION_ITEM.

Previously pcbnew used to iterate through entire lists of items making point comparisons to determine whether items or features contained the given connection point. This was extremely wasteful of resources and time consuming. Also, the approach did not scale. The connection map design pattern was used to avoid this waste.

The connection map design pattern tracks connection points on the board (both inside and outside board edges). The map is represented as an unordered multimap that uses a key that

contains the point location of a connection point for an item and a layer specification indicating the layers on which the connection point exists. The value contains a pointer to the connection item that has the connection point.

A connection item is an object that can connect at a number of points. It contains a pointer to the connection map to which it belongs, and a set of references to the points (entries in the connection map) at which the item can connect. Whenever a connection item moves or is added to or removed from a board, the connection map is updated with the new points (or, when being removed, the points are removed).

To check whether a given point connects to other items is a simply matter of looking up the point in the connection map. Once these items that can connect at the point are iterated, the precise relationship between the point and the items may be determined (located). The lookup of connection map points uses a constant-time hash function that is not dependent upon the number of items on the board (only the density of items) and scales well.

This pattern is very useful for determining the intended connection between items for determining copper connectivity, connectivity of board edge segments, connection of draw segments, etc.

It took a significant effort to replace all of the location functions in pcbnew with the optimizations provided by the connection map pattern.

2.1.3 Dirty Classes

The file include/class_dirty.h contains defintitions and implementation for class DIRTY, a class for things that can get dirty and need to defer updates, or clean themselves.

Previously pcbnew used to maintain a number of dirty flags on a board-wide basis. At many points in the program this resulted in every item on the board board being updated in a very inefficient way. For boards that contained upwards of 20,000 features, the resulting performance was unacceptable. Much of the performance hit had to do with sorting pads and tracks and updating ratsnests.

The addition of a full connection graph and proximity maps helped to removed the need for sorted pads and tracks. Nevertheless, it is not necessary to update the entire connection graph or proximity map with each change in a board item: only the dirtied item needs to be updated the lists, directed graphs and maps, and only when the item is required to be updated. This prompted the addition of a design pattern for deferred update of board items for various goals. The DIRTY class in this file implements that design pattern.

eeschema does not suffer from the same performance problems, or, at least the performance problems associated with eeschema can be managed by the user. eeschema allows the schematic to be organized into a hierarchal set of sheets. By limiting the maximum number of items on a hierarchal sheet, eeschema basically implements a design pattern equivalent to proximity maps. By limiting the editing focus to one hierarchal sheet at a time, the range of updates is also limited. Therefore, an eeschema schematic can grow quite complex without impacting performace.

pcbnew is a different situation. All of the board items are contained on the same sheet. Thus a proximity map, connection graph and deferred updates using this DIRTY class are required.

DIRTY Class Class DIRTY is a class for things and their dependent children that can get dirty. This class is currently used by pcbnew and friends. It is not yet used by eeschema as eeschema's performance difficulties are manageable by the user.

Each dirty class is organized into a rooted tree of dirty classes. For the most part in pcbnew, the root of the tree is the BOARD class. When an item within the class heirarchy becomes dirty, it places itself on it's parent's dirty children list and the parent is marked dirty. The parent then places itself on its parent's dirty list, and so on until the root of the tree. When an item is to be cleaned, all of its dirty children are cleaned first and then the item itself is cleaned (by default, a virtual function is provided for alternate behavior). When an item cleans itself, it removes itself from its parent's dirty children list and marks itself clean. If the parent has no more dirty children, it too is marked clean. In the process of cleaning itself, other items to which the cleaning item has a relationship might be marked dirty. Thus, when a board is cleaned, only its dirty children are cleaned, and any other items that require cleaning as a result of the cleaning process itself.

The design pattern is further enhanced by considering different categories of dirtiness, as follows:

Presence (0x00001), presence has changed, the item needs to be fully added or deleted.

Geometry (0x00002), geometry (but not position) of the item has changed.

Position (0x00004), position of the item has changed.

Clearance (0x00008), clearance of the item has changed.

Proximity (0x00010), proximity of the item to other objects has changed.

Layer (0x00020), layer of the item has changed.

Display (0x00040), display of the item has changed.

Connect (0x00080), connections to or from the item must be rebuilt.

Fullrats (0x00100), full ratsnests need to be updated for the item.

Ratsnest (0x00200), active ratsnests need to be updated for the item.

Subnet (0x00400), subnets need to be recalculated for the item.

Net (0x00800), nets need to updated (propagated) for the item.

Filling (0x01000), zone filling needs to be udpated.

Names (0x02000), name change needs to be updated.

Classes (0x04000), net class change needs to be updated.

Appearance (0x08000), the displayed appearance of the item has changed.

Drc (0x10000), the items Drc has changed.

2.1.4 Connection Graphs

The file include/class_graph.h contains definitions and implementation of the CONNECTED class.

Previously pcbnew used to iterate through net-code sorted lists of copper layer items using connectivity analysis to determine whether board connected items where attached. In a number of cases these iterations where wasteful of resources and time consuming. The solution did not scale. A particular problematic case was that for ground nets. When large BGA devices are contained on the board, the number of ground and power net connections can soar into the thousands. Reiterating through the net-code sorted list of connections exhibited really poor performance. Adding dog-bone breakouts to ground on a large BGA field, would stagger and delay longer with each addition. To avoid this problem, the connection graph design pattern was used.

The connection graph design pattern simply collects all connectivity items into a directed graph structure. Each connected item has a set of connections for each possible connection point.

Whenever a connected item moves, is added, or deleted, the connection graph is updated. A set of non-looping, depth-first, recursive traversal functions where used for updating the connection graph on demand.

This pattern is very useful for maintaining the connectivity of the board. It is useful both for copper layer features, as well as connectible items on other layers.

It took a significant coding effort to replace all of the connectivity functions in pcbnew with the optimization provided by connection graphs.

2.2 Fabrication Output

2.3 High-Speed Design

2.4 DRC

Historically, pcbnew has performed only the most rudimentary copper clearance DRC checks. These checks were not sufficient for the application of simple space and trace rules. Many deficiencies exist.

2.4.1 DRC Deficiencies

There were a number of deficiencies in DRC calculations. These deficiencies were not the fault of the method of calculation nor the internal units, but resulted from having insufficient information retained as to all of the clearances and DFX rules that are necessary for full plot-and-go board fabrication.

The following clearances are required:

Hole to Pad.

Track to Hole Edge (non-plated hole).

Track Width.

Track to Track Spacing.

Track to Pad Spacing.

Pad to Pad Spacing.

Track to Routed Board Edge.

Track to Scored Board Edge. It is important to provide track-to-edge clearance on all layers for multiple reasons:

1. To avoid parital or complete removal of tracks. Apart from the obvious problem associated with such an occurence, partial removal will result in exposed copper on the edge of the board causing:
   - possible shorting, if in contact with metal case or guide.
   - possible shorting of inner layer ground and power plane from resulting burrs.
   - oxidized copper over time especially in corrosive or humid environments.
2. To enable complete encapsulation of the tracks with solder mask which also improves the adhesion of the mask on the edge of the board.

Track to Bevel.

Hole Edge to Routed Board Edge.

Hole Edge to Scored Board Edge.

Hole Size.

For internal power and ground planes:

Hole to Pad.

Connection Track (thermal relief).

Plane to Hole Edge (plated).

Plane to Hole Edge (non-plated).

Plane to Pad.

Plane to Routed Board Edge.

Plane to Scored Board Edge.

For Legend:

Legend to Hole.

Legend to Pad (SMT).

Line/Text width

For Solder Mask:

Mask to Track.

Mask to Pad.

Mask Web (Dam, Wall).

Mask to Routed Board Edge.

Mask to Scored Board Edge.

Same Net Clearance. The clearance between a track and itself is considered a feature. pcbnew only checks clearance between different nets. It is possible to place a track to close to itself (or another same-net pad or track) to be properly resolved during etching or plating.

Board Edge, Unplated Slots, NPTH to Trace or Pad Clearance. There are a number of reasons why conductors should not be placed close to board edges, unplated slots and NPTH, particularly if any of these items is exposed to the plating process (e.g. board edges are plated). This is because electroplating can travel in the gaps between resin and glass in the internal laminate and form shorts between the conductor and any metal layer inserted into the hole.

On outer layers, the primary mask must be withdrawn from the routed or scored board edge to prevent primary mask debris from being generated that might stick to pads or peel from the board. Traces and via pads on outer layers must be withdrawn from this primary mask opening to avoid unwanted exposure of traces and vias to finish, HASL, solder wave, reflow or rework.

Board Edge, Unplated Slots, NPTH to Primary Mask Clearance. The main reason for keeping primary mask away from routed or scored board edges, unplated slots, cutouts and wells, and non-plate-through holes is to avoid peeling and flaking of the primary masks during or after milling or drilling. To keep the tool off of the edge of the primary mask, it is necessary to compensate for both tool deviation and registration as well as primary mask registration. Both must be included in the clearance.

I have several boards here in the office (PCI add-in cards). Examining them, only one in four pull the primary mask back from the routed or scored board edge. This is a task that normally performed by fabricators in under full-service, but must be performed by the CAD system under plot-and-go service.

Whenever the primary mask is retracted from some part of the board, copper features on outer layers must also be kept out of this area; otherwise, they will be exposed beyound the primary mask causing problems (shorting) during the application of finish, and during the assembly process.

SMT Pad to Via Clearance. There is a clearance between an SMT (Surface Mount Technology) land requiring conductive paste and reflow, and a via or other hole that can suck the liquified paste down the hole under reflow or rework. Also, a pad represents an opening in the primary mask and so does the via hole. Misregistration of the primary mask could cause solder bridging with finishes such as HASL, wave soldering, or during reflow or rework.

NSMD Pad to Via Clearance.

Contact to Via Clearance.

Contact to Track Clearance.

Track to Via Clearance.

2.4.2 DRC Enhancements

On of the most annoying things that I found when initially using pcbnew was its inability to perform proper DRC calculations (considering round-off). When setting differential pairs for a 7-mil separation (5-mil gap), DRC would let the traces be laid on the board (using the intrusion feature) and the when a DRC check was run, it would complain that every on of these tracks was too close together at the ends. Obviously DRC calculations were not being performed correctly, or the approach to round-off error was unwise, or too conservative, or in error.

Several changes were made to address this.

1. Internal unit was increased in precision from 1/10,000th of an inch to 1/1000,000th of a millimeter.

   This change was necessary to have high precision on smaller baord items. A 5-mil separation was only 50 units in the old internal unit system, but is 127,000 units in the new internal unit system. Consider this size of rounded end on a 5-mil trace. The radius is 25mil. The point at a 45-degree angle from the center of the end is at point (35.355339, 35.355339), so the choice is either to round up or down. Rounding up could result in a DRC failure where there is none; down, miss where there is one. This is either a 12% error or a 24% error depending on which direction is chosen.

   The nanometer internal unit changes this to point (89,802.561, 89,802.561) providing much greater precision. The difference when rounding up or down is 5 ppm.

2. Integer trigonometric functions changed to perform full double floating-point calculations.

   This change was largely made possible by the increase in precision of internal units as well as the use of proximity maps (so trigonometric functions are no longer required on every item on to board to find a single item that is "close").

2.5 Ratsnests

2.6 Layers

Previously pcbnew was using a 32-bit mask for layers. This restricted the total number of copper and technical layers. Only 3 layers were remaining as undefined and there were 8 or more layers to add.

2.6.1 Layer Identifiers

In considering expanding the internal layer management, one of the options considered was to simply expand the mask to a 64-bit mask for layers. One of the difficulties with a 64-bit mask is that the long long type necessary to represent 64-bit numbers on 32-bit architectures is: not part of the C++ standard (nor the upcoming revision) and 64-bits would simply introduce a new limit not far from the last one. Therefore, this approach was discarded.

The approach taken was to represent a layer identifier as two 16-bit integers. The most significant 16-bits form the layer class and the least significant 16-bits form the layer index within the layer class.

2.6.2 Layer Classes

Historically, pcbnew did not provide a number of fundamental technical layers and imaging capabilities. These required the addition of new layer classes. The new layer classes are:

- *Dielectric.* A layer class where each layer represents a dielectric layer in a laminate.
- *Resistance.* A layer class that provides imaging for embedded resistors.
- *Plating.* A layer class that provides imaging for flash and button plating for internal layers.
- *Hole Fill.* A layer class that provides imaging of hole filling stencil apertures.
- *Keep Out.* A layer class that identifies keep out and keep in areas for components, traces, vias, thieving, venting, and other PWA (Printed Wiring Assembly) features.
- *Via Plug.* A layer class that provides imaging for secondary mask via capping.
- *Contacts.* A layer class that provides imaging for gold and carbon contacts.
- *Peel Mask.* A layer class that provides imaging for removable masks (Corfin and peel-off mask).
- *Finish.* A layer class that provides imaging for selective or secondary finishes.
- *Heat Sink.* A layer class that provides imaging for board heat sink stencil apertures.
- *Courtyard.* A layer class that provides design courtyards for components.
- *Component.* A layer class that provides component outlines and primary assembly drawings.
- *Coating.* A layer class that provides imaging for conformal coatings.
- *Probe.* A layer class that provides imaging of pads and test points for ICT or bare-board testing.
- *Fixture.* A layer class for the imaging of test probe fixtures for bed-of-nails and clamshell testing.

These new layer classes and the roles of the existing layer classes are discussed below:

Dielectric:– The dielectric layer class (in the board fabrication category) is used to describe all of the dielectric layers of the board. Attributes include thickness, dielectric constants, thermal conductivity and other characteristics.

Dielectric layers were not historically represented by pcbnew.

See *Sec. 5.1(59)* for details on dielectric layers.

Resistance:– The etch layer class (in the board fabrication category) represents copper that is etched away from copper layers to expose the resistance ply underneath for embedded resistors. Plotted data includes etch masks for secondary copper etching.

Resistor layers were not historically represented by pcbnew.

See *Sec. 5.2(61)* for details on plating layers.

Copper:– The copper layer class (in the board fabrication category) is used to describe all of the copper layers on the board. Additional attributes associated with this layer include the copper weigth, thickness, roughness, xy-plane registration, and other attributes not previously tracked by pcbnew.

pcbnew has historically represented these layers; however, it has not represented the process characteristics nor capabilities necessary to fully fabricate layers in this class.

See *Sec. 5.3(62)* for details on copper layers.

Plating:– The plating layer class (in the board fabrication category) represents plating that is applied to copper layers, such as button, flash-and-button, or flash plating to internal copper layers during blind and buried via construction. Plotted data includes plating masks for button plating.

Plating layers were not historically represented by pcbnew.

See *Sec. 5.4(65)* for details on plating layers.

HoleFill:– The hole filling layer class (in the board fabrication category) represents hole filling that is applied to holes through the laminate or sub-laminate constructions. Plotted data includes selective stencil openings for hole-filling processes.

Hole filling layers were not historically represented by pcbnew.

See *Sec. 5.5(66)* for details on hole filling layers.

KeepOut:– The keep out and keep in layer class (in the board fabrication, board test, assembly, assembly test and fixture categories) represent keep out and keep in areas for traces, via, thieving, components, and other items. Keep out and keep in layers do not normally result in post-processed fabrication or assembly outputs. Data from this layer can be imported from or exported to mechanical CAD systems. It is also possible to plot data from this layer on assembly drawings.

Keep out and keep in layers were not historically represented by pcbnew.

See *Sec. 5.6(69)* for details on keep-out layers.

ViaPlug:– The via plugging or capping layer class (in the board fabrication category) represents secondary mask that is applied to vias on the top or bottom side to tent, cap or plug the via. Plotted data includes photoimagable masks for the secondary mask process.

Via plugging (or capping) layers were not historically represented by pcbnew.

See *Sec. 5.7(73)* for details on via plugging layers.

SolderMask:– The primary mask or solder mask layer class (in the board fabrication category) represents the primary mask applied to either side of the board.

pcbnew has historically represented this layer; however, it has not represented the process characteristics and capabilities necessary to fully fabricate layers in this class. Also, pcbnew has not represented the characteristics of the layers necessary for performing embedded microstrip impedance calculations.

See *Sec. 5.8(77)* for details on primary mask layers.

Contacts:– The contacts layer class (in the board fabrication category) represents edge-finger, gold plated or carbon contacts. The contact layer class is separated from the finish class because they are treated separate from finished pads. Plotted data for this layer class includes plating masks for contacts or mask for dry film for the application of carbon. Mutliple contact finishes can be supported for a given board.

Contact layers were not historically represented by pcbnew.

See *Sec. 5.9(79)* for details on contacts layers.

PeelMask:– The peel mask or removable mask layer class represents the removable (peelable) mask applied to either side of the board. Peelable mask can be used to protect areas of the board from wave solder or other processes during assembly. This can be used to protect edge-fingers during board finish application as well as providing protection from wave soldering during the assembly process. Peelable mask can be removed during board fabrication, or left to protect items during assembly, to be removed at an intemediate or final assembly stage.

Plot data from this layer provides artwork for photoimagable peelable masks. This layer data is positive layer information by default.

Peel mask or removable mask layers were not historically represented by pcbnew.

See *Sec. 5.10(81)* for details on removable mask layers.

Finish:– The finish layer class represents the finish or finishes that are applied to the exposed copper surfaces of the board. Plotted data includes plating masks for application of selective finishes. Multiple finishes can be supported for a given board.

Board finish layers were not historically represented by pcbnew.

See *Sec. 5.11(82)* for details on selective finish layers.

SilkScreen:– The legend or silkscreen layer class represents the legend applied to either side of a board.

pcbnew has historically represented this layer; however, it has not represented the process characteristics and capabilities necessary to fully fabricate layers in this class.

See *Sec. 5.12(83)* for details on legend layers.

HeatSink:– The heat sink layer represents board level heat sinks formed with stencil printed non-conductive thermal paste. Plotted data includes the stencil apertures for printing the thermal paste.

See *Sec. 5.13(89)* for details on heat sink layers.

SolderPaste:– The paste or solder-paste layer class represents the paste applied to either side of the board for reflow assembly.

pcbnew has historically represented this layer; however, it has not represented the process characteristics and capabilities necessary to fully fabricate layers in this class.

See *Sec. 5.14(90)* for details on solder paste layers.

Adhesive:– The glue or adhesive layer class represents the adhesive or glue dots applied to either side of the board to retain components for wave soldering or double-side reflow assembly.

pcbnew has historically represented this layer; however, it has not represented the process characteristics and capacbilities necessary to fully fabricate layers in this class.

See *Sec. 5.15(96)* for details on adhesive layers.

Courtyard:– Courtyard layers do not normally result in post-processed fabrication or assembly outputs. Data from this layer can be imported from or exported to mechanical CAD systems. It is also possible to plot data from this layer on assembly drawings.

Courtyard layers were not historically represented by pcbnew.

See *Sec. 5.16(97)* for details on courtyard layers.

Component:– Component layers do not normally result in post-processed fabrication or assembly outputs. Some data from this layer can be imported from or exported to mechanical CAD systems. It is also possible to plot data from this layer on assembly drawings.

Component layers were not historically represented by pcbnew.

See *Sec. 5.17(104)* for details on component layers.

Coating:– *Coating layers were not historically represented by pcbnew.*

See *Sec. 5.18(106)* for details on coating layers.

Edges:– The edges layer class represents the board edges. This includes the exterior board edges (route, scores, bevels, perforations) as well as any interior cutouts, slots or wells. Data plotted for this layer includes artwork for board edges, cutouts, slots and wells, following the industrty practice of developing route paths from Gerber plots.

pcbnew has historically represented these layers; however, it has not represented the process characteristics nor capabilities necessary to fully fabricate layers in this class.

See *Sec. 5.19(107)* for details on edges layers.

Eco:– *pcbnew has historically represented these layers; however, it has not represented the process characteristics nor capabilities necessary to fully fabricate layers in this class.*

See *Sec. 5.20(112)* for details on ECO layers.

Probe:– The probe layer class represents pads and other probable areas. Data plotted for the probe layers include the "pad-master" for both sides of the board. The pad master could include exposed vias and other exposed coppper areas.

Probe layers were not historically represented by pcbnew.

See *Sec. 5.21(113)* for details on probe layers.

Fixture:– *Fixture layers were not historically represented by pcbnew.*

See *Sec. 5.22(114)* for details on fixture layers.

Comment:– The comment layer in pcbnew had an unknown purpose. This layer was never used for library modules and was typically not plotted. The only use that can be conceived is for designer annotations or callouts.

pcbnew has historically represented these layers; however, it has not represented the process characteristics nor capabilities necessary to fully fabricate layers in this class.

See *Sec. 5.23(115)* for details on comment layers.

Drawing:– The drawing layer provided the fabrication print, but did not permit detail sheets, nor assembly drawings.

pcbnew has historically represented these layers; however, it has not represented the process characteristics nor capabilities necessary to fully fabricate layers in this class.

See *Sec. 5.24(116)* for details on drawing layers.

2.6.3 Layer Handling

Although all plot formats used by pcbnew support scratching, clip area, clear overlay or knockout capabilities, pcbnew has only historically provided a single dark positive image for any layer.

(R) 1 (pcbnew) *Enhancements have been provided to permit pcbnew to support negative overlays on all layers.*

2.7 Internal Units

Previously pcbnew was using 1/10th of a mil (1/10,000th of an inch) as an internal unit. This caused several problems. All modern PCB CAD systems use the fact that any grid point on an imperial scale can be represented precisely with two more digits of precision in metric units. A decimil (the common unit of RS-274D 2:4 plots), can be represented precisely with a 10 nanometer precision (254 ten-nanometer units is precisely one decimil).

(R) 2 (pcbnew) *The internal unit will be changed to 1 nanometer making precise 3:6 metric RS-274X plots and precise 2:5 imperial RS-274X plots possible.*

2.7.1 Rectilinear Units

Historically pcbnew has used 1/10,000th of an inch (0.1mil or 1 decimil) as an internal unit. While this is accurate enough, it causes problems for components that are built to a metric grid. Also, it causes problems for DRC and subsequent DFM calculations. All modern CAD packages use the fact that there are an integer number of tenths of a millimetre in an inch, and 1 decimil is 2,540 nanometres. This means that an internal unit of 10 nanometres will handle both a hard metric and hard inch grid.

(R) 3 (pcbnew) *pcbnew's internal unit has been changed from deci-mils to nanometres.*

This change is not directly visible to the user: it is a question of the internal representation of dimensions and coordinates; however, the following annoying behaviours have disappeared:

- Items set on points on a metric grid would not be at the precise location, causing differences in distances between, for example, BGA pads laid out in a metric grid pattern.
- Metric grids would be represented inexactly when displayed in imperial units. (Metric grids are now always displayed in metric units: duh!)
- Imperial grids would be represented inexactly when displayed in metric units. (Imperial grids are now always displayed in imperial units: duh!)
- Typing in a precise value for coordinates (such as module x or y position coordinates) would change to a different value (nearest rounded off decimil) after entering.
- Gerbers were always plotted with imperial 2:4 precision. (Now they are plotted 2:6 for imperial and 3:5 for metric, by default.) NC drill files were always generated with imperial 2:4 prevision. (Now they are generated 2:6 for imperial and 3:5 for metric, but always use decimal notation, by default.)

Although this change sounds simply, poor coding practise within the KiCad code base made the change extremely difficult. Many explicit non-symbolic constants were codified and the expectation that the units would be deci-mils. It may take some time to rout out the last occurrences.

2.7.2 Angular Units

Historically pcbnew has used 1/10th of a degree (0.1 degrees) as an internal unit for angles. While this is accurate enough for some purposes, it is really not accurate enough to represent general arcs using the centre-end-angle approach.

There are two ways that arced traces are represented in files:

1. *Start-End-Centre-Direction:* Four items are specified:
   (a) The start point of the arc.
   (b) The end point of the arc.
   (c) The centre point or pivot point for the arc.
   (d) The direction (clockwise or counterclockwise) in which to rotate about the pivot.

   The practical difficulties with the start-end-centre-direction specification of an arc are as follows:
   - When using fixed place representations of coordinates, very small arcs can result in the centre point overflowing the maximum coordinate representation in on or both axes.
   - When representing very small arcs of high included angle, round-off error in the position of the centre makes a drastic difference in the included angle.
   - When using fixed place representations of coordinates, it is hard to accurately represent 45° arcs because, although the start and centre point can be precise, the accuracy of the end point is complicated by an irrational number ($\sqrt{2} = 1.414...$).
2. *Start-End-Angle:* Three items are specified:
   (a) The start point of the arc.
   (b) The end point of the arc.
   (c) The included angle of the arc when moving from start to end. A positive angle is a CW or CCW rotation; negative is the opposite rotation.

The practical difficulties with the start-end-angle specification of an arc are as follows:

- When using fixed place representations of angles and included arc angles are very small, but the distance between start and end is quite large, the round-off error in the angle can be quite significant.
- When using fixed place representations of coordinates, it is hard to accurately represent 45° arcs because, although the start point and angle can be precise, the accuracy of the end point is complicated by an irrational number ($\sqrt{2} = 1.414...$).

The difficulties of the start-end-centre-direction approach make it impossible to represent all arcs using the approach (because of the difficulty with overflow in the centre location). Therefore, internally, pcbnew should represent arcs with the start-end-angle approach: which it already does. However, the resolution of the angle (0.1 degrees of arc) is not sufficient for representing very long, low included angle, arcs.

2.8 Modules

pcbnew historically has duplicated every module on the board. When there are over 500 1uF ceramic chip decoupling capacitors in a board design, this is rather wasteful of resources. Also, the mechanism for placing items within a module results in a significant duplication of code for items such as draw segments (module edges), text and other components of a module. The only significant difference, for example, between draw segments (which are placed on a board) and module edges (which are placed within a module), is the placement reference. The board items are placed with respect to board coordinates and origin, whereas the module items are placed first with respect to the module and then the module is placed with respect to the board. Absolute board coordinates are also maintained in the module item and conversion between module-local coordinates and board-local coordinates is routinely performed. The historical approach presents the following problems:

1. The historical approach is wasteful of memory. Yes, I know that computers have logs of memory nowadays, however, the larger the memory footprint, the slower the application.
2. The historical approach is wastes board file contents by replicating redundant information. Not only is the information redundant but consider hand-editing the footprints in a pcbnew board file (yes, I've done this): changing a footprint for CAPC1005L can require edits in hundreds of locations.
3. Some file formats (such as [GenCAM], [GenX], [258X], [Kehmeier and Makowski, 1998]) factor modules into packages and placement. Generating these file formats would require factoring the old pcbnew modules into placement instances and unique footprints anyways. This was my primary reason for splitting the module into placement and relative position information.

(R) 4 (pcbnew) *With a view toward supporting fabrication panelization, it should be possible for all board items to be placed with respect to a given reference position, side and orientation. Placement information and geometry information will be separated into template and instance.*

The following are implementation considerations for factoring modules into placements and packages:

- When a old version 1 formatted board file is read, a comparison function will be performed on the local coordinates of modules to determine whether the modules are identical.
- Field placement and text characteristics will not be considered when comparing modules.
- Module of the same footprint name that are different from each other will be given a version code. The version code can be displayed in the footprint name by appending a tilde and then the footprint version code.
- When version 2 board files are written, the version code will be appended to the footprint name to form a package name when there is more than one version of the footprint.
- Version 2 board files will also generate a $PACKAGE$ definition that will be placed into the board file. The $PACKAGE$ definition will only contain local information.
- The MODULE class will be modified to contain only placement information (placement of fields and absolute placements of module items).
- The MODULE class will refer to an instance of the PACKAGE class for local footprint definitions read from a library or unique within the board file.

This approach has ramifications on the module editor, particularly when editing a module within a board. The key issue is whether to apply the changed module to all instances of the footprint, or only the one from which the editor was launched. While this might require some internal juggling, the user interface need not change.

3 New Features

General. General features are detailed in *Sec. 3.1(17)*.

Pads. Pads have required a number of feature enhancements including pad removal, complex shapes, additional basic pad shapes, plotting of thermals, plotting of moirés, plotting of drill marks, and solder-mask-defined vs. non-solder-mask-defined. Pad features are detailed in *Sec. 3.2(17)*.

Vias. Vias have required enhancement to through vias, blind vias, buried vias, microvias and the addition of back-drilled vias, depth-control-drilled vias, subcomposite vias, and via-through-pad. Via features are detailed in *Sec. 3.3(26)*.

Tracks. Tracks have required enhancement for etch compensation, teardrops, sub-composite vias, isolation traces, arced track, and push/pull tracks. Track features are detailed in *Sec. 3.4(34)*.

Zones. Zones have required enhancement for impedance cutouts, copper islands, isolation zones, module zones, cross-hatching, thieving, venting, NSMD vs. SMD pads. Zone features are detailed in *Sec. 3.5(40)*.

Modules. Modules have required enhancements for courtyards, rework, wave soldering, component outlines, keep-outs, thermal vias, press-fit fixtures, non-plate-through holes, routed slots and cutouts, and other features. Module features are detailed in *Sec. 3.6(42)*.

Nets. Although historically pcbnew has provided basic net classes, these are inadequate to model high speed buses, stub-terminated logic, impedance controlled lines, back-drilling and other via fabrication to reduce stubs, reference-plane cutouts under pads and contacts, tight and loosely coupled differential pairs, and other net-based requirements. Net features are detailed in *Sec. 3.7(50)*.

Layer. Aside from the specific layer requirements and the addition of new layers, layers have required enhancements for thieving, venting, crosshatch, and other plating and lamination features. Layer features are detailed in *Sec. 3.8(50)*.

Stack-up. The lamination sequence and stack-up of multilayer boards have not previously been modelled in pcbnew. Many advanced features require a model of the stack-up. Therefore a stack-up model was added to pcbnew. Stack-up features are detailed in *Sec. 3.9(54)*.

Panels. In fitting with *design-to-spec* and *plot-and-go* objectives, panel features were required by pcbnew. pcbnew did not historically provide any panel features, and all data was *one-up*. Therefore, panel features were added to pcbnew. Panel features are detailed in *Sec. 3.10(55)*.

3.1 General

General changes demanded for the architecture are as follows:

Defaults. Changes related to the application of defaults and process limits are detailed in *Sec. 3.1.1(17)*.

3.1.1 Defaults

The layout of a board can be a tedious and repetitive process. The approach to the application of process limits and defaults should be designed so as to remove as much repetition and tedium from the process as possible. When a board is near completion, a simple change could require iterating though hundreds or thousands of items. Also, when fabrication capabilities change due to selection of a different fabricator, process limits (minimums and maximums) may change, causing many features on the board to fail fabricator DRC checks. An design pattern for defaults and the application of process limits in general can alleviate most of the tedium and repetition of these kinds of changes.

One of the few places that pcbnew has historically provided a default that would automatically change many items was the default drill and size for vias. However, even this feature is lacking in applying process limits.

A design pattern was devised to provide for two views of the same set of parameters. The two view are:

As designed. The "as designed" view presents the information as it was specified by the designer. That is, if a 2mil pad size is specified, a 2mil pad size will be viewed.

As built. The "as built" view presents the information as it was specified by the designer, but with process limits applied. That is, if a 2mil pad size is specified, and the minimum pad size is 10mil, then a 10mil pad size will be viewed.

3.2 Pads

This subsection addresses the changes and feature enhancements that were made to pads and pad construction, as follows:

Pad Attributes. Pads are the basic way of generating holes (other than vias) in the PCB. Pads are also used for specialized structures in pcbnew. However, many common uses of pads in pcbnew have no direct attribute support. This feature addresses additional attribute support for pads. See *Sec. 3.2.1(18)* for details.

BGA Pads. BGA pads have unique requirements that have not historically been supported by pcbnew. This feature addresses attributes and specifications necessary for full BGA pad support. See *Sec. 3.2.2(18)* for details.

Pad Removal. Most fabricators now want to remove non-functional pads from inner layers, or to removal all non-functional pads, from PTH. Removal of non-functional pads is as applicable to via pad stacks as it is to component pad stacks. Non-functional pads can be removed by fabrication note; however, for plot-and-go service, this feature addresses this removal. See *Sec. 3.2.3(18)* for details.

Complex Shapes. Fabrication data (plots) should flash pads, regardless of their complexity. RS-274X, DPF and all CAM formats have a way of specifying complex pads. This features addresses complex pad shapes. See *Sec. 3.2.4(20)* for details.

More Basic Pad Shapes. Some basic pad shapes that are supported by RS-274X, DPF and CAM formats have not historically been implemented in pcbnew. This feature addresses a fuller set of basic pad shapes. See *Sec. 3.2.5(20)* for details.

Plotting of Thermals. RS-274X, DPF and CAM formats have a way of directly plotting thermals. pcbnew has historically plotted thermals by creating the clearance ring and bridge lines separately. This feature addresses using thermal features of plotting to formulate thermals. See *Sec. 3.2.6(21)* for details.

Plotting of Moirés. See *Sec.* **??***(??)* for details.

Plotting of Drill Marks. Most fabricators do not permit direct drill on ground plane (without thermals) although this is a common way of providing via connections in under dense BGA fields. To avoid tool breakage caused by direct drill on ground plane, a clearance hole in the artwork should be plotted. This feature addresses plotting of drill marks. See *Sec. 3.2.8(21)* for details.

SMD vs. NSMD. Historically pcbnew has not distinguished between solder-mask-defined and non-solder-mask-defined pads. This features addresses the distinction and the post-processing necessary to support both. See *Sec. 3.2.9(22)* for details.

Non-Plate-Through Holes (NPTH). Historically pcbnew (the module editor) has permitted the definition of NPTH; however, proper support for NPTH once specified has little to no support in pcbnew. This feature addresses full support for NPTH and other mechanical or tooling holes. See *Sec. 3.2.10(25)* for details.

Tooling Holes. See *Sec. 3.2.11(25)* for details.

Hole Tolerance. See *Sec. 3.2.12(25)* for details.

Fiducial Marks. See *Sec. 3.2.13(25)* for details.

Contacts. See *Sec. 3.2.14(26)* for details.

3.2.1 Pad Attributes

3.2.2 BGA Pads

BGA pads, especially on high-pitch BGAs, have special requirements that need to be separately addressed.

SMD vs. NSMD. BGA pads can be solder-mask-defined or non-solder-mask-defined. See *Sec. 3.2.9(22)* for more information on the differences and handling of SMD and NSMD pads in general.

SMD BGA pads. The advantages of SMD pads is that the copper area of the land is larger and encompassed on all sides by primary mask. This provides better adhesion of the pad to the laminate, and thus better resists popping the pads of the laminate for large CBGA packages. CBGA packages also have a higher solder volume requirement to form a stronger mechanical joint. The disadvantage of SMD pads is that the copper pad is physically larger, resulting in reduced room for routing traces between lands during BGA breakout. Another disadvantage of SMD pads is that the solder mask registration is not so good, so the resulting pads may be slightly displaced from the necessary location. SMD pads typically withstand the stresses of rework better than NSMD pads.

NSDM BGA pads. The advantages of NSMD pads is that the land is copper defined and therefore can be smaller: permitting additional room for BGA breakout traces and vias without resorting to expensive VTP or VBP constructions. Because the land is copper-defined, the size can be adjusted so that the resulting joint overlaps the edges of the copper pad, providing additional mechanical strength and resisting peeling. A disadvantage of NSMD pads is the larger primary mask opening and the need to properly avoid bridging and wicking with vias under the BGA field. NSMD pads generally do not withstand the stresses of rework as well as SMD pads. One of the ways to mitigate some of the disadvantages of NSMD BGA pads is to add teardrops to the BGA pads to increase the copper area exposed by the primary mask opening without reducing the room for routing of traces. Increasing the exposed copper area, increases the joint area and strength, as well as reducing the stress riser at the junction of the trace and the pad. For more general information on teardrops, see *Sec. 3.4.2(35)*.

There has been much debate about the approach to BGA pads. The current trend is to always defined BGA pads as NSMD pads of a reduced size and to always add tear drops. Nevertheless, pcbnew should support all aspects of both approaches.

BGA pad masks and clearances. Although BGA pads are subject to many of the same requirements that are applied to any SMT pad, there are special requirements for BGA pads. This is because, unlike SMT that has a terminal that directly contacts the copper, solder balls are unique things, and when solder is bridged or wicked from the joint, very poor yields can result. Also, CBGA pads are critical in that the rigidity of the package combined with board bow can cause teeter-tottering of the BGA. Because of these effects, the solder paste mask must be treated specially as well as clearances between the BGA land and other board features.

To address the requirements of BGA pads, the following needed to be done:

- An attribute is necessary for the pad itself to identify it as a BGA pad so that specialized rules and post-processing can be performed on these lands. See *Sec. 3.2.1(18)* for more information on enhanced pad attributes.
- An attribute is necessary for the module to identify the type of BGA module so that specialized post-processing can be performed on the module. Of particular interest is CBGA, CSBGA, FPBGA and PBGA distinctions. See *Sec. 3.6(42)* for more information on necessary module attributes.
- The pitch of the BGA should be identified so that specialized post-processing can be performed (such as the adjustment necessary to solder paste stencil apertures).
- All pads that exist on any outer layer should have an attribute that determines whether the pad is an SMD pad or and NSMD pad. See *Sec. 3.2.9(22)* and *Sec. 3.2.1(18)* for more information on SMD and NSMD pad attributes.
- All pads should have an attribute that determines whether teardrops should be applied to NSMD pads to increase the exposed copper area without otherwise affecting routing densities. See *Sec. 3.4.2(35)* for more information on attributes applicable to teardrops.

Some additional rules for NSMD BGA pads are as follows:

- The opening between the pad and solder mask should have at least 50 microns (2mil) clearance.
- The trace width of traces connecting to the pad should be no more than 60% of the pad diameter. Otherwise, the trace could expose too much copper where the trace is exposed in the mask aperture.
- The trace should have a fillet radius at the point it meets the pad. This prevents the stress riser that would otherwise occur at the pad/trace intersection.

3.2.3 Pad Removal

There are a number of reasons and approaches to removing pads from pad stacks for a number of purposes. The categories of removing pads are as follows:

Removal of unconnected pads. It is possible to remove the land associated with unconnected solder balls or bumps on BGA devices. The removal can increase assembly yields and reduce rework, as well as increasing the signal integrity associated with connections to large BGA, and primarily under FPGA devices. Signal integrity is improved largely by the reduction of parasitic capacitance. Enhancements required to address the removal of unconnected pads are considered in *Sec. 3.2.3(19)*.

Removal of non-functional pads. For boards with low to moderate aspect ratios, it has long been known that removal of non-functional via and TH pads (pads that serve no electrical or solder join purpose) actually increases the reliability of the PTH barrel without significantly affecting yields. From the standpoint of signal integrity, removal of non-functional pads removes parasitic capacitance between the non-function pad and adjacent reference planes. Enhancements required to address the removal of non-functional pads are considered in *Sec. 3.2.3(19)*.

Removal of pads from NPTH. Enhancements required to address the removal of pads from NPTH are considered in *Sec. 3.2.3(20)*.

Removal of Unconnected Pads Some component manufacturers (in particular FPGA manufacturers) recommend that unconnected BGA (or SMT pads simply be removed. Previously in pcbnew the only way to accomplish this was to manually remove pads from each of the footprints on the board using the module editor from within the board. This was a tedious an error prone process. Also, because the corresponding solder balls might not be removed from the BGA, no ghost pad would keep tracks from being routed under the solder ball. On reflow, solder bridging between balls could occur due to the irregularity of the track under the solder mask.

An attribute indicating whether to remove unconnected pads was added to the plot dialogue, board design rules, module definition, and pad definition. The attribute is multivalued and can have one of the following values:

- 0, the default behavior.
- 1, pads are not to be removed.
- 2, pads are to be removed.
- 3, pads are to be masked over. This results in removal of the solder-mask opening for the pad as well as removal of any paste stencil opening. It does not; however, remove the copper pad.

Treatment of these values is different depending on where the attribute occurs. For plots, a setting other than default overrides all other settings on the board. For board design rules, the setting is only consulted when the module and pad settings are default. For module descriptions, the setting is only consulted when pad settings are default. When all settings are default, the default behaviour is to not remove unconnected pads.

A removed pad is never considered removed for DRC/DFX checks. This is so that tracks and vias are not placed under a solder-ball over mask.

Removal of Non-Functional Pads Most fabricators require that non-functional pads be removed. This in fact increases the reliability of the board for all except the highest of aspect ratios. Furthermore, it reduces the amount of copper drilled, particularly on high layer-count boards, which reduces drill bit wear and breakage. High-speed signal integrity is increased by removing non-functional pads from vias, particularly in dense via fields under BGA where higher anti-pad clearances can cause excessive loss of ground plane. This is so much so that it is not even necessary to have an option to leave in non-functional pads.

For example:

> Pads that do not connect to traces are typically not required on inner layers and should be removed to increase drill life (reduce drill cost). This is called non-functional pad removal. Removal of unused pads will not affect the board reliability. Non-functional pad removal should be performed at the design stage or can also be specified on the fabrication drawing or specification. Care should be taken while removing non-functional pads for blind/micro via jobs. Non-functional pads are typically not removed from thick backplane and some high aspect ratio designs.

And:

> Drilled holes are a substantial percentage of the fabrication cost. The product cost will be minimized by removing all unnecessary holes and removing all internal pads that do not connect to a trace (non-functional pads). One of the predominate factors that reduce the quantity of holes drilled per drill bit is the amount of copper that it must drill through. There is a significant amount of reliability data that has been performed over the last 30 years that demonstrate the removal of non-functional pads will not detrimentally affect the plated through hole reliability. There will be exceptions to this rule for very thick backplanes and PCB's with very high aspect ratios.

And:

> One of the 'cleanup' procedures that a fabricator will carry out on a layout is the removal of any non-functional pads (NFPs) or tracks. [There may exist a] 'non-terminated vector', where it is not clear from the design whether this track is a remnant from a previous design iteration, or is an intentional feature.
>
> The fabricator will query every unterminated track, so time can be saved if the designer tells the fabricator which NFPs are intentional...
>
> In order [sic] to provide the best reliability of construction, it is usual for through-holes and vias *not* to have associated pads on internal layers except where they are actually electrically connected to tracks on the layers.
>
> There is, however, a case to be made for the designer setting the CAD program to create such pads automatically, and for the board specification requiring the fabricator to remove unused pads as part of the CAM process. This may appear somewhat bizarre, but having internal pads forces the CAD program to allow sufficient clearance between internal copper tracks and through-holes, preventing problems...

Many fabricators prefer the removal of non-functional pads, where a non-functional pad is a pad on a layer that has no electrical connection. (Note that for a TH, or pin, hole, pads on the outer layers are considered functional: that is, the form a solder joint connection between the pad and the pin.) This is to reduce the amount of copper through which the drill passes and reduce tool wear and breakage. Studies over the last 30 years show that the barrel is not compromised by removal of non-functional pads for typical aspect ratios.

From a signal integrity standpoint, for high-speed board design, removing non-functional pads reduces the parasitic capacitance between ground and reference planes thorugh which the via passes but does not connect, permitting a smaller antipad for the same electrical characteristics. A smaller antipad also provides for fewer traces over reference plane voids for BGA breakouts, also increasing signal integrity without going to the extent of expensive full VTP (Via Through Pad) construction.

Note that from the standpoint or drill wear and breakage there is usually no point to removing non-functional pads from exterior layers or outer sub-laminate layers for buried or blind vias. This is because holes are drilled *before* etching on these layers, so pad removal cannot reduce drill wear or breakage. Also, removing pads from the outer layers of the laminate or sub-laminate can cause a detached via barrel. If the via is then hole-filled, the hole filling process can break away the detached barrel.

For multi-core laser-drilled vias, the outer pad is necessary to reflect laser energy and to keep the laser-drilling process from ablating too much dielectric material on the outer layer. Also, the last inner pad is always necessary as a reflection target. However interior pads can be removed.

Another exception to non-functional pad removal is back-drilled vias. It is advantageous to place a thieving pad of the primary drill diameter minus 4mil on the layers through which the back-drill passes. The purpose of these pads is to stabilize the laminate at this position and make z-axis location of layers more consistent.

Previously in pcbnew there was no way to remove non-functional pads. The only way to remove non-functional pads was to provide fabrication notes requesting remove by the fabricator: (most CAM systems can automatically remove non-functional pads). Some fabricators recommend leaving non-functional pads in the CAD system because many CAD systems are not capable of ghost-pads or track clearance to PTH calculations necessary

for proper DRC/DFX. A simple fabrication note can be provided, for example:

```
1. REMOVE ALL NON-FUNCTIONAL PADS.
```

The problem with fabrication notes is that removal of non-functional pads cannot be selectively applied. Further, parasitic capacitance calculations and high-speed via tuning cannot be correctly calculated by the CAD system. DIY (Do It Yourself) fabrication require manually editing Gerber file (a poor proposition at the best of times).

An atribute was added to address removal of non-functional pads. The attribute indicating whether to remove non-functional pads was added to the plot dialog, board design rules, netclass rules, via definition, module definition and pad definition. The attribute is tri-state and can have one of the following values:

- tri-state, the default behavior.
- on, non-functional pads are to be removed.
- off, non-functional pads are not to be removed.

Treatment of these values is different depending on where the attribute occurs. For plots, a setting other than default overrides all other settings on the board. For board design rules, the setting is only consulted when the module, pad, via and netclass settings are default. For modules descriptions, the setting is only consulted when pad and netclass settings are default. For vias, the setting is only consulted when netclass settings are default. When all settings are default, the default behavior is to not remove non-functional pads.

A removed pad is considered removed for DRC/DFX checks because DRC/DFX checks have been enhanced to include checks for track, pad, via and zone to PTH clearance.

Pad Removal from NPTH Some fabricators may remove pads at the same location as drill hits for NPTH. Some fabricators might not. To accommodate NPTH an attribute was added to pads to define mechanical or NPTH holes. Copper for these holes will always be removed. Previously in pcbnew the only way to define a NPTH was to specify a zero-sized pad, or a pad with the same diameter as the hole (also used by fabricators that remove NPTH pads). To accommodate keep-out clearances for mechanicals such as screw heads or nuts, a pad can be defined on a keep-out layer. One argument for keeping these pads is as a drill mark for manual drill presses. Removing or keeping these pads is therefore only a plot option associated with drill marks.

The fabrication note:

```
2. REMOVE ALL NON-FUNCTIONAL PADS.
```

should be sufficient for removing clearance pads from NTPH as well; however, it performed by fabrication print notes only, a note such as:

```
3. REMOVE ALL CLEARANCE PADS FROM NPTH.
```

should appear on the fabrication print.

3.2.4 Complex Shapes

When providing lands for modules or contacts for edge connectors, it is often necessary to generate a land or contact of a complex shape that is not represented by one of the simple standard pad shapes. Typically, these contacts or lands are made using a composite of standard shapes (rectangles and discs). For example, *Fig. 1(20)* illustrates the composition of a complex shape for edge tips (gold finger) for a board edge connector.

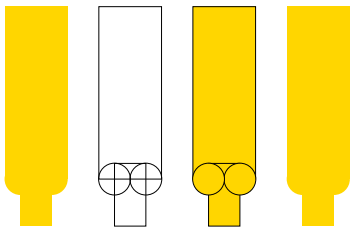

Figure 1: Complex edge tips

Historically in pcbnew pads of a complex shape were defined by defining multiple pads that are connected together to form the necessary shape in just such a manner. There are, however, several deficiencies with pcbnew's historical approach:

- When generating plots, each component of the complex shape were flashed separately. This makes it difficult for CAM software to determine that the individual flashes are intended on forming a single pad of complex shape.
- DRC checks within pcbnew can cause false negatives because the software did not understand that the components of the complex pad were intended on forming a single pad.
- Annular ring checks can cause false negatives due to the size of a single component of the complex shape.

To avoid these difficulties requires a method of determiing the composite shape and collecting its components together. The complex shape can then be communicated as a whole. This can be done in Gerbers by defining a aperture macro; in DPF, a "complex" aperture.

Because it is undesirable to change the module definition file format, there must be some way for pcbnew to determine that the individual components of the complex shape belong together. This can be accomplished by ensuring that all of the components of the complex pad have the same (non-null) pin number. Any pads belonging to a module that meet this criteria will be considered as as a group of pads making up complex shape.

Display: When displaying the pads under pcbnew or the module editor, they will be display in the same manner as if they were independent pads. Therefore, there are no special considerations for displaying the complex shape. There is a question whether to only permit traces to terminate (magenetic pads) to the origin of the complex, or whether to permit traces to terminate to any component of the shape. Currently, traces can terminate to the centre of any component. Note, however, that the origin of any of the components can be moved by applying an offset to the component. It is even possible to make the connection point (origin) of each component the same with the creative use of pad offsets.

Plotting: When plotting, the complex shape will be plotted as a complex shape instead of as the individual components. This is so that downstream CAM systems do not mistake the pads as independent. When plotting Gerbers, the aperture macro will be used to formulate a special aperture. When plotting DPF, the "complex" aperture type will be used.

3.2.5 More Basic Pad Shapes

Historically, pcbnew has only supported a basic list of pad shapes. That included:

- A filled circle (or disc).
- A filled rectangle (or square).

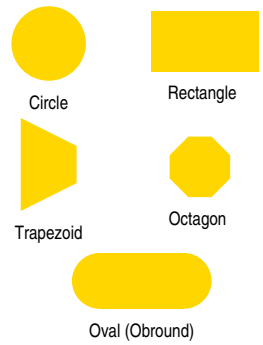

Figure 2: Basic pad shapes

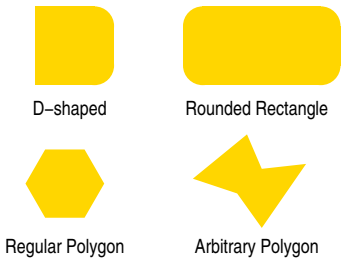

Figure 3: New pad shapes

- An oval (obround).
- A trapezoid.
- A octagon.

These basic shapes are illustrated in *Fig. 2(21)*. Some additional pad shapes that can easily be generated using Gerber aperture macros, or DPF complex apertures, are:

- A D-shaped (bullet shaped) pad.
- A rectangle with rounded corners.
- A regular polygon.
- A arbitrary polygon.

These additional shapes are illustrated in *Fig. 3(21)*.

Display: When displaying these pads they can be displayed by composing their integral shapes.

Plotting: When plotting the new pads under pcbnew, they will be displayed using a Gerber aperture macro (for all but regular polygons, which can use a standard aperture).

3.2.6 Plotting of Thermals

Historically pcbnew plotted thermals by flashing the pad and drawing the bride lines. This violated the rule that apertures and complex pads should always be flashed. pcbnew should use the aperture macro when plotting Gerbers, and should use the "thermal" aperture when plotting DPFs. This has been changed as an feature enhancement to pcbnew.

Note that overlapping thermals are not permitted in accordance with DFM rules. pcbnew has not historically checked whether thermals overlap.[4]

Figure 4: Thermal shape

Figure 5: Moiré shape

3.2.7 Plotting of Moirés

Historically pcbnew also ploted moirés by drawing the circles and constituent lines. This violated the rule that apertures and complex pads should always be flashed. pcbnew should use the aperture macro when plotting Gerbers, and the "target" aperture when plotting DPFs. This has been changed as a feature enhancement.

3.2.8 Plotting of Drill Marks

Historically pcbnew has not plotted drill marks on Gerbers. Strangely, though, it would plot drill marks on HPGL and PS plots. The only plot format that I know of that will not support *scratching*[5] is DXF [DXF].

There are a number of reasons for adding drill marks to plots:

1. Most fabricators discourage direct drill on ground plane. Direct drilling on copper plane results when a via or PTH passes through a ground or reference plane without thermals being added to the pads. Direct drilling on an internal layer copper plane results in increased tool breakage. One of the ways to avoid this is to plot a *scratch* layer of clearance holes on internal copper layers on ground planes. The clearance holes are typically defined to be one-half of the primary drill size.
2. Clearance holes can be defined on all layers to assist with manual drill press operation. Historically, pcbnew has allowed these drill marks to be small, or full sized.

Note: care must be taken when assigning drill mark size. The drill mark must not be larger than the F.H.S. minus the xy-plane registration of the layer and drilling process. Also, the drill mark size must not be less than the minimum feature size for the copper layer on which it appears. This means that some drill marks for small drills on thick copper might have to be removed. Thermals or increasing the drill size are the only options in these cases.

Plotting Drill Marks. There are several ways to add drill marks to Gerber plots:

1. Define one or more %LPC*% information layers and place flashes for each of the drill marks in those layers. The only problem with this approach is that it might confuse CAM

4. Check this out: I seem to remember thermals being calculated in some "optimal" fashion in the case of overlap.

5. Removal of the dark positive portion of the image is called "scratching" from the old days of film emulsions where the only way to make an area in a dark area clear was by scratching off the emulsion.

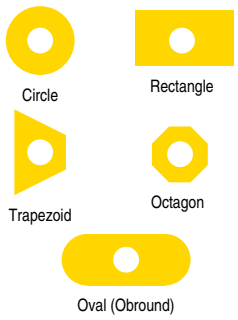

Figure 6: Drill marks

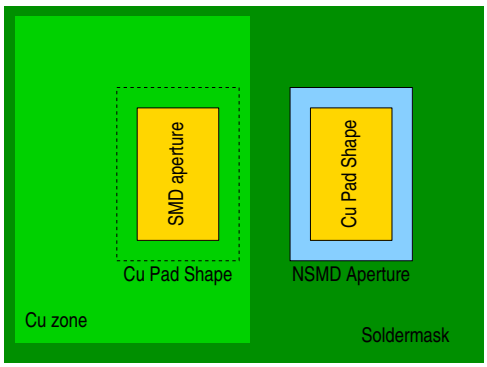

Figure 7: SMD vs. NSMD lands

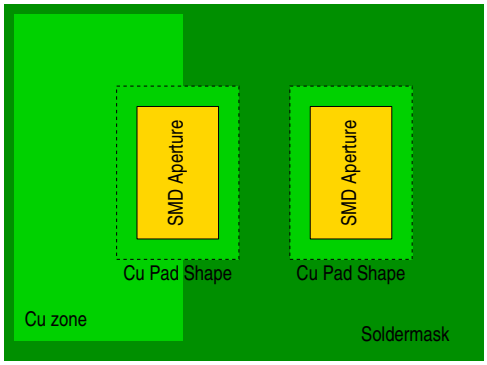

Figure 8: Including SMD pads in a zone

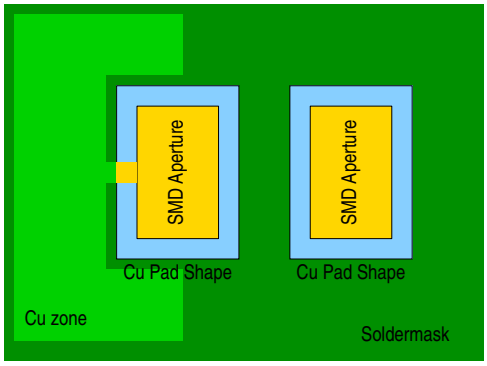

Figure 9: Excluding NSMD pads from a zone

systems. CAM systems normally only expect clear overlays for cutouts on zones.

2. Define an %LPC*% layer for each pad flash and flash a scratch overlay. This too might confuse a CAM system.
3. Define a "hole" (as illustrated in *Fig. 6(22)*) for each aperture so that the hole appears when it is flashed. In this case, the flahses must always be after the drawing of the traces and zone, otherwise the drill mark will not appear. The advantage of this approach is that it compoese the drill mark and the single flash.

3.2.9 SMD vs. NSMD

SMD (Solder Mask Defined) pads is a technique for defining a pad using its solder mask opening rather than its copper outline. SMD is a way of getting a better pad definition when the pad is flooded on one or more sides by a copper plane. If the pad is not defined by its solder mask opening, but by the size of its copper extents, the size of the resulting pad would be increased by the solder mask registration (to avoid solder mask encroaching on the pad), resulting in exposed copper that is too large. This effect is illustrated in *Fig. 7(22)*. To avoid this problem, where the pad should only be defined by solder mask, the opening is not adjusted by the solder mask registration clearance to achieve a pad definition that is the proper dimensions (albeit its true position is affected by the registration of the solder mask). This is illustrated by the dashed line opening on the left of *Fig. 7(22)*.

NSMD (Non-Solder Mask Defined) pads is a technique for defining a pad using its copper outline. It is also, therefore, referred to as a "copper-defined" pad. The size of the pad or contact is defined by copper and the solder mask opening is expanded by the xy-plane registration of the solder mask. This is illustrated to the right in *Fig. 7(22)*.

Surrounding, or flooding, a pad with copper is a technique typically used in power circuitry to reduce the resistance and increase the current handling capacity of the resulting connections, particulary for low ESR (Equivalent Series Resistance) surface mount MLCC capacitors. Component pads can also be partially surrounded or flooded with copper as illustrated in *Fig. 8(22)* and *Fig. 9(22)*.

pcbnew has not historically distinguished between these two types of pad construction and, therefore, can result in solder-mask defined pads that have too large an opening. pcbnew does have a attribute that can be set for zones that permits same-net pads to be defined as "included" or "excluded."

When same-net pads are "included" in a zone, they are flooded with copper as illustrated in *Fig. 7(22)*. However, when same-net pads are "excluded" from a zone, they have copper clearance created between the pad and the zone and are, therefore, copper-defined. This is illustrated in *Fig. 8(22)*.

When "excluded", it is necessary to extend a trace between the pad and the zone to affect a connection. Excluded pads of this sort are illustrated in *Fig. 9(22)*. Therefore, historically in pcbnew, the only way to avoid the problems associated with NSMD pad formation has been to always "exclude" pads from copper zones. This is not sufficient, of course, because it defeats the purpose of flooding the pad with copper in the first place. This is illustrated in *Fig. 9(22)*.

There are several difficulties with not properly accomodating solder-mask defined pads:

- Where the exposed copper is too great, there is a chance

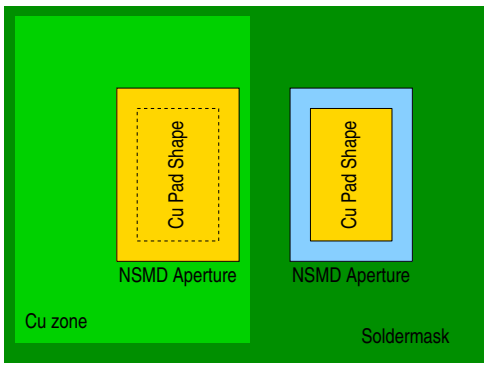

Figure 10: Incorrect flooded NSMD pad

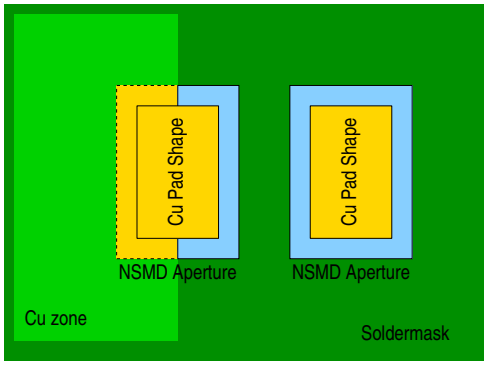

Figure 11: Incorrect partially flooded NSMD pad

that mounted components will float away from their proper position during reflow.

- Because the solderpaste stencil opening will be proper, yet the exposed copper too great, there is a chance that the resulting joint will be starved of solder.
- Because other terminals not flooded by copper will be properly defined by copper, but the erroneous pad is defined by solder mask, the alignment between the two terminals can be incorrect, particularly for small devices where solder mask clearance is a significant percentage of the pad size.

These incorrect applications of an NSMD aperture to a copper flooded pad is illustrated in *Fig. 10(23)*.

Difficulties with the SMD approach. There are a number of difficulties associated with defining lands using soldermask openings:

Registration of the resulting land: Because the land is now defined by the soldermask opening, its positional tolerance is that of the soldermask, not just the tolerance of the copper pattern as was the case for NSMD lands. Soldermask opening positional tolerance is worse than that of copper image and the resulting etch or plating pattern registration. The way of circumventing this difficulty is to apply the soldermask positional tolerance to lands that are soldermask-defined, instead of the copper pattern registration tolerance. This has the ramification, however, that clearances must be derived directly from positional tolerances and not from other rules of thumb or absolute requirements.

Partially flooded lands: Where a land is partially flooded by copper, but not completely, there is a connundrum whether to use solder mask to define the land. Depending on the degree of flooding, the resulting pad could be too large or too small. This is illustrated in *Fig. 11(23)*. An approach to circumventing the difficulties associated with partially flooded lands is to calculate the percentage of flooding and adjust the soldermask opening proportionally. That is, the opening would be somewhere beteen the SMD and NSMD opening, chosen so that the exposed copper pattern has the copper area of the defined pad. This is illustrated by the dotted line on the left of *Fig. 11(23)*.

Communication of design: One technique of providing soldermask design information to fabricators is to plot lands 1:1 (essentially a pad-master) and allow the fabricator to adjust clearances as required. Where lands are plotted 1:1, it can no longer be distinguished from the plot as to whether the design intent was to define the land with soldermask or not. Therefore, it is not possible for the fabricator to properly adjust soldermask openings without manual intervention and possibly consultation with the designer.

Criteria for SMD vs. NSMD. Therare a number of challenges associated with developing a criteria for NSMD vs. SMD determination.

Whether the difference between the land defined as SMD or NSMD can be determined from the size of the land and the clearance of the soldermask opening. Typical 0402 land patters are 15mil × 30mil. Typical soldermask opening clearances are 3mil on each side. This means that the area of the SMD land is 450mil^2. The NSMD openging is 21mil × 36mil, or 756mil, which is 68% larger. A pad of an 1206 package is about 45mil × 90mil, or 4050mil^2. The NSMD opening is 51mil × 96mil, or 4896 square mils, which is 21% larger. As the land gets larger, the difference becomes less significant. At some size, the difference is not important; however, for the most part the difference is significant.

Determining whether solder-mask definition should be used or not, and the degree to which it is used, depends on whether the pad is intended to be flooded with copper, and the extent to which it is flooded. Under pcbnew it can easily be determined whether an SMT or TH pad is in a copper zone, whether it is of the same net as the copper zone, and whether the copper zone is set to "include pads" or "exclude pads." The degree to which a pad is flooded by a zone is a more complex calculation. Under pcbnew it is possible to use the *kbool* engine to determine the amount of flooding of a pad by taking the intersection of the full-clearance soldermask opening minus the pad extents and the copper zone, an calculating the area of the resulting shape. This area can be compared to the full area of the clearance to determine the proportion of flooding. This proportion an then be used to adjust the soldermask opening smaller, resulting in the precise area of exposded copper. This caluclation can be made for all pads in a zone when filling or refilling the zone.

This criteria for automatic SMD calculations could been added as an enhancement to pcbnew, but is is likely too complex for the designer to control and could introduce errors. Perhaps this calculation is better incorporated as a DFM check. For NSMD pads, a calculation can be made to determine the amount of copper that occupies the area between the outline of the copper and the edge of the solder mask aperture, and this area compared to the area of the land alone. When the ratio exceeds a specified percentage, a DFM error and marker could be placed.

I have seen manufacturer DFM rules that state that a trace joining a BGA pad that is wider than 60% of the pad size (diameter) is an issue. This is illustrated in *Fig. 12(24)*.

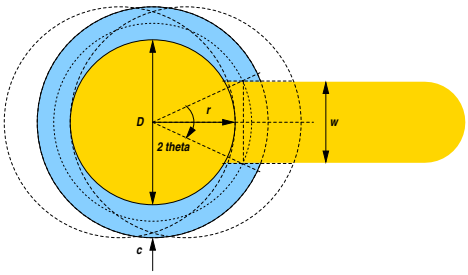

Figure 12: Exposed Copper at Junction

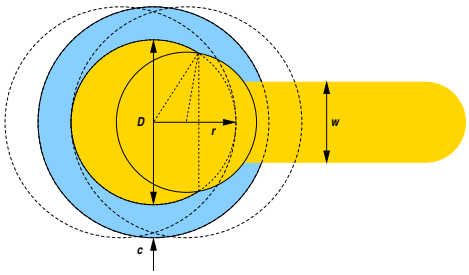

Figure 13: Exposed Copper with Teardrop

The minimum amount of copper area exposed is almost zero. The maximum amount of copper exposed is approximately twice the nominal amount. The nominal amount of copper exposed, where D is the size (diameter) of the pad, and w is the width of the trace, can be calculated approximately as:

$$\theta = arcsin(\frac{w}{D}) \quad (1)$$

$$A \approx \frac{2\theta}{2\pi} \times (\pi(r+c)^2 - \pi r^2) = \theta(2rc + c^2) \quad (2)$$

$$A \approx arcsin(\frac{w}{D}) \times c \times (c + D) \quad (3)$$

also, for rectangular pads,

$$A = c \times w \quad (4)$$

For an example of a $D = 40$mil pad with a $c = 3$mil clearance, a trace of 60% width, $w = 0.6D = 24$mil, $\theta = arcsin(0.6) = 0.6435$ radians, and the approximate area is $A_{nom} = 0.6535((2)(20)(3) + (3)(3)) = 83mil^2$, or $(3)(24) = 72mil^2$. The area of the pad is $A_{pad} = \pi r^2 = 400\pi$, so the ratio of the exposed copper to the pad area is $A_{nom}/A_{pad} = 6.6\%$. The maximum area of exposed copper is approximately twice the nominal, so $A_{max}/A_{pad} \approx 13.2\%$ This is not so bad. But substituting a pad size of $D = 20$mil, $w = 12$mil, $A_{nom} = 44.4mil^2$, $A_{nom}/A_{pad} = 14.1\%$ and $A_{max}/A_{pad} \approx 28.2\%$, which is significant. Reasonably one should expect a copper joint area of $\pm 10\%$. Achieving this for a ratio of 20% would require reducing the copper pad by 10% in area.

For BGA lands, this analysis can become critical. One of the ways of mitigating the variation in exposed copper between the copper-defined pad and the solder mask opening is to provide the same fillet or teardrop at the connection site. A typical teardrop to use is a circular sub-land sized 2mil smaller than the land and offset 3mil from the centre of the land. This is illustrated in *Fig. 13(24)*. The application of a teardrop to the BGA land also serves to remove the stress riser at the junction between the trace and the land. Calculating the additional exposed copper in the sub-land is an exercise in circle geometry and can be expressed as the difference between the circle segment areas:

$$A = \frac{r'^2}{2}(2\theta' - sin(2\theta')) - \frac{r^2}{2}(2\theta - sin(2\theta)), \text{where} \quad (5)$$

$$r = \text{the radius of the land,} \quad (6)$$

$$r' = r - 1\text{mil, the radius of the sub-land,} \quad (7)$$

$$2\theta = arcsin(\frac{h}{r}), \text{the central angle of the chord on the land, and,} \quad (8)$$

$$2\theta' = arcsin(\frac{h}{r'}), \text{the central angle of the chord on the sub-land.} \quad (9)$$

DFX Rules for NSMD Pads. For SMT pads, it is advantageous to develop a rule for when the exposed copper in the primary mask clearance of the NSMD pad significantly affects the copper area of the pad. As a basis, we use 10% nominal as a significant value (because this corresponds to -0 or +20%). The rule is applied to NSMD pads that are not defined for teardrops (such as undersized BGA pads with defined teardrops). Whenever an NSMD pad is partially or completely included in a zone, the entire solder mask opening is compared to the copper pad. The following steps can be performed:

1. Calculate the area of the pad. For a circle, this is $A_{pad} = \pi D^2$, for a rectangle, $A_{pad} = h \times w$.

2. Calculate the mask opening area. For a circle this is $A_{mask} = \pi(D + 2c)^2$, and for a rectangle, $A_{mask} = (h + 2c) \times (w + 2c)$.

3. Calculate the maximum ratio of the exposed copper to pad copper, $\alpha_{max} = \frac{A_{mask}}{A_{pad}} - 1$.

4. When $\alpha_{max} \leq 0.1$, then precise calculations do not need to be performed. When $\alpha_{max} \geq 0.1$, precise calculations need to be performed.

5. The precise calculation for a circle with an adjoining trace is:

$$\alpha = \frac{A_{exp}}{A_{pad}} = \frac{arcsin(\frac{w}{D})c(c + D)}{\pi D^2} \quad (10)$$

and for a rectangle is:

$$\alpha = \frac{A_{exp}}{A_{pad}} = \frac{c \times w}{H \times W} \quad (11)$$

When $\alpha \geq 0.1$, then a DRC error should be raised, or the pad should be set for SMD.

For example, consider a 45mil pad with a 24mil trace adjoining it with a 3mil mask clearance.

$$\alpha_{max} = \frac{(D + 2c)^2}{D^2} - 1 = \frac{51^2}{45^2} - 1 = 28.4\% \geq 10\% \quad (12)$$

so it is necessary to consider exact calculations. Calculating exactly:

$$\alpha = \frac{arcsin(\frac{w}{D})c(c + D)}{\pi D^2} = 11.251\% \geq 10\% \quad (13)$$

and a DFX rule violation needs to be reported.

For example, consider a 12mil trace adjoining a 12mil × 40mil rectangular land with a 3mil mask clearance.

$$\alpha_{max} = \frac{(H + 2c) \times (W + 2c)}{H \times W} - 1 = 91.25\% \geq 10\% \tag{14}$$

So, exact calculations need to be performed.

$$\alpha_{max} = \frac{c \times w}{H \times W} = 7.5\% \leq 10\% \tag{15}$$

and no DFX rule violation need be reported.

Data Requirements for Library Modules. The following are the data requirements for library modules:

1. Pads defined on copper layers should always be defined as the NSMD (that is, copper-defined) shape. It should be the responsibility of the board editor and post-processing to expand the copper shape as necessary where SMD pads are required. It should also be the responsibility of the board editor and post-processing to shrink the copper pad and apply teardrops where required.
2. Solder mask apertures should always be defined as the same size as the SMD (that is, copper-defined) pad. It should be the responsibility of the board editor and post-processing to expand the solder mask aperture as necessary where NSMD pads are required.
3. For SMT pads, it should be possible to specify either SMD or NSMD for the pad at the pad level, the module level and the board level.
4. BGA pads should be separately identified as such.

3.2.10 Non-Plate-Through-Holes (NPTH)

3.2.11 Tooling Holes

3.2.12 Hole Tolerance

See also *Sec. 3.2.3(20)* for more information on pad removal from NTPH.

3.2.13 Fiducial Marks

Fiducial Marks, sometimes simply called *fiducials*, are marks placed on a production panel or board to assist vision or machine vision systems in establishing a reference plane for the panel, board, or locality on a board. Another purpose for fiducial marks is to establish whether a panel or board within a panel is good or bad (has passed or failed quality control tests).
Fiducial marks can serve these several purposes as follows:

Panel Fiducial Marks: (Also called *Global Fiducial Marks*, where the term "global" refers to the production panel.) Panel fiducial marks intended to provide positioning information for layers and the panel in general are typically placed in three or four corners of a production panel. They are placed in the margins of the production panel outside of any board outline, and should not normally occupy shared assembly rails. This is illustrated by the black fiducial marks in *Fig. 14(25)*.

Bad Panel Fiducial Marks: Bad-panel fiducial marks intended to distinguish between good and bad panels in a production line. A panel positioning fiducial could provide a dual purpose as a bad-panel fiducial mark. This is illustrated by the red fiducial marks in *Fig. 14(25)*.

Board Fiducial Marks: (Also called *Local Fiducial Marks* from the perspective of the production panel, or *Global Fiducial Marks* from the perspective of the board.) Board fiducial marks intended to provide pick and place and other vision processes with primary reference axes for a board. They are, in general, placed in three or four corners of a board within a production panel. They are placed within the board outline and must, therefore, share space with other board features. This is illustrated by the blue fiducial marks in *Fig. 14(25)*.

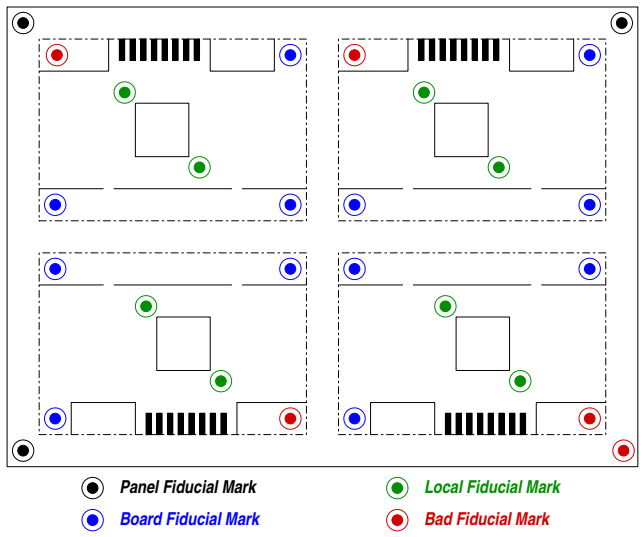

Figure 14: Fiducial Mark Placement

Bad Board Fiducial Marks: Bad-board fiducial marks intended to distinguish between good and bad boards within a production panel. A board positioning fiducial could provide a dual purpose as a bad-board fiducial mark. This is illustrated by the red fiducial marks in *Fig. 14(25)*.

Local Fiducial Marks: Local fiducial marks intended to provide pick and place and other vision processes with local reference axes for specific high-pitch components. They are typically placed in two diagonally opposite corners of the critical component. Components that normally need local fiducial marks include large or high-pitch BGA devices, large or high-pitch QFP/QFN devices, large or high-pitch SOP devices. This is illustrated by the green fiducial marks in *Fig. 14(25)*.

The placement of the various fiducial marks are illustrated in *Fig. 14(25)*.

Almost all assemblers use the standard fiducial marks defined by SMEMA [SMEMA 3.1]. A standard fiducial mark is illustrated in *Fig. 15(26)*. Fiducials require the following:

1. A circular copper pad $R = 1$mm in diameter.
2. A circular clearance area of bare laminate $4R = 2$mm in diameter concentric with the pad.
3. A circular solder mask aperture, $4R + 2c = 2.15$mm in diameter.
4. A circular anti-pad $4R + 2c + 2w = 2.3$mm in diameter.

It is necessary that there be a clearance, c, between the clearance area and the anti-pad so that there will be no exposed copper from an adjacent trace, pad, or zone.

The pcbnew module editor has historically provided much of the functionality necessary to define fiducial marks in this manner. A pad with a 1mm diameter can be defined. A "pad" on the soldermask layer of 2mm in diameter can be defined to clear the solder mask back, or, alternately, the 1mm copper pad can be

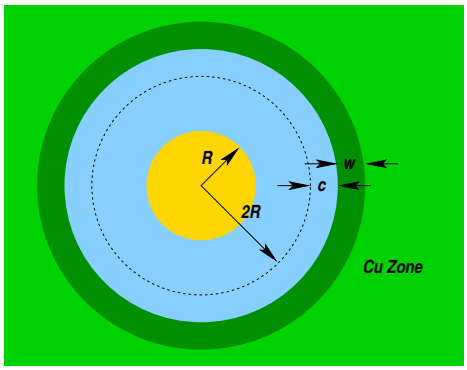

Figure 15: Fiducial Construction

given a solder mask clearance of 1mm. The 1mm diameter copper pad can be given a (copper) clearance of 1.075mm to provide the antipad. The antipad size will not, in general, be adjusted properly for solder mask registration, but that's OK. The copper pad could be treated as SMT; thus, the fiducial would be included in the pick-and-place files. The solder-paste layer could be unchecked for the copper pad, making sure that no solder-paste opening would be provided. The footprint could even be given the name "FIDUCIAL" so that its purpose is apparent in pick-and-place files.

The capabilities for making and placing fiducial marks are somewhat lacking, however, in the following respects:

1. It should be possible to define the difference between a fiducial mark for placement, and a fiducial mark that is used as a bad-panel or bad-board mark.
2. The component or pad should have a distinct type identifying it as a fiducial mark so that it does not appear in BOM lists. Historically, under pcbnew the choice was an SMD [sic], requiring pick and place, or as a virtual, not requiring pick and place. What is needed is a fiducial type so that it can be excluded from BOM like a "virtual", yet included in the pick and place data like a "SMD."
3. Although one could describe the solder mask aperture of a sufficient size using a non-copper layer pad, it was not possible to properly identify the clearance zone. The clearance zone must be devoid of *all* copper features. This is because a copper feature, even a trace on an inner layer, can cause UV shadows and reflections to be cast, or can cause an irregularity in the planarity of the fiducial mark, confusing the precise location of the fiducial to vision systems. This is of particular importance for the board and local fiducial marks that must be used by automated pick and place machine vision systems. Convincing pcbnew to clear out copper planes to a sufficient anti-pad size, as well as teaching it that traces and other copper features must clear this zone is difficult using the historical approach.
4. Some manufacturers recommend placing copper planes underneath fiducials on an inner layers beneath the fiducial to prevent solder mask ghost images (UV light transmission through the laminate). Copper on inner layers can improve the local planarity of the laminate at the location of the fiducial. Convincing pcbnew to provide clearances on inner layers and add copper to inner layers was not possible with the historical approach.

(R) 5 (pcbnew) *A specialized module type will be assigned to identify a fiducial mark, so that the appropriate rules can be applied.*

3.2.14 Contacts

Contacts are board external copper layer features that are intended on forming a mechanical contact with some part of the final product. This includes edge tips (gold finger) for edge contacts on add-in cards as well as carbon contacts for keyboard integration. Because of their specialized purpose, contacts typically require a separate finish from the primary board finish. In the case of gold fingers this is usually nickel/hard gold; for keyboard contacts, carbon ink.

Historically pcbnew has grouped contacts in with other "virtual" board components, where a "virtual" component is one that is drawn on the board using some special shape. Edge tips have been successfully modelled as pads in a component with a "virtual" attribute. This excludes the module from both the BOM and the pick-and-place data. To define tie-bars for electroplated nickel/hard gold edge fingers in a vertical plating line, it was possible to use draw segments on copper layers. Unfortunately, these approaches have the following limitations:

1. The contacts are simply added to external copper artwork. There should be a way to provide a separate artwork detail for contacts. This can be accomplished using the new *contacts* technical layer. See *Sec. 5.9(79)* for more detail on the *contacts layer*.
2. There was no way to distinguish the surface finish of the contacts. This has been accomplished by adding a *material* to the pad definition.
3. There was no way to define removable (Corfin, Kapton tape) or peel-off mask to protect the contact surfaces during the application of primary mask or during wave soldering. This has been accomplished using the new *peelable mask* technical layer. See *Sec. 5.10(81)* for more detail on the *peel mask layer*.
4. There was no way to define *gold masks* for selective electroplated nickel/gold contacts when the contacts where not suitable for a gold finger electroplating line. This has been accomplished using the new *finish* technical layer. See *Sec. 5.11(82)* for more detail on the *finish layer*.
5. DFM rules for gold fingers were not incorporated, resulting in the possibility of defining contacts that can not be manufactured. A set of DFM rule checks were incorporated into the new *contacts* technical layer to accommodate this requirement. See *Sec. 5.9(79)* for more detail on the *contacts layer*.
6. There was no way to provide a fabrication print detail sheet for gold finger and carbon contacts. This has been accomplished using the new *contacts* and *finish* technical layers that may be plotted on fabrication masters or detail. See *Sec. 5.24(116)* for more detail on the *drawing layer*.

Accommodating these requirements demands a new attribute so that pads in a "virtual" module can be identified as being contacts, and the material that is used to surface the contact must also be identified. The "contact" attribute is simply another data point in the pad type. The material is a new attribute added to pads indicating the surface or contact finish associated with the pad.

3.3 Vias

This subsection addresses the changes and feature enhancements that were made to via construction. The primary via constructions are illustrated in *Fig. 16(27)*.

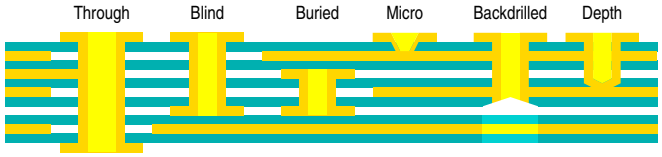

Figure 16: Via constructions

Through vias. Through vias are regular vias that are constructed with PTH that traverse through the entire board. Through vias have historically been represented by pcbnew. Enhancements required for through vias are considered in *Sec. 3.3.1(27)*.

Blind vias. Blind vias are vias that traverse from an exterior layer to an interior layer of the board. Blind and buried vias have historically been represented by pcbnew. Enhancements required for blind vias are considered in *Sec. 3.3.2(28)*.

Buried vias. Buried vias are vias that traverse from on interior layer of the board to another. Blind and buried vias have historically been represented by pcbnew. Enhancements required for buried vias are considered in *Sec. 3.3.3(28)*.

Microvias. Microvias are small vias, normally laser-drilled, but can be chemically etched, that traverse through an external layer. Microvias have historically been represented by pcbnew. Enhancements required for microvias are considered in *Sec. 3.3.4(29)*.

Back-drilled vias. Back-drilled vias start out life as a regular through-via, but which have a secondary drilling process applied to remove most of all of the stub remaining on the via. Back-drilled vias have not historically been represented by pcbnew. Enhancements required for back-drilled vias are considered in *Sec. 3.3.5(30)*.

Depth-control-drilled vias. Depth-control-drilled vias are vias that are drilled from one side of a the laminate or sub-laminate to form a blind or buried via. Depth-control-drilled vias have not historically been represented by pcbnew. Enhancements required for depth-control-drilled vias are considered in *Sec. 3.3.6(31)*.

Sub-composite vias. Sub-composite vias are vias that transition layers in a step. The sub-composite is a feature that extends between a microvia and a buried via. Sub-composite vias have not historically been represented by pcbnew. Enhancements required for sub-composite vias are considered in *Sec. 3.3.7(32)*.

Via-through-pad vias. Via-through-pad vias are constructed using one of the other methods, but which pass through a pad on an exterior layer. Via-through-pad vias have not historically been represented by pcbnew. Enhancements required for via-through-pad vias are considered in *Sec. 3.3.8(33)*.

Thermal vias. Thermal vias are a type of via whose primary purpose is to provide thermal conduction through the board in the z-axis. The connection between thermal vias and other copper areas are for the primary purpose of thermal conduction rather than electrical conduction. Enhancements required for thermal vias are considered in *Sec. 3.3.9(33)*.

Ground plane stitching. Ground or voltage plane stitching is a technique used to control the flow of currents in and out of a copper plane, or between copper planes. Often they do not require (not should connect) a trace at either end of the via. Enhancements required for ground plane stitching is considered in *Sec. 3.3.10(33)*.

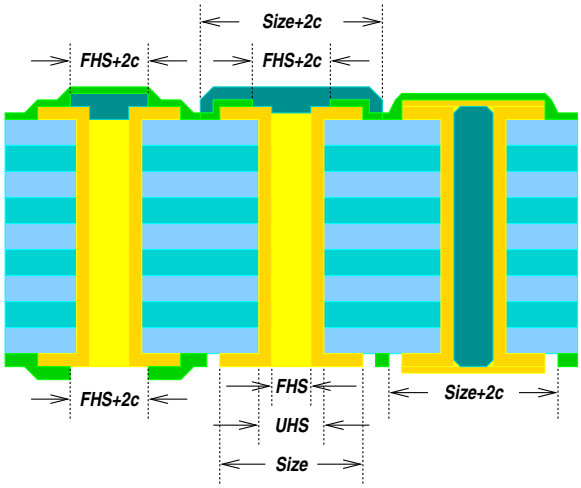

Figure 17: Through Via Construction

Isolation vias. Isolation vias are neutral copper vias intended on providing shielding in the xy-plane on multiple layers. They can be used to encase sensitive signal traces. Enhancements required for isolation vias are considered in *Sec. 3.3.11(33)*.

3.3.1 Through Vias

A through via is illustrated in *Fig. 17(27)*. The historical treatment of through-vias under pcbnew is well established. Little change needs to occur with the exception of thermal vias, ground plane stitching and isolation vias that are addressed separately in *Sec. 3.3.9(33)*, *Sec. 3.3.10(33)* and *Sec. 3.3.11(33)*. Through via basic construction is well handled in pcbnew with the exception of the application of "as-designed" and "as-built" parameter treatment (see *Sec. 3.1.1(17)*). Where pcbnew has been lacking is in the identification of through-vias, the definition of the via pad-stack, and the application of a usable set of DFM rules for via construction.

Identification of through vias.

DFM rules for through vias. The fabrication rules and parameters for through-vias, illustrated in *Fig. 17(27)*, are as follows:

Minimum drill size, UFS_{min}: Because machine drilling of hole sizes larger than about 6mil is the least expensive drilling process with the least capital expenditure, it is the most common drilling process for through vias and predominates current PCB fabrication. Machine drilling has an absolute minimum drill size of about 6mil. With spindle wobble and wander, most fabricators limit the minimum drill size to 8 or 10mil. The minimum drill size will be specified by fabricators.

Note that, although fabricators often require their customers to provide designs specifying only the Finished Hole Size (F.H.S.) for all NPTH, they seem to break their own rules by specifying the minimum *drill* size.

Maximum aspect ratio, α_{max}: The plating process, above a minimum drill size, is limited by aspect ratio. There is a maximum aspect ratio that will result in adequate yields. Above this maximum aspect ratio, the middle of the barrel

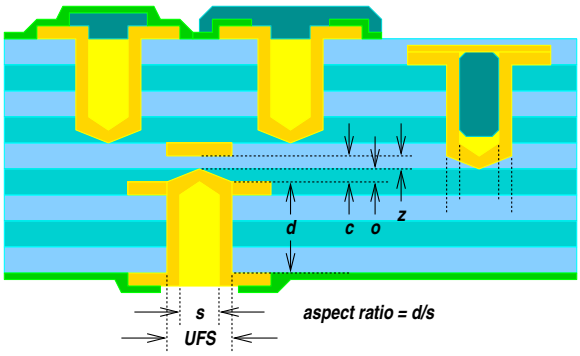

Figure 18: Blind Via Construction

of the hole will not plate adequately. The maximum aspect ratio for PTH is specified by the manufacturer. Aspect ratios of 6 or 7 are very easy to fabricate. Maximum aspect ratios are in the range from 10 to 12. Some fabricators might be capable of aspect ratios as high as 20.

Minimum annulus, A_{min}:

In addition to the above, there are the normal clearances for the fabrication of PTH:

- PTH to feature clearances.
- PTH to edge or NPTH clearance.
- PTH to copper clearances on inner layers.
- Hole to hole clearances for minimum laminate web.
- Via to contact clearance.
- Via to pad clearance (for blind and through-vias).

pcbnew like, many early CAD systems, only historically addressed copper to copper clearance values.

Fabrication data requirements for through vias.

3.3.2 Blind Vias

A blind via is illustrated in *Fig. 18(28)*. Blind vias are vias that traverse from an exterior layer to an interior layer of the board. Blind and buried vias have historically been represented by pcbnew, however, the fabrication outputs for these vias has been lacking.

Blind vias can be constructed using a number of methods; some more expensive than others. The term "blind" relates to the final appearance of the via and does not directly dictate a specific fabrication process. A *Laser-Blind Via* is a blind via that is fabricated using laser drilling. See *Sec. 3.3.4(29)* for more information on laser-blind microvias. However, when the laser-blind via is $\geq$ 6mil, is no longer qualifies as a microvia and is a regular laser-drilled blind via. A *Machine-Blind Via* is a blind via that is fabricated using machine drilling. The minimum size of a machine-drilled via is normally 6mil. This section discusses blind vias that are not microvias. This section discusses only blind vias that are fabricated by through-drilling a sub-laminate. For blind vias are fabricated by depth-control-drilling, see *Sec. 3.3.6(31)*.

Blind vias can be constructed by drilling and plating through a sub-laminate in the buildup. The distinction between blind and buried vias relates the the final location of the via. Blind vias appear on an exterior layer. Blind vias may also be constructed using depth-control-drilling of the laminate or sub-laminate. Depth control machine-drilling or laser-drilling the laminate always results in a blind via.

Which sub-laminates are drilled and whether blind vias are constructed precisely to the functional pair or not has a strong effect on the cost of fabrication. It is possible to reduce the number of sub-laminates and the number of depths achieved by through-drilling or depth-control-drilling by leaving a stub on the top or bottom (or both ends) of the via. The largest acceptable stub is a matter of signal integrity. More stubs is less expensive; fewer, more expensive. Because of the effect on fabrication cost, the designer must have good control over the number and kind of vias fabricated. It is simply not possible to achieve the data requirements and the control necessary without modelling the stack-up of the multilayer board.

Identification of blind vias.

DFM rules for blind vias. The fabrication rules for blind vias are no different than through vias or depth-control-drilled vias, except that the laminate drilled is always the laminate (depth-control-drilled) or sub-laminate in the build-up. Blind vias are also commonly plated closed, filled with a non-conductive material and plated over, or filled with a conductive material. For more information on hole-filling, see *Sec. 5.5(66)*.

Fabrication data requirements for blind vias. The data requirements for these vias depends only partly on the method of construction.

When constructed using through-drilling of a sub-laminate in the build-up, the sub-laminate that is drilled must be identified. Tis is normally done for NC drill files by naming and labelling the file with the first and last layer through which the bit passes. The first and last layer will also correspond to the outer layers of the sub-laminate being through-drilled and plated. For CAM formats, the layers through which the drill passes is identified by the layer-set of the sub-laminate. Most CAM formats have a way of describing the stack-up and build-up for lamination.

When constructed using depth-control-drilling of the laminate or sub-laminate, the sub-laminate must be identified. The depth of the drilling operation must also be identified. This is normally done for NC drill files by naming an labelling the file with the first and last layer through which the bit passes, as well as an indication that the file is a depth-control-drilling file (e.g. by placing DCD in the name). Only one of these layers will correpond to the outer layer of the laminate or sub-laminate (the other will be interior). Therefore, the file must also identify the sub-laminate by layer-pair. For CAM formats, the sub-laminate being drilled is identified by the layer-set of the sub-laminate; however, most CAM formats are unable to describe a depth-control-drilled hole. See *Sec. 3.3.6(31)* for more information on blind vias created using depth-control-drilling.

Historically, pcbnew generated NC drill files for each an every layer pair through which the drill passed. It did not identify the first and last layer for any hole. It did not identify the sub-laminate being drilled. This is inadequate data for fabrication and would have to be manipulated by off-line tools.

3.3.3 Buried Vias

A buried via is illustrated in *Fig. 19(29)*. Buried vias are vias that traverse from on interior layer of the board to another. Blind and buried vias have historically been represented by pcbnew, however, the fabrication outputs for these vias has been lacking.

Buried vias can be constructed using number of methods; some more expensive than others. The term "buried" relates to the final appearance of the via and does not directly dictate a specific fabrication process. The distinction between blind or buried in this case relates to the final location of the via. Blind vias appear on an exterior layer; buried vias do not. Buried vias may be constructed either by through-drilling or depth-control drilling a sub-laminate in the build-up. A through-drilled buried via is a buried via that is fabricated by through-drilling a sub-laminate

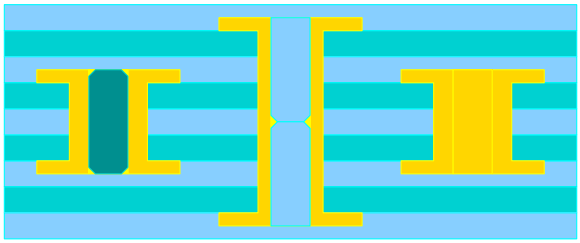

Figure 19: Buried Via Construction

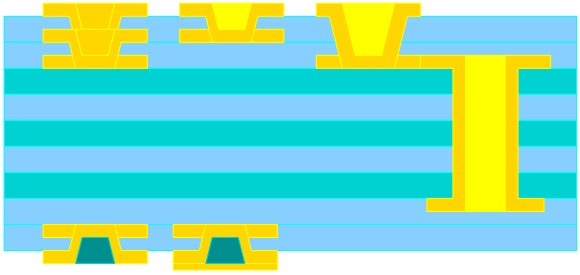

Figure 20: Microvia Construction

as part of the build-up.[6] This section only discusses through-drilled buried vias. A depth-control-drilled buried via is a buried via that is fabricated by depth-control-drilling a sub-laminate as part of the build-up.[7] For more information on depth-control-drilled buried vias, see *Sec. 3.3.6(31)*.

Which sub-laminates are drilled and whether blind and buried vias are constructed precisely to the functional pair or not has a strong effect on the cost of fabrication. It is possible to reduce the number of sub-laminates and the number of depths achieved by through-drilling or depth-controlled-drilling by leaving a stub on the top and/or bottom of the via. The largest acceptable stub is a matter of signal integrity. More stubs is less expensive; fewer, more expensive. Because of the effect on fabrication cost, the designer must have good control of the number and kind of vias fabricated. It is simply not possible to achieve the data requirements and the control necessary without modelling the stack-up of the multilayer board.

DFM rules for buried vias. The fabrication rules for buried vias are no different than through vias or depth-control-drilled vias, except that the laminate drilled is always a sub-laminate in the build-up. Buried vias are also commonly filled with resin from the lamination process, non-conductive filling, or conductive filling. For more information on hole-filling, see *Sec. 5.5(66)*.

Fabrication data requirements for blind and buried vias. The data requirements for these vias depends only partly on the method of construction.

When constructed using through-drilling of a sub-laminate in the buildup, the sub-laminate that is drilled must be identified. This is normally done for NC drill files by naming and labelling the file with the first and last layer through which the bit passes. The first and last layer will also correspond to the outer layers of the sublaminte being through-drilled and plated. For CAM formats, the layers through which the drill passes is identified by the layer-set of the sub-laminate. Most CAM formats have a way of describing the stack-up and buildup for lamination.

When constructed using depth-controlled-drilling of the sub-laminate, the sub-laminate must be identified. The depth of the drilling operation must also be identified. This is normally done for NC drill files by naming and labelling the file with the first and last layer through which the bit passes. Only one of these layers will correspond to the outer layer of the sub-laminate (the other will be interior). Therefore, the file must also identify the sub-laminate by layer-pair. For CAM formats, the sub-laminate being drilled is identified by the layer-set of the sub-laminate; however, most CAM formats are unable to describe a depth-control-drilled hole.

Historically, pcbnew generated NC drill files for each an every layer pair through which the drill passed. It did not identify the first and last layer for any hole. It did not identify the sub-laminate being drilled. This is inadequate data for fabrication and would have to be manipulated by off-line tools.

3.3.4 Microvias

A microvia is illustrated in *Fig. 20(29)*. Microvias are small vias, normally laser-drilled, but can be chemically etched, or machine drilled, that traverse through one or two HDI layers. Microvias have historically been represented by pcbnew, but their representation has been lacking in a number of respects:

1. It only allowed microvias to extend from an exterior layer through one layer. That is, only 1[C]1 HDI constructions were modelled.
2. NC drill file information was provided for each HDI layer, which was acceptable, but again only modelled 1[C]1 HDI constructions.
3. Microvias were not permitted to stack on top of other microvias, nor on top of other vias.
4. Sub-composite vias (steps from microvia to buried via) were not supported.
5. Filling and plating over of microvias was not supported.

Correcting these deficiencies was not possible without providing a model for the buildup of the laminate. Enhancements were made to permite defining dielectric layers as HDI layers. The definition of microvias was extended to permit the via to extend across multiple HDI layers and stack on top of other microvias or regular vias. Sub-composite via constructions were also defined. This permits n[C]n HDI constructions.

Identification of microvias. Microvias were historically simply identified by type in pcbnew. A check was performed before creation to ensure that the layer-pair for the via was appropriate for microvias. Once created, a microvia could not be transformed into another type of via without deleting the microvia and adding the intended type of via. This is inadequate for manipulating the board design under cost considerations. Enhancements were provided to permit the determination of the construction (fabrication) of a via to proceed automatically.

Blind microvias that, indeed, extend from an outer layer of the board through to the next layer can be fabricated either as a laser-drilled microvia, or as a mechanically drilled via (provided that the size of the microvia is within the limitations of mechanical drilling). A right-click option was added to permit transforming a via from one type to another, based solely on its layer pair, and, in the case of stacked vias, the layer pair of the stack. Global operations for convering vias are also provided. This permits the vias on a fully populated board to be transformed and the cost factor of the board adjusted accordingly.[8]

Fabrication data requirements for microvias. Data for microvias consists of NC drill files that identify the layer-pair (and thus the HDI dielectric) through which the microvia passes. For

6. Through drilling the laminate creates a through-via.
7. Depth-control-drilling the laminate creates a blind via.
8. This is true for all vias, not just microvias.

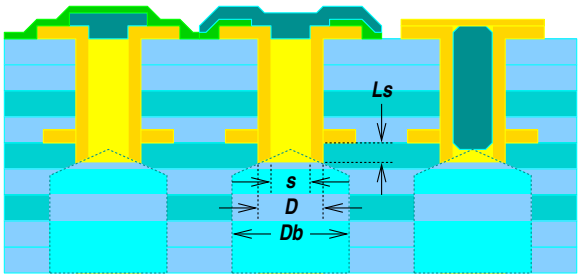

Figure 21: Back-Drilled Via Construction

CAM formats, the layer-pair identifies the sub-laminate through which the laser-drilling or checmical etching process is perfomed. From the NC drill file information, the fabricator can normally generate the aperture mask necessary for removing copper prior to a low-power CO2 laser drilling process; however, the capability for pcbnew to also generate this etching mask was also included.

Note that there are additional requirements when microvias are used for via-through-pad constructions. The requirements for VTP are discussed in *Sec. 3.3.8(33)*.

3.3.5 Back-Drilled Vias

A back-drilled via is illustrated in *Fig. 21(30)*. Back-drilled vias are a form of through-via that have a stub between the last functional layer and the board surface drilled out in a secondary drilling operation performed after plating (but before primary mask and finish). The purpose of back-drilled vias is to remove all or at least a portion of the stub that would affect the electrical performance of high-speed connections.

> Backdrilling removes the unused portion of a PTH or subcomposite via barrel. It can be done from either side of a PWB or subcomposite. The typical backdrill diameter is 0.008" larger than the drill bit diameter used to create the original hole. The backdrill diameter needs to be considered when designing the PWB. Backdrill depth control capability is ±0.004" although a larger tolerance should be specified if allowed by the design.
>
> Backdrilled hole locations are typically identified by using a slightly larger PTH diameter than non backdrilled PTHs of the same size. For instance, if some 0.012" finished PTHs are to be backdrilled, they can be identified by making them 0.01201" PTHs. If another group of 0.012" PTHs are to be backdrilled to a different depth or from a different side, they can be identified as 0.01202" PTHs and so on. The details of each group of backdrilled PTHs, including the identification of the starting side and the do not pierce layer should be defined within the FAB drawing. Some physical design systems include backdrill definition capability. [EI-2010]

Back-drilled vias are typically performed on high-aspect-ratio backplanes where the additional cost of back-drilling is more than compensated for by the signal integrity gains of removing overly-long stubs. Nevertheless, back-drilling is also a viable approach for high-speed boards for addin cards with lower aspect ratios as an lower cost alternative to sub-laminate construction of blind and buried vias.

It is typical to not mix blind and buried vias with back-drilling. Back-drilling is normally reserved for non-HDI constructions where otherwise the entire board would be made up of PTH vias passing through the entire board.

Back-drilling has a number of capability attributes that must be satisfied for the back-drilled via to be fabricated. Pertinent attributes are the amount of oversize, the minimum remaining stub, clearance to the back-drilled hole, registration, etc.

The cost of back-drilling increases with the number of depths back-drilled. This is because some setup is required for each depth. Therefore, there must be some way to allow the designer to control the number of depths and to automate the selection of back-drilling depths for individual vias.

The cost of back-drilling increases with the number of vias back-drilled. There is simply no point to back-drilling a via in a low-speed net, or, for example, in a power or ground net. Therefore, there needs to be a mechanism for specifying which vias need to be back-drilled. It is possible to simply right-click on the via an set whether the particular via is back-drilled or not; however, that is a tedious manual process. The net info data has been enhanced to include the frequency of the net, and the acceptable stub length. This permits the back-drilling process to be automatically selected for through-vias that have stubs that are too long. A set of back-drill depths may also be defined by the designer and the program can select the available back-drill depth that best meets the minimum stub length requirement.

To accommodate this control over back-drilled vias, the stack-up model has been provided to specify for each sub-laminate sequence whether back-drilling is permitted, the process side(s) on which it is permitted, and the permitted drilling depths (in terms of layers) for each process side.

Where a non-back-drilled via is permitted (e.g. signal integrity stub length permits), a non-back-drilled via will always be preferred over back-drilling. This is because through-drilling can be performed on a panel book (or deck), wherease back-drilling must be performed on a single panel. This means that the cost of a back-drilled via is much greater than that of a regular through-via in production. This also means that although back-drilling is permitted, that a particular design might not require any back-drilling.

DFM rules for buried vias. The fabrication rules and parameters for back-drilled vias, illustrated in *Fig. 21(30)*, are as follows:

Back-drill size, D_b: The drill size of the back-drill. This is typically the via drill size (unfinished hole size) plus 10mil.

Minimum stub length, L_S: The minimum distance from the last layer of the via to avoid disconnecting the via from its last functional layer. This value is typically 5mil. The availability of values beneath 5mil depends heavily on the drilling equipment used.

Additional requirements for DFM are as follows:

- Proper clearances to the back-drilled hole must be considered. These are the normal clearances to NPTH or second-drilled holes. It is interesting to note that the clearance of copper to NPTH is usually greater than the clearance of copper to via-pad or PTH, so back-drilled vias cannot increase the design density on the basis of clearance.
- As back-drilled holes are drilled in a separate drilling process, their xy-plane positional tolerance can combine with that of the first drilling process. Therefore, twice the drill table tolerance must be considered for inter-hole spacing to maintain a minimum web of laminate between drilled holes to avoid tool breakage.

Aside from the requirements above, back-drilled vias must also follow the rules of the through-drilled via.

Fabrication data requirements for back-drilled vias. The data requirements for back-drilled vias are similar to the requirements for blind or buried via construction: a separate NC drill file is required for each drilling depth. Although performed at about the same process step, NC data for back-drilling and second-drilled NPTH and tooling cannot normally be combined. One of the reasons for this is that the Excellon-2 NC drill file format has no industry adopted way of specifying depth-control. The way to specify drilling to a given depth is to provide a separate NC drill file for each depth.

When there are back-drilled vias, they are not back-drilled during build-up and are not back-drilled before plating: they are back-drilled after plating. A separate drill file is required in this case and it should imply depth-control-drilling (z-axis) information. A separate NC drill file is required for each depth. It is not necessary to provide drill files for all depths. It is cheaper to under-back-drill some vias (leaving a larger stub) so that a smaller set of depths can be provided. Not all vias need to be back-drilled: only high speed signals and high-speed digital grounds and power rails and these can be determined by net and netclass or subnet and pad definitions. The algorithm for determining the set or zones of drill depths should allow up to the maximum permissible stub length for any via.

Typically the number of tools is restricted to the limited tool sizes used for drilling the primary vias plus the necessary oversize for back-drilling. Note, however, that the back-drilling process is performed typically from the bottom side of the board, and so the drill map will be mirrored by the fabricator to reflect the bottom process side.[9]

The back-drilling process also has additional requirements placed on the artwork. Although fabricators can generate the necessary features using the NC drill files and depth information, for plot-and-go services these features must be added by the CAD system. The necessary artwork consists of clearance pads placed in the back-drilled portion to act as resin venting and stabilizers for depth-control-drilling solutions. Feedback systems require a copper pads on the back-drilled portion so that the feedback mechanism can sense the pads. These clearance pads are typically 4mil larger than the primary drill size. Non-feedback systems are calibrated on a corner of the panel and rely on clearance pads in the stub portion to act as resin venting as well as stabilizing the z-axis depth calibration. These clearance pads are typically 4mil smaller than the primary drill size.

Note that when non-functional pads are removed from the initial through-via construction, clearance pads of the appropriate size should be added to each non-functional layer to stabilize the laminate.

3.3.6 Depth-Control-Drilled (DCD) Vias

A depth-control drilled via is illustrated in *Fig. 22(31)*. Depth-control-drilled vias are vias that are drilled from one side of a the laminate or sub-laminate to form a blind or buried via.[10] Depth-control-drilled vias are one of the two ways of fabricating blind or buried vias. DCD vias beyond a given maximum aspect ratio are not normally used for buried vias due to the inability to fill the via. Depth-control-drilled vias have not historically been represented by pcbnew.

The fabrication cost of a board increases dramatically with depth-controlled-drilling. The cost increases with the number of drill hits and the number of depths drilled.[11] Therefore, there must be some way for the designer to control where DCD vias are placed and the depth drilled. This need is very similar to back-drilling. It is again possible to simply right-click the via and transform it from DCD to another via construction; however, that is a manual process. The net info data has been enhanced to include the frequency of the net, and the acceptable stub length.

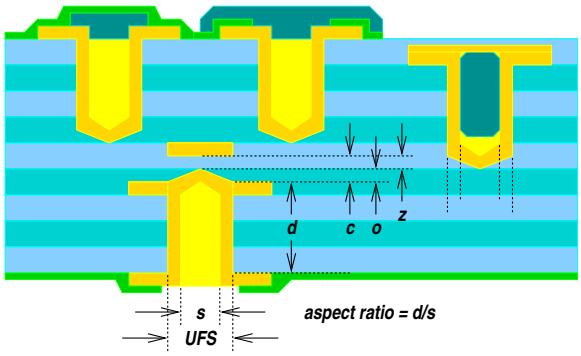

Figure 22: Depth-Control-Drilled Via Construction

This permits the DCD process to be automatically selected for through-drilled laminate or sub-laminate vias that have stubs that are too long. A set of DCD depths may also be defined by the designer and the program can select the available DCD depth that best meets the minimum stub length requirement.

Nevertheless, unlike back-drilling, the purpose of DCD vias is not just signal integrity. The additional purposes of DCD vias are as follows:

1. Increasing the density of a board design while trading off the cost of through-drilled sub-laminate vias. Generating blind vias in a sub-laminate can avoid breaking the sub-laminate into smaller sub-laminates. The material handling of creating a sub-laminate can easily be more expensive than the DCD process.
2. Generating blind or buried vias across a number of layers that could not otherwise be generated due to sub-laminate buildup sequences. Because sub-laminates must be built up using symmetrical lamination sequences, it might not be possible to through-drill a pair of sub-laminates that must be backed by another for symmetry. In this case, it is possible to DCD drill the sub-laminate from one or both sides to construct the necessary vias. The via construction might be demanded by board density, signal integrity, or both.

Note that a DCD created blind or bured via could have a stub on either or both sides of the via.

(R) 6 (pcbnew) *To accommodate this control over DCD vias, the stack-up model has been provided to specify for each sub-laminate sequence whether DCD drilling is permitted, the process side(s) on which it is permitted, and the permitted drilling depths (in terms of layers) for each process side.*

Where through-hole drilling is permitted for the sub-laminate, and the characteristics of the via permit (e.g. DRC and signal integrity stub length), through-hole drilling will always be preferred over DCD drilling. This is because through-hole drilling can be performed on a panel book (or deck), whereas DCD must be performed on a single panel. This means that the cost of a DCD via is much greater than that of a through-hole drilled via in production. This also means that although DCD vias are permitted, that a particular design might not generate any.

9. Most fabricators want all data as viewed from the top to bottom through the PCB and do not want any layers mirrored (or reflected), whether artwork or drill files.

10. Note, however, that using depth-control-drilling to form a buried via would be quite expensive. Nevertheless, in some configurations it might be the only way to create the necessary buried vias.

11. Some setup time is required for each drill depth, particularly with calibarated depth drilling versus feedback depth drilling.

DFM rules for depth-control-drilled vias. The critical fabrication parameters, illustrated in *Fig. 22(31)*, are as follows:

Maximum aspect ratio, α_{max}: There is a maximum aspect ratio, $\alpha = d/s$, (dictated by the plating process), beyond which plating is low-yield. This maximum aspect ratio is specified by the board fabricator. It applies to all depth-control-drilled mechanical-blind vias on the production panel. Typical maximum aspect ratios are 0.5–0.75.

Minimum (finished) hole size, s_{min}: There is a minimum hole size, (dictated by the surface tension of the plating solution), beneath which plating is low-yield. This minimum hole size is specified by the board fabricator. It applies to all depth-control-drilled mechanical-blind vias on the production panel. Typical minimum drill size is 10–6mil for a minimum finished hole size of 8–4mil.

Maximum (finished) hole size, s_{max}: There is a maximum hole size, (dictated by the ability to plate the floor of the via), above which plating is low-yield. This maximum hole size is specified by the board fabricator. It applies to all depth-control-drilled mechanical-blind vias on the production panel.

Depth clearance, c_{min}: There is a minimum depth (z-axis) clearance between the last layer drilled and any copper feature on the layers beneath the last layer drilled. This is dictated by the drilling machine (its ability to and predict depth). This minimum depth clearance is specified by the board fabricator. It might vary somewhat over the full range of possible finished hole sizes. It applies to all depth-control-drilled mechanical-blind vias on the production panel. A typical controlled depth margin is 5mil.

Annular ring calculations follow normal layer rules. Also, xy-plane clearances follow normal rules.

Optimization of depth-control-drilled vias. There is a cost associated with depth-control-drilled vias. An additional cost is associated with each depth drilled due to the setup costs associated with each depth. Therefore, a reduction in the number of depths drilled can reduce cost. To give the designer some control over the number of depths drilled, or the set of allowable depths, would be advantageous. When the set of allowable depths are provided, pcbnew could fabricate vias to meet the specification where possible and provide a DFX error otherwise.

Fabrication data requirements for depth-control-drilled vias. The data requirements for depth-control-drilled vias is similar to that of through-drilled vias for constructing blind or buried vias: the NC drill files must identify the first and last layer through which the bit will penetrate. An additional requirement for DCD vias, is an identification of the laminate or sub-laminate being drilled, and thus the process side of the sub-laminate.

There will be a separate NC drill file, drill map and drill report for each depth drilled. NC drill files, maps, and reports, should be named in such a way that the layer-pair drilled is identified. The name will, therefore, include the top layer drilled and the bottom layer drilled. NC drill files, maps and reports are always generated as viewed from the top of the board. Some CAD tools used *.ncd or *.drl for the NC drill file, *.gbr or *.dpf for the drill map, and *.rep for the ASCII drill report. These tools used the same base filename for all three files. For example: board-DCD-L1-L4.ncd, board-DCD-L1-L4.gbr, and board-DCD-L1-L4.rep.

Non-functional pads: As for other drilling processes, it is necessary to create clearance holes of 1/2 the F.H.S. on ground-planes that do not have thermals, to reduce drill breakage. Also, on fabrication plots, all non-functional pads should be removed.

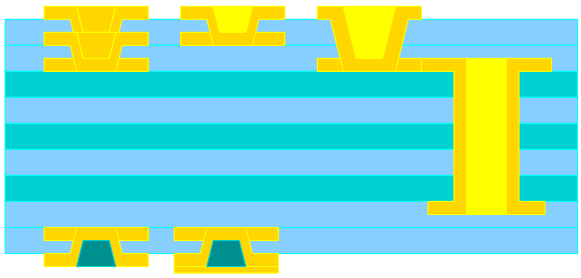

Figure 23: Sub-Composite Via Construction

```
4. ALL VIAS HAVE BEEN DESIGNED WITH NON-FUNCTIONAL PADS
   REMOVED.
```

Clearance pads: The depth-control-drilling process also has additional requirements placed on the artwork. Although fabricators can generate the necessary features using the NC drill files and depth information, for plot-and-go services these features must be added by the CAD system. The necessary artwork resin venting and stabilizers for depth-control-drilling solutions. Feedback systems require copper pads on the depth-control-drilled portion so that the feedback mechanism can sense the pads. These clearance pads are typically 4mil smaller than the drill size. Non-feedback systems are calibrated on a corner of the production panel and rely on clearance pads in the depth-control-drilled portion of the via to act as resin venting as well as stabilizing the z-axis depth calibration. These clearance pads are typically 4mil smaller than the drill size.

Note that, where possible, non-functional clearance pads (neutral, no-net) should also be added on the copper layers through which the drill does not pass, as well as any layers internal to the via where non-functional pads are removed.

In addition to these requirements, it is common to place a circular clearance pad with a diameter of 1/2 the F.H.S. on each layer through which the drill passes, where no other pad or plane exists, to vent resin in the laminate and better stabilize the layers for depth-control-drilling. These clearance pads need to be plotted on the artwork. A note explaining their presence should appear on the fabrication print or drill map detail. For example:

```
5. ALL DEPTH-CONTROL-DRILLED VIAS HAVE BEEN DESIGNED
   WITH A CLEARANCE PAD OF ONE-HALF THE F.H.S. ON EACH
   INTERMEDIATE LAYER TO LOCALLY STABILIZE THE LAMINATE
   FOR DRILLING.
```

3.3.7 Sub-Composite Vias (SCV)

Sub-composite vias are vias that transition layers in a step. Unlike other via constructions, the sub-composite via is a compound construction of two vias and a track. The sub-composite is a feature that typically extends between a microvia and a buried via. Sub-composite vias have not historically been represented by pcbnew.

Identification of sub-composite vias. SCVs can be identified by their intervening track. When the track has a requirement for teardrops (see *Sec. 3.4.2(35)*) a both ends, and the length of the trace is beneath a given length, it is a candidate for a sub-composite via construction. In this case, a trapezoid connecting the two via pads can be formed where the opposite sides of the trapezoid are tangent to each of the pads. This is a sub-composite via construct.

Fabrication data requirements for sub-composite vias. Data requirements for sub-composite vias are simple artwork to

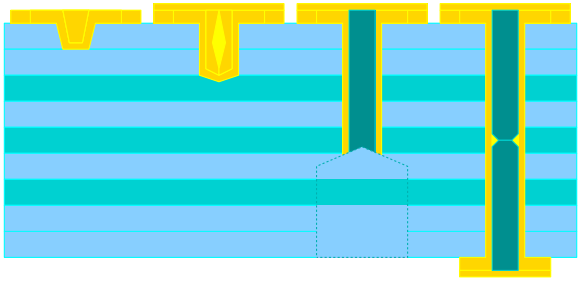

Figure 24: Via Through Pad Construction

describe the intervening trace as a trapezoid connecting to two via pads. For artwork plots (such as Gerber and DPF), this is a simple matter of plotting the trapezoid instead of the track. For CAM formats (such as GenCAM and 258X) there is no way of describing a sub-composite via parametrically, so, again, a simple plot of the resulting trapezoid is performed.

Note that, unlike teardropping,[12] forming the composite trapezoid and the resulting composite shape might violate DRC rules in comparison to the simple trace. This is because the additional copper fills a relatively large area with respect to the microvia pad. It is quite possible that another copper feature intrudes on this area. Because of this, it is unlikely that fabricators can perform the job of including sub-composites on an arbitrary board. Therefore, fabricators ask that designers consult before requesting sub-composites. We make the CAD system do the heavy lifting and the resulting flashed feature can easily be DRC checked by the CAM system on input data.

3.3.8 Via-Through-Pad (VTP)

Via-through-pad (VTP) is a form of via construction used on high-density card designs where the via is permitted to pass through an SMT land or BGA (solder bump or ball) pad on the board's surface. Several via-through-pad constructions are illustrated in *Fig. 24(33)*.

There are two types of VTP vias: those that are centred on the pad and those that are adjacent to the pad. Those that are adjacent or tangent to the pad are often distinguished by the term VBP (Via Beside Pad).

To avoid deteriorating the performance of the joint between the component and the pad, VTP vias are typically either plated closed, filled with a non-conductive material that is plated over with a conductive surface, or filled with a conductive material. The process of filling a VTP via has limits on the aspect ratio of the via hole and the minimum and maximum hole size. Also, VTP constructions can be limited by the type of surface finish. VTP constructions can be left open when beneath a critical hole size (dependent upon the surface tension of the solder material). For BGA solder ball, it is best to place such a VTP beside the pad (VBP) instead of directly under the pad. VTP for BGA solder balls should otherwise be centred. In general, the same approach should be taken for all VTP vias on a board, although pcbnew should permit customization on a via-by-via basis.

Historically, pcbnew has permitted placement of any via on or near a same-net pad, and has not performed the necessary checks to determine whether the placement can be fabricated. An additional attribute is required for objects of class SEGVIA to determine whether VTP or VBP construction is permitted and, when present, which type of filling is to be performed.

Identification of VTP and VBP vias. VTP vias can be detected by checking whether the via hole (centre of the via) meets the criteria of VTP construction. That is, whether the via is in proximity of a same-net pad on the board surface. VTP and VBP can be distinguished by whether the via is centred on the pad (VTP) or simply close to it (VBP), where "close" is defined as too close for non-VTP construction. A class BOARD design rule is necessary to specify whether VTP or VBP construction is permitted on a board at all.

Fabrication data requirements for VTP and VBP vias.

3.3.9 Thermal Vias.

Thermal vias are vias of special construction that are intended to transfer heat from a component into the copper planes in a multilayer board. They are used for specialized components that have thermal grounding pads under the component. Texas Instruments, Linear Devices, and Analog Devices have a number of components with this requirement. The construction of thermal vias inexorably include filling the hole of the via with a conductive material (electrically conductive materials also have better heat-transfer characteristics than the usual non-conductive fillings).

Rather the using via editing to create thermal vias, it was easier to provide for a special pin construction as part of a module. Thermal vias are directly linked to a module (component) and their placement is always constrained relative to the component. A number of thermal vias might be placed through the a single ground pad.

Fabrication of thermal vias. Fabricating thermal vias requires a process to fill the via hole with conductive material. Thermal vias can be filled as part of the board fabrication process, or as part of the board assembly process. When performed by the board fabricator, a process is required to fill the vias. This is likely the normal filing process for filling vias with conductive material. When performed by the assembler of components, these vias can be filled with (solder) paste by including an opening for the hole in the paste stencil.

Data requirements for thermal vias. Thermal vias must be identified so that they may be filled with conductive material. Otherwise, thermal vias are normal PTH described for vias. The attribute associated with the class D_PAD object describing the thermal via identifies the via as a thermal via. Process characteristics identify the filling process and material. The filling process can be a board fabricator process, in which case the holes to be filled must be identified. This can require the generation of a separate stencil for the filling process. When the process is performed as a part of assembly, the via hole openings can either be included on the normal paste stencil, or artwork for a separate stencil can be provided.

3.3.10 Ground Plane Stitching.

Ground plane stitching consists of vias that are used to connect multiple ground planes together. Historically, pcbnew could not create an isolated via. The track and via editing process under pcbnew required that a track start with a trace before a via can be created. This causes some difficulties in convincing the program to generate vias for ground stitching.

Ground plane stitching vias are no different in construction than other vias. They can actually be any of the via types (through, blind, buried, etc.). There are no data requirements to identify ground plane stitching vias to fabricators that are different from that for other vias.

3.3.11 Isolation Vias.

Isolation vias are a special type of via that is used to isolate or contain high-frequency signals within the board construction. They are typically stitched along the sides of internal traces to

12. See *Sec. 3.4.2(35)*.

place a metalized picket-fence around the traces internal to the board. They are used in a similar way as board edge plating, that is, their purpose is one of shielding and EMC containment. Isolation vias are similar to other vias with the following exceptions:

1. The vias are neutral copper and are not connected to a net.
2. The vias cannot be connected together with traces, otherwise eddy loops will be formed that can deteriorate the signals that they are intended to isolate.

Item (1) causes the difficulty that the vias will not pass DRC check muster; causes the vias will be removed by a number of track operations; and DRC checks will deem the vias as "unconnected." The global board operations to delete dead-end traces will likely remove these vias in their current trim.

(R) 7 (pcbnew) *Enhancements will be provided to pcbnew to permit the use of isolation vias. Likely the simple way to accommodate this is to define a "neutral" copper net.*

3.4 Tracks

This section addresses the enhancements considered for tracks. Enhancements include the following:

Etch compensation. Etch compensation is not performed for the purpose of dictating fabrication of the board. The fabricator will apply etch compensation to meet their own processes. The reason etch compensation is used to more accurately estimate press-out of B-stage sub-laminates, impedance of controlled traces, and provide as-fabricated 3D display. . Enhancements required to address etch compensation are discussed in *Sec. 3.4.1(34)*.

Teardrops. The application of teardrops can increase the density of a board, maintain or enhance its yield, and increase the signal integrity and power transfer characteristics of vias and pads. When performed by the CAD system, teardrops do not add cost to the fabrication processes for a board. Enhancements required to address teardrops are discussed in *Sec. 3.4.2(35)*.

Sub-Composite Vias. Sub-composite via constructions can increase the density a board and increasing yield. When performed by the CAD system, sub-composite vias do not add cost to the fabrication processes for a board. Enhancements required to address sub-composite vias are discussed in *Sec. 3.4.3(37)*.

Isolation traces. Isolation traces can be used to isolate fast rise-time signal traces (such as non-spread-spectrum clocks) from other sensitive traces, increasing the signal integrity of the protected traces. Isolation traces do not add cost to the fabrication processes for a board. In fact, isolation traces can remove some need, or improve the yield, for plate thieving or resin venting. Enhancements required to address isolated traces are discussed in *Sec. 3.4.4(38)*.

Arced tracks. Arced tracks can be used on high-speed traces to increase the signal integrity of the trace, affecting neither the board fabrication yield nor the board fabrication cost. In fact, arced tracks can increase yields by avoiding blossoming or mushrooming of the etch pattern at sharp interior corners. Enhancements required to address arced track segments are discussed in *Sec. 3.4.5(38)*.

Push/pull tracks. The pushing and pulling of tracks is a method that is useful for enabling the editing of tracks during manual routing to assist with increases in the density of a board as well as assisting in increasing the signal integrity associated with the affected traces. Enhancements required to address the pushing and pulling of tracks during editing are discussed in *Sec. 3.4.6(39)*.

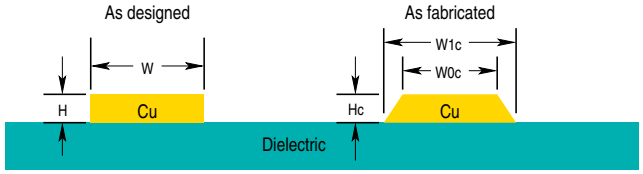

Figure 25: Etch Compensation

Neck-downs. Enhancements required to address neck-downs of traces are discussed in *Sec. 3.4.7(39)*.

3.4.1 Etch Compensation

Etch Compensation is the process of altering CAD widths of traces to compensate for the etch factor of the trace. The etching process is more rapid at the top of a trace than at the bottom of the trace. The resulting cross-section is roughly trapezoidal instead of rectangular. The thicker the copper, the more the width of the top of the trapezoid deviates from the bottom. A trace as designed compared to the same trace as fabricated is illustrated in *Fig. 25(34)*.

The reasons that etch compensation is calculated is to perform the following:

1. Estimate more accurately the press factor of the prepreg between layers. The press factor is illustrated in *Fig. 45(54)*.
2. Estimate more accurately the impedance of impedance controlled traces.
3. Provide an as-fabricated 3D display of the board.

Etch compensation is typically calculated and performed as illustrated in *Fig. 26(35)*. Etch factor, e, is a constant determined by the etching process, for inner layers, or the plating and etching process, for outer layers. An etch factor of $e = 0$ would be a perfect process. Etch factors are typically in the range from $e = 0.2$ to $e = 0.4$. The etch-back width, w, can be calculated from the copper thickness, h, and the etch factor, e, by the simply expression $w = e \times h$. The compensated trace width, W_C, the trace width that is imaged in photoresist, can be calculated as $W_C = W + w$, or $W_C = W + e \times h$.

Example:
$e = 0.3$
$W = 0.005$"
$h = 0.002$"
$W_C = W + e \times h = 0.005 + 0.3 \times 0.002 = 0.0056$"
$err = W_C/W - 1 = 0.0056/0.005 - 1 = 12\%$

As can be seen from the example, the significance of the change to the cross-section is greater as the line width becomes narrower.

Note that the cross-sectional area of the compensated trace and the uncompensated trace is the same: $A = W \times h$. This is the characteristic of this compensation method that allows press factor to be calculated using uncompensated traces.

Where the difference between the uncompensated trace and the compensated trace is significant is in the calculation (estimation) of the trace impedance.

Display. When displaying trace widths in pcbnew on the 2D screen, CAD widths are always displayed and compensated

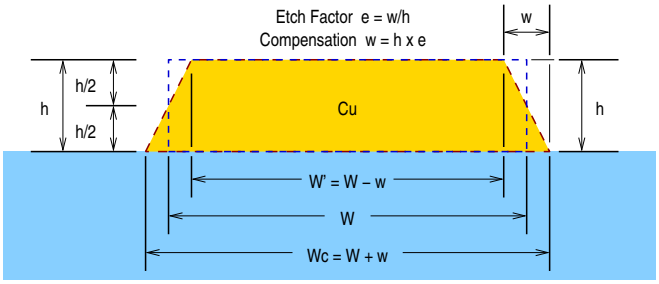

Figure 26: Etch Compensation Detail

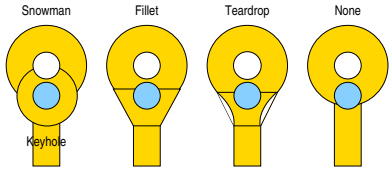

Figure 27: Teardrops

widths are never displayed. This avoids any confusion. Compensated widths and etch factor are only used in impedance calculations and are accounted for in the estimated impedance values displayed in a message box.

3D-Display. For the pcbnew 3D display, however, the intention is that the display be as close to the as-fabricated view as possible. Therefore, etch compensation (the compensated width and the etch-back width) will be displayed resplendent with trapezoidal cross-section in the 3D view. In addition to an as-fabricated view for the designer, the VMRL export function should generate as-fabricated traces in the event that the VMRL model can be used as input to a 3D, FEM, full wave simulation. An option will be provided to suppress etch factor in the event that 2.5D FEM input is desired instead.

Post-Processing. For fabricator outputs, no etch compensation is performed on plots. Fabricators will perform the etch compensation appropriate for their processes when creating the photo-resist images. The only widths present in design files, therefore, are CAD widths and not compensated widths. However, when the KiCad user *is* the fabricator, such as when plots are being developed for use in the toner-transfer method or other techniques such as PCB milling, there needs to be an option to perform etch compensation on the fabrication outputs. Therefore,

(R) 8 (pcbnew) *An option will be provided in the dialogues for generation of fabrication outputs to perform etch compensation on plots of copper layers.*

3.4.2 Teardrops

Teardrops are a track features that is used on the end of a track where it connects to a via or pad where the via or pad is of minimal size causing the hole to potentially break-out from the pad at the position of the trace; or where the solid NSMD pad might be slightly too small to form a complete joint. This construction is termed a "*teardrop*" because one of the possible constructions takes the shape of a teardrop.

The break-out of drilled holes results in the complete or partial disconnection of the trace from the pad. A partial disconnection can affect the signal integrity of high-speed traces, or reduce the current-handling capacity of power or ground traces. This break-out effect can be mitigated by placing additional copper extending from the pad toward and along the attaching trace as illustrated in *Fig. 27(35)*.

In addition to adding teardrops to vias that have drill holes, teardrops can also be added to BGA lands. NSMD BGA pads can be made somewhat smaller when a sub-land is added to the trace side of the circular pad in the same fashion as a via. The most common dimensions for this construction is the sub-land one shape (see *Fig. 30(36)*), where the radius of the sub-land, A, is the radius of the pad, R, less 1mil, $A = R - 1$, and the offset from the centre of the pad is 3mil, $B = 3$, to provide an additional annulus of $B - A = 2$mil on the keyhole side. This arrangement should permit shaving a full 2mil off of the size of the BGA land. In this case, what the sub-land serves to do is to remove the stress riser from the junction of the trace and the land, as well as providing some additional copper for the solder joint. This is a somewhat different use of teardrops, but is common and appropriate. The difference between the BGA goals and the via goals stresses that the teardrop simply manages to add a little more copper at the position of the trace, where it is least likely to cause problems with DRC, DFM and layout.

Placement of teardrops does not add to the cost of the fabrication of a board because the constructions are a normal part of the artwork associated with copper layers. This is a low-cost approach to increasing the density of designs while maintaining yields.

The constructions used for teardrops have the characteristic that they do not affect DRC checks, because they simply fill the two little depressions between the pad and the trace at the entry point to the pad. This is a localized depression that can be filled without the possibility of encroaching upon another board feature. Typical constructions are 'fillet', where a trapezoidal shape is used; 'keyhole', where a circular shape is used, and 'teardrop', where a flute is used.

The criteria used for determining the requirement and placement of teardrops on vias is well known. It is the familiar annular ring analysis. The analysis proceeds by determining the potential for break-out of the drilled hole using the positional tolerance of the hole and the surrounding artwork in the xy-plane. Where a break-out would occur and a trace is attached, the attaching trace can have a teardrop applied using one of the appropriate methods.

(R) 9 (pcbnew) *An attribute will been added to class TRACK to specify whether teardrops are required and applied to either end of the track, and to identify the construction and parameters of the teardrop.*

(R) 10 (pcbnew) *An attribute will be added to class D_PAD to specify whether teardrops are required and applied to to the pad, and to identify the construction and parameters of the teardrop.*

(R) 11 (pcbnew) *A general track option will be provided to specify whether teardrops are to be automatically applied to traces that connect to pads where breakout and disconnection of the trace could otherwise occur.*

Fillet. *Fig. 28(36)* illustrates the geometry of a fillet-shaped teardrop. This teardrop geometry is taken from the GenCAM specifications [GenCAM, GenX]. The shape of the fillet is described by the two parameters, A and B, where A is the height of the fillet triangle, and B is $\frac{1}{2}$ of the base of the fillet triangle.

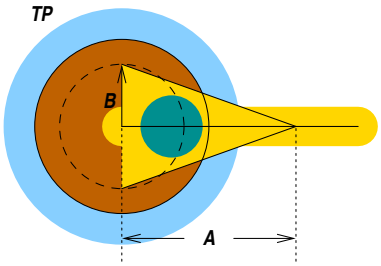

Figure 28: Fillet

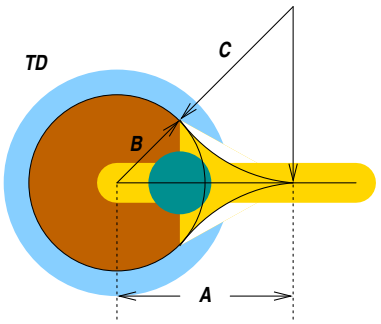

Figure 29: Teardrop

This shape can be rendered as a triangle flash aperture and can be combined with the pad into an aperture macro.

Teardrop. *Fig. 30(36)* illustrates the geometry of a teardrop-shaped teardrop. This teardrop geometry is taken from the GenCAM specifications [GenCAM, GenX]. The shape of the teardrop is described by two parameters, A, B, and C, where A is the distance of the teardrop radius from the centre of the pad, and, B is the radius of the circle tangential to the teardrop radius. The radius, R, can be calculated as:

$$\begin{aligned} A^2 + C^2 &= (B + C)^2 && (16)\\ &= B^2 + 2BC + C^2 && (17)\\ A^2 &= B^2 + 2BC && (18)\\ C &= \frac{A^2 - B^2}{2B} && (19)\\ &= \frac{1}{2}(\frac{A^2}{B} - B) && (20) \end{aligned}$$

This shape is rather difficult to render as a RS-274X flashed aperture.

Sub-Land One. *Fig. 29(36)* illustrates the geometry of a sub-land-1-shaped teardrop. This teardrop shape is also known as a "snowman" or "keyhole." The smaller non-pad disk is called a "sub-land." This is perhaps the most common teardrop shape cited by manufacturers. It is also the easiest to represent with RS-274X flashes, aperture draws, or complex aperture macros. This teardrop geometry is taken from the GenCAM specifications [GenCAM, GenX]. The shape of the teardrop is described by two parameters, A and B, where A is the radius of the sub-land, and B is the distance from the centre of the pad to the centre of the sub-land.

Approach 1: The sub-land should extend 5mil into the trace from the pad; therefore, if R is the radius of the pad, and r is the radius of the drill, and w is the width of the trace,

$A + B - R = 5\text{mil}$, or $B = R - A + 5\text{mil}$.

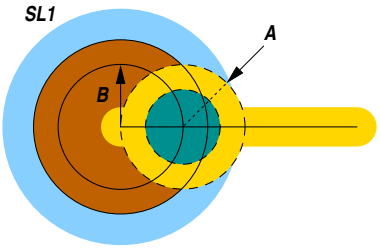

Figure 30: Sub-Land One

The amount of annulus about the drill hole should be one-half of the track width, so,

$$A = r + \frac{1}{2}w. \quad (21)$$

Substituting,

$$B = R + 5 - r - \frac{1}{2}w\text{mil}. \quad (22)$$

For an example pad of 20mil with a 10mil hole size and a 6mil line width, $R = 10\text{mil}$ and $r = 5\text{mil}$, so

$$\begin{aligned} A &= r + \frac{1}{2}w = 8\text{mil, and} && (23)\\ B &= R - A + 5 = 7\text{mil}. && (24) \end{aligned}$$

This shape can be rendered as a 8mil circle flashed aperture, flashed 7mil along the line of the trace, and can be combined with the pad into a complex aperture.

Approach 2: Another way to calculate the parameters is to set B to the maximum offset of the finished hole from the centre of the pad, and, assuming that the hole is tangential to the pad, set $A = r + 5\text{mil}$ so that the sub-land extends 5mil into the trace. Then, for the above example,

$$\begin{aligned} A &= r + 5 = 5 + 5 = 10\text{mil}, && (25)\\ B &= R - r = 10 - 5 = 5\text{mil}. && (26) \end{aligned}$$

which would result in flashing the same pad 5mil along the line.

From the above, it can be seen that the calculation of the sub-land radius and position can be complex. TI recommends [TI-2010] a sub-land that is 2mil smaller than the pad size, offset 3mil along the trace. This would only provide 2mil of additional annular ring at the breakout point. TI recommends this for BGA pads whether they have a via in them or not.

Endicott recommends [EI-2010] the following calculation for sub-lands for vias with a possible breakout (IPC Class 2 board):

$$\begin{aligned} R &= \text{radius of the pad;} && (27)\\ r &= \text{radius of the hole;} && (28)\\ A &= \text{radius of the sub-land;} && (29)\\ B &= \text{offset of the sub-land;} && (30)\\ s &= \text{trace to pad space;} && (31)\\ A_{min} &= r + 1\text{mil} && (32)\\ B_{pad} &= s - 2 + R - A + 1.5\text{mil} && (33)\\ B_{pin} &= s - 4 + R - A + 1.5\text{mil} && (34) \end{aligned}$$

Following the previous example,

$$\begin{aligned} R &= 10, r = 5, s = 6, \text{ then} && (35)\\ A &= r + 1 = 6\text{mil}, && (36)\\ B_{pad} &= s - 2 + R - A + 1.5 && (37)\\ &= 6 - 2 + 10 - 6 + 1.5 && (38) \end{aligned}$$

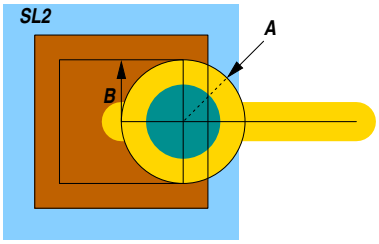

Figure 31: Sub-Land Two

$$= \quad 9.5\text{mil, and} \tag{39}$$

$$B_{pin} = s - 4 + R - A + 1.5 \tag{40}$$

$$= 6 - 4 + 10 - 6 + 1.5 \tag{41}$$

$$= 7.5\text{mil.} \tag{42}$$

As can be seen from the calculations, this is quite different from the case of 0.002" smaller than the pad, offset by 0.003". This example has a sub-land that is 0.008" smaller than the pad, offset by 0.0075" to 0.0095", resulting in an extension into the line of 0.0035" to 0.0055". This is, however, in fitting with the Samina-SCI recommendations [SS-2007] to extend the teardrop into the trace by 5mil.

Sub-Land Two. *Fig. 31(37)* illustrates the geometry of a sub-land-2-shaped teardrop. It is used for square pads. This teardrop geometry is taken from the GenCAM specifications [GenCAM, GenX]. This shape can be rendered as a circle flash aperture and can be combined with the pad into a complex aperture.

Fabrication data requirements for teardrops. Because of the characteristic that teardrop application does not affect DRC calculations enough to require movement of features, it is possible for a CAM system to apply teardrops where necessary based on limiting criteria.[13] Nevertheless, for plot-and-go services, it is essential that the CAD system also be capable of applying teardrops using similar criteria. When the fabricator is expected to add teardrops where necessary, there are no special data requirements. However, where the CAD system is adding teardrops, it is necessary to express the teardrops in the format required by fabrication outputs.

For data formats not specialized enough to permit identification of teardrops and teardrop construction, it will be necessary for pcbnew to create the artwork necessary for the construction. This applies to outputs unable to express teardrops in any other form, such as *Gerber*, *DPF*, and *GenCAD* outputs. The new attribute for class TRACK that identifies the presence and construction of the teardrop is used by the plot routines to generate artwork for the teardrop. The same approach is used to display teardrops on pcbnew's canvas and for the print routines, making the application of teardrops a WYSIWYG operation. Plotting teardrops on production *Gerber RS-247X*, *DPF* and *GenCAD* outputs cause some potential issues as follows:

1. Pads are normally flashed and other items are not. Flashing apertures to form sub-lands, fillets or tapers, can cause problems for import functions in CAM systems. This may be the real reason why fabricators normally apply their own teardrops.
2. The normal way to generate a sub-land for sub-land one and sub-land two constructions would be to flash a circular aperture at the location of the sub-land. Another method would be a zero-length draw with the same aperture.

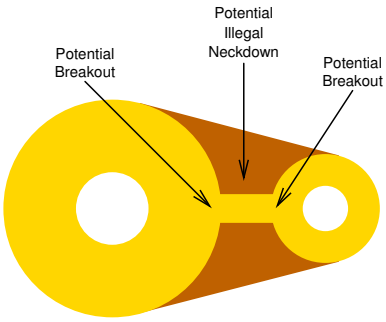

Figure 32: Sub-Composite Via

3. The normal way to generate a sub-land for fillet would be to define a complex aperture for the triangle and then flash it at the appropriate location and orientation. Another method would be a zero-length draw with the same aperture. Yet another method would be to draw the outline of the fillet at the necessary location using polygon fill commands.
4. The normal way to generate a sub-land for teardrop would be to define a complex aperture for the taper and then flash it at the appropriate location and orientation. Another method would be a zero-length draw with the same aperture. Yet another method would be to draw the outline of the teardrop at the necessary location using polygon fill commands.

Fabricators may wish to define their own teardrops, or designers may wish to address teardrops with a callout instead of with design data. Therefore,

(R) 12 (pcbnew) *There will be an option to suppress generation of teardrops on plots, as well as a collection of teardrops into their own information layer within the RS-274X file for ease of deletion.*

In some CAM formats (such as *GenCAM*, *GenCAM-XML* and *258X*) it is possible to identify traces that require teardrops, the end of the trace upon which the teardrop is applied, and the particular construction for the teardrop. These formats do not require additional artwork to form teardrops, but require that the specification of a teardrop at the end of a trace be represented using the native format.

3.4.3 Sub-Composite Vias

A Sub-Composite Via construction is a copper feature in which the perimeter of the feature is blended into the circumference of the pads of two vias as illustrated in *Fig. 32(37)*. Typically one via is a microvia and the other is a buried via. This is because the structure is most useful when there is a significant difference between the sizes of the vias and illegal neck-downs must be avoided. Historically, pcbnew has not addressed this feature. Fabricators will add sub-composite via features to a board in the same fashion as they add teardrops (see *Sec. 3.4.2(35)*). However, for plot-and-go services, as well as reducing product dependence on one fabricator, we have added sub-composite via constructions to pcbnew. The SCV structure is treated as a track feature in a way similar to teardrops.

13. A fabricator may add teardrops to via holes where the track enters the pad wherever the annular ring is small enough to cause potential breakout and cannot be increased in size. Note that teardrops are typically added *by the CAM operator* making it a per-run fixed cost item to be avoided.

Figure 33: Arced vs. Line BGA Breakout

(R) 13 (pcbnew) *When a track has teardrops set for both ends of the track, and the vias are beneath a maximum distance, SCV structure will be applied.*

(R) 14 (pcbnew) *A general track option will be provided to allow or disallows the automatic application of sub-composite via structures.*

The criteria and thresholds for applying SCV structures should be based on minimum space and trace rules for the layer to which it is applied in the same fashion as illegal neck-downs are detected, see *Sec. 3.4.7(39)*.

3.4.4 Isolation Traces

Isolation traces are neutral (no net) traces that are placed between an aggressor trace and a victim trace to reduce the near- and far-end cross-talk between traces, or to isolate a trace of a differential pair. They are typically used where a full clearance cannot be obtained between, say, a fast rise-time clock trace and a trace of a differential pair. Historically, isolation traces caused problems for pcbnew in a number of respects:

1. pcbnew does not distinguish between neutral copper (no net) traces and any other traces. Neutral copper traces are lumped into the default net class.
2. pcbnew applies the same clearance rules to all neutral copper traces as are applied to other, non-neutral traces.
3. pcbnew basically treats neutral copper (no net) traces as errors and removes them during redundant track operations and when automatically removing unconnected traces during net-list import.

Providing support for isolation traces is a non-fabrication design issue. That is, it supports the electrical design specification of the board rather than the tailoring of board features to a particular process.

(R) 15 (pcbnew) *pcbnew will be enhanced to support isolation traces. Likely the easiest way to accomplish this is to define a special net for "neutral copper."*

3.4.5 Arced Tracks

It is well-known that beyond a given frequency, 45-degree angle tracks are no longer preferred. At almost all frequencies requiring controlled impedance, arced (or swept) traces are preferred. Arced tracks are preferred for the following reasons:

1. *Signal integrity:* Arced tracks improve the signal integrity of high-frequency critical traces as follows:

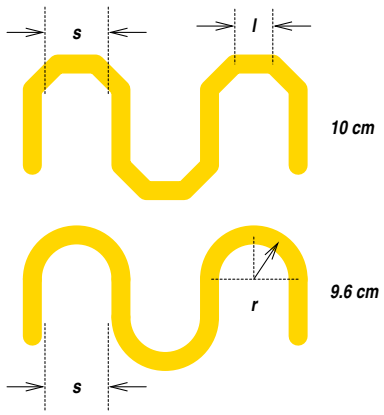

Figure 34: Arced vs. Line Meander

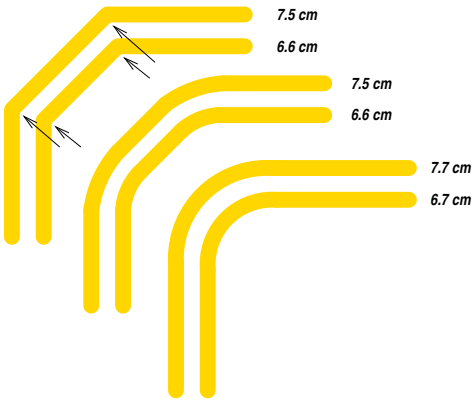

Figure 35: Arced vs. Line Bend

(a) Arced tracks provide a smoother transition of impedance from one direction of travel to another when routing traces. See *Fig. 35(38)*.

(b) The minimum arc radius can easily be calculated to meet the signal integrity requirements of a given net.

(c) Arced tracks provide for a less disruptive meandering (or accordion) pattern for length-matching of differential or bussed nets. See *Fig. 34(38)*.

(d) Arced tracks provide for easier length matching in general.

(e) Let's face it: those sharp corners are freakin' antennae at some frequency.

The shape of features that are close to $\frac{1}{4}$ the wavelength of the fundamental or harmonic frequencies of a signal are significant.

2. *Etch definition:* Sharp interior angles can cause problem with the process for etching and plating of copper traces. The interior angle can under-plate when plating, and can over-etch when etching, causing a blossoming or mushroom deformation of the trace that reduces yields and affects the width (and thus the impedance) of the resulting traces. Removing interior angles by using arced traces can *increase* yields. Problem areas are illustrated with arrows in *Fig. 35(38)*.

3. *Zero cost feature:* Adding arced traces to a board can increase signal integrity and yields without adding any fabrication cost. The CAD system is completely responsible for all of the complexity of arced traces. Provided that the CAD system does not approximate arcs with line segments, the size of Gerber file is not increased.

For these reasons, work was done to add support for arced tracks.

There are several significant problems with 45-degree angle tracks under pcbnew when laying out high-speed (10Gbps) and very-high-speed (100Gbps) boards. For high-speed traces, it is necessary to observe minimum radius and same-track separations (e.g. for meanders) even when not using arced traces. None of the DRC calculations covered this, and it is very easy to manually lay out erroneous tracks (of course, an auto-router can lay out erroneous tracks much faster than a human). Arced tracs improve the layout as follows:

1. Signal integrity for high-speed, high-frequency and fast rise-time traces is improved.

2. Tracks can be laid at a somewhat higher density where necessary. When two traces meet at a central point, an arc can curve away from the point allowing the other track to curve away in the other direction also. When only line segments are used, the only turning point is the same for both tracks.

3. Minimum arc radius can be enforced.

4. same-trace separation (especially in meanders) can be more easily controlled by the minimum arc radius.

5. Meanders can be more compact and yet do not violate any rules.

Specifying an arced track.

Two segment build. Two-segment build becomes one or two segments and a 45° arc. The arc radius is taken from the minimum arc radius for the net begin edited. When the minimum arc radius is zero, no arc will be used.

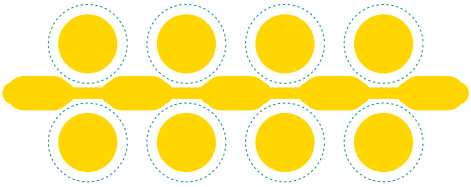

Figure 36: Example application of a Neck-down trace

3.4.6 Push/Pull Tracks

To fit a new traces between a couple of immovable objects along side other traces can be a challenge for manual routing. There are two conflicting desires: (a) to have the traces spaced as widely as possible to improve signal integrity; and, (b) to include as many traces as possible, adding them one at a time. One ways to accommodate this (currently with pcbnew) is to drag some traces to one side, one-by-one, until there is enough room for the new trace, and then laying the new trace between the immovable objects. The issue with this way is that when 64 to 100 traces are between the immovable objects, the shifting of lines can be time consuming and error prone. If nothing else, it takes a lot of mouse-clickin'.

The solution, of course, is to have the new track push or shove its way into the new position, dislodging the existing tracks in the process. This is the same as the manual process; however, the calculations are done by the computer and the function can be performed in real-time so that the operator instantly sees the results as the mouse cursor is moved. The opposite is also possible, where removal of one of the tracks causes the remaining tracks to slide back to the position they would have been in before the removed track was added. This can also be performed when the track is pulled back while editing. This is another good job for a computer.

These features are not too difficult to effect. There is already a check for intrusions while editing so that one track can be laid along another. The trick is that when the intrusion is found, to move it also (drag while keeping slope) until it bangs into the next intrusion.

This feature is not so high on my list because it affects the user experience during the *editing process* and not so much the resulting *quality of the board*. In contrast, a feature such as arced tracks increases the *quality of the board*, and might even degrade the user experience *during editing*.

3.4.7 Neck-Downs

Neck-downs is a situation where a trace of one thickness is connected directly to a trace of narrower thickness, and then connected, again, directly to a track of greater thickness. Another neck-down condition is where two traces of the same thickness are not directly connected, but overlap at their rounded ends by some fraction of the circles.

A typical intentional use of neck-downs is power busing through a BGA field. This is illustrated in *Fig. 36(39)*.

Neck-downs can easily violate DRC rules for trace and space (or feature size). Two such conditions are illustrated in *Fig. 37(40)*. Historically, pcbnew has not handled neck-downs. This enhancement provides DRC rules for neck-downs. In the first case illustrated in *Fig. 37(40)*, same-net DRC checks can be performed to determine that the two copper traces are touching, but are not physically connected. In the second case, without full geometric analysis, it can be determined that there is a violation because the traces are connected and the length of the neck trace is too short to meet trace/space rules for space.

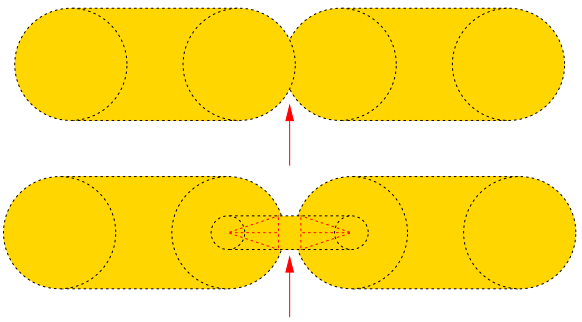

Figure 37: Illegal neck-downs

The length of neck trace that is consumed in the connection to the larger trace at the left, L_l, right, L_r, and neck-down, L_n, where L is the length of the necked-down trace, is given by:

$$L_l = \sqrt{(W_l/2)^2 - (W_n/2)^2} \quad (43)$$

$$L_r = \sqrt{(W_r/2)^2 - (W_n/2)^2} \quad (44)$$

$$L_n = L - L_l - L_r \quad (45)$$

$$= L - \sqrt{(W_l/2)^2 - (W_n/2)^2} - \sqrt{(W_r/2)^2 - (W_n/2)^2} \quad (46)$$

When the neck length, L_n, is negative, it represents the amount of overlap between the segments at the left and right. In this case, the second violation illustrated in *Fig. 37(40)* is indistinguishable from the first violation. When L_n is less than the copper-to-copper space permitted by the copper layer, but not less than $-W_l/2$ or $-W_r/2$, the neck-down is a violation of DRC rules.

This calculation of *Eqn. 46(40)* will be added to DRC checks to identify illegal neck-downs. It will be invoked both when performing bulk DRC checks, as well as when the track widths of a track are changed and DRC is active. The criteria for deciding when to apply the equation to a particular track segment is when the track segment connects to other track segments of different widths, and where those widths are larger than the width of the track being considered.

(R) 16 (pcbnew) *The test for illegal neck-downs on tracks will be added to the DRC calculations for board items of class TRACK.*

3.5 Zones

This subsection addresses the enhancements considered for zones.

Polygon Clipping. The computational cost of filling copper zones appears quite high and does not scale very well, particularly under large BGA fields where the antipad cutouts for a single zone can grow toward a thousand. The polygon clipping algorithm is considered in *Sec. 3.5.1(40)*.

Arced outlines. Zones under pcbnew have historically only supported line segments forming the outline. To fully support some import and export file formats (e.g. IDF[Kehmeier, 1993]), it is necessary that zones permit arcs in the outline. Enhancements required for arced outlines are considered in *Sec. 3.5.2(40)*.

Impedance cutouts. Impedance cutouts are a way of providing cutous in reference planes beneath an SMT land, an edge-finger, or the end of an blind or buried via. In general, impedance cutouts do not increase the fabrication cost of a board and are therefore an excellent way of increasing the signal integrity of a high-speed board without increasing the fabrication cost of the board. Enhancements required for impedance cutouts are considered in *Sec. 3.5.3(40)*.

Copper islands. Enhancements required for copper islands are considered in *Sec. 3.5.5(41)*.

Isolation zones. To isolate high-frequency signals and clocks often requires the use of isolation copper zones that are otherwise non-functional (that is, they are not connected to any other net). Enhancements required for isolation zones are considered in *Sec. 3.5.6(41)*.

Module zones. Historically, modules could not create zones; neither on copper nor technical layers. To support some import and export file formats requires that zones be defined in modules. These zones are used for courtyards, keep-outs and other closed contours, but should also be able to be used for copper zones. Enhancements required for module zones are considered in *Sec. 3.5.7(41)*.

Zone plotting. Historically, pcbnew has painted all zones in Gerbers and other plots. Fabricators do not want painted zones in RS-274X. Enhancements required for zone plotting are considered in *Sec. 3.5.8(41)*.

Zone Plotting.

3.5.1 Polygon Clipping

There are a number of polygon clipping algorithms [Peng et al., 2005b, Kim and Kim, 2006]. Some have been implemented in open-source:

- The kbool library has been used by KiCad, wxArt2D, and libarea.
- libarea also uses another approach implemted by clippoly [Schutte, 1995], but it is unclear whether this approach handles degenerative cases [Modi, 1996].
- kicadocaml uses a home-brew mechanism that it is unclear whether it handles degenerative cases properly or whether it is simply another labelling approach.
- Farseer uses a refined approach [Peng et al., 2005a].
- OpenCASCADE appears to use yet another approach.

From all of the confusion on polygon clipping, it is not clear whether KiCad is taking the best approach by using kbool or not.

3.5.2 Arced Outlines

All of the open-source polygon clipping code, and, in particular, kbool, supports arced segments to form the concave contour. Editing and representing arced segments in KiCad is well supported in both eeschema and pcbnew. There does not seem to be any reason not to permit zones to have arced segments in their outlines.

From the perspective of wanting to properly represent all contours and cutouts of extrusions that can be exported or imported using IDF, it is desirable to have Zone outlines include arced segments in the outline as part of their internal representation, even if the routines to create them by editing are not yet supported.

This feature adds arced segments to zone outlines.

3.5.3 Impedance Cutouts

To increase the signal integrity associated with high-speed contacts (gold-finger), SMT lands, and blind or buried vias, it is a normal procedure to create cutouts in ground planes that are beneath the contact, land or via end. The cutouts serve to remove parasitic capacitance between the large pad or contact land and reference plane, compensating for the impedance mismatch where a high-speed trace meets the land.

Making cutouts under SMT pads, gold-finger contacts, and blind or buried via ends, does not significantly affect the fabrication of the board, nor increase the fabrication cost. There are a few issues associated with cutouts under gold-finger contacts (in

that the reference plane adds to the stiffness and planarity of the board); however, they are not normally significant.

pcbnew has long provided a mechanism for manually creating cutouts in ground planes; however, the manual process is tedious, error prone, and makes it difficult to tweak the cutout dimensions. This enhancement provides the ability to specify the criteria and construction of cutouts so that they may be automatically applied by the CAD system. It allows the user to specify the nets and pads that should have automatic applied, and the construction that should be applied to the particular net.

There are several places that automatic cutouts can be applied:

1. *SMT pads.* Where high-speed signal SMT pads exist, and particularly for pads that terminate impedance controlled single-ended or differential pairs, it is typical to apply cutouts to the reference planes beneath the pads. Cutouts typically follow the shape of the pad and are proportional in size to the pad. Because of the curved nature of the EM fields involved, it is typical to spin a proportional pad about the longest axis of the real pad and cutout the intersection of this pad rotation with the reference planes. The purpose here is to trade-off two effects:

   (a) the parasitic capacitance (and resulting impedance mismatches) between the pad and the underlying reference plane;

   (b) avoiding creating reference plane voids that affect nearby or traversing signals; and,

   (c) reducing isolation of inner board traces by creating holes in isolation planes.

   The method for specifying the construction of the cutouts and the number of reference planes to which the cutout applies must be flexible enough to allow the board designer to choose between these trade-offs.

2. *Contacts.* Where high-speed signal contacts (edge-fingers) exists, and particularly for contacts that terminate impedance controlled signle-ended or differential paired signals, it is typical to apply cutouts to the reference planes (to some depth) beneath the edge fingers. Again, cutouts typically follow the shape of the contact, are proportional to the size of the contact, and may vary in size depending on the depth of the reference plane cut. Trade-offs between reference plane voids and the removal of parasitic lumped capacitance are not as important for edge fingers as they are for SMT pads, due to the relative distance of traces from the edges of the card. The methods of construction should also accommodate edge fingers.

3. *Via ends.* Where a via carrying high-speed, impedance-controlled single-ended of differential pair traces pass through the board, and terminate on a signal layer adjacent to a reference plane, there can exist a parasitic lumped capacitance between the pad at the internal end of the via and the ground plane beneath. It can be advantageous to also create cutouts beneath the via end pad to remove this capacitance. Cutouts of this sort should be specified on a per-net basis; however, it should be possible for the designer to specifically enforce or remove a cutout formed in this way.

3.5.4 Via Stitching

Via stitching consists of vias that are used to connect multiple ground planes together for the purpose of thermal conduction or the control of electrical currents. A typical use of reference plane stitching is to control the current feed into the plane from a power supply. Another typical use is to provide a route for return currents for a single-ended or differential-pair controlled impedance signal as it passes through layers. In this later case, the stitching vias are placed beside the signal vias, particularly for a tuning fork arrangement of differential pair vias where the common mode voltage and current is of significance.

Historically, pcbnew could not create an isolated via. The track and via editing process under pcbnew require that a track start with a trace before a via could be created. This causes some difficulties in convincing the program to generate vias for ground plane stitching.

Ground or reference plane stitching vias are no different in construction than other vias. They can actually be an of the via types (through, blind, buried, etc.). There are no data requirements to identify reference plane stitching vias to fabricators that are different from that for other vias. The requirements for reference plane stitching vias are discussed in *Sec. 3.3.10(33)*.

3.5.5 Copper Islands

3.5.6 Isolation Zones

Isolation zones are neutral (no net) copper zones that are placed between an aggressor trace or pair and a victim trace or pair to reduce the near- and far-end cross-talk between traces, or to isolate a trace of a differential pair. They are typically used where a full clearance cannot be obtained between say, a fast rise-time clock trace and a trace of a differential pair. Historically, isolation zones caused problems for the pcbnew in a number of respects:

1. pcbnew does not distinguish between neutral copper (no net) zones and other zones. Neutral copper zones are lumped into the default net class.
2. pcbnew applies the same clearance rules to all neutral copper as are applied to other, non-neutral copper.
3. pcbnew basically treats neutral (no net) copper as errors and refuses to fill them.
4. pcbnew would not fill a copper zone that was not connected, regardless of to which net it belonged.

Providing support for isolation zones is a non-fabrication design issue. That is, it supports the electrical design specification of the board rather than tailoring of board features to a particular purpose.

(R) 17 (pcbnew) *pcbnew will be enhanced to support isolation zones. Likely the easiest way to accomplish this is to define a special net for "neutral copper."*

3.5.7 Module Zones

Module Zones is an enhancement to both modules and zones. Historically in pcbnew it has not been possible to create a zone as part of a module. This enhancement makes it possible. Zones can be created on copper layers or non-technical layers. Zones created on copper layers are given a pin-number connection for the purposes of assigning net values.

3.5.8 Zone Plotting

Historically, pcbnew has painted all zones with line segments when plotting to Gerber RS-274X and other plot formats. Fabricators expressly prefer copper zones to be created with polygon (contour) fills instead of being painted with line segments. The reason is so that fabricator CAM software can identify the difference between copper zones and traces. One of the first jobs of the CAM software when importing Gerbers or DPFs is to convert all painted pads to flashes and to convert all painted zones to contours. Using a polygon fill approach necessitates the use of clear overlay areas to form cutouts in the polygon shapes. Historically, pcbnew has not used overlays at all. In fact, pcbnew's Gerber plots have been more like the approach taken many decades ago for RS-274D Gerber and uses none of the features of RS-274X.

(R) 18 *This feature addresses changing plot routines to use polygon fill instead of painting zones.*

3.6 Modules

Historical support for modules in pcbnew has be sparse and basic. Many features required by specialized packages are missing from the module editor. Some of the primary features requiring support in the module editor are as follows:

Solder Dams. Some fine-pitch components require solder dams be placed between component contacts. Where the solder mask opening provides an insufficient web, the choices are "gang relief", where solder mask is removed completely from between lands; or, "solder dams", where the solder mask is permitted to overlap slightly with the land so that a minimum web thickness of solder dam can exist between the lands. Solder dams and gang relief are discussed below in *Sec. 3.6.1(42)* and detail in the primary mask layer in *Sec. 5.8(77)*.

IPC Footprint Identification. Most assembler prefer the use of IPC footprints for SMT components, particularly for chip capacitors, chip resistors, BGAs and QFP/QFN. pcbnew needs to identify the IPC footprint name [PCB-2010] and possibly the JEDEC package name so that it can be provided as the package name for pick-and-place files as well as for BOM (parts-lists) for manufacturing outputs. IPC Footprint identification is discussed below under *Sec. 3.6.2(43)*.

IPC Component Types. To be able to apply rules for component-to-component spacing that is dependent upon the package-type relations, it is necessary to identify the package type to which a component belongs. Part of this can be accomplished with the IPC footprint name; however, non-IPC footprints still need to identify the package type. IPC package types will be used to avoid confusing systems of naming. IPC component types are discussed below under *Sec. 3.6.3(43)* and requirements detailed in the new "Courtyard" layer in *Sec. 5.16(97)*.

Component Courtyards. Component courtyards act as a guide for placement of components (both manual and autorouting). Also, component courtyards provide the definition of the manufacturing zone to which DFA rules for wave soldering or reflow can be applied. Courtyards are summarized in *Sec. 3.6.4(43)* and detailed under the new "Courtyard" layer in *Sec. 5.16(97)*.

Rework Spacing. For components on a board that require the ability to be reworked, additional clearances to other components are required to support the ability to de-solder and replace the component. Rework zones are related to component courtyards for placement and wave or hand soldering. Rework spacing is discussed below in *Sec. 3.6.5(44)* and the new "Courtyard" layer is detailed in *Sec. 5.16(97)*.

Wave Soldering Spacing. Wave soldering, selective wave, and miniwave, soldering places additional separation demands on component placement. Two forms of selective wave soldering is possible: in a regular full wave soldering line with the use of conformal solder stencils that protect reflow components from being washed away, or the use of selective wave (dipping) or miniwave (drag) soldering. In general, this soldering is only applied to TH components, and provides the necessary clearances for hand-soldering as well. Wave soldering spacing is discussed below in *Sec. 3.6.6(45)* and the new "Courtyard" layer is detailed in *Sec. 5.16(97)*.

Component Outlines. Component outlines are used to create drawings of the components installed on the exterior sides (or interior layers) of the board. Their purpose is documentation (repair manuals, rework instructions, user manuals, module documentation) and assembly drawings. Outlines are discussed below in *Sec. 3.6.7(47)* and the new "Component" layer is detailed in *Sec. 5.17(104)*.

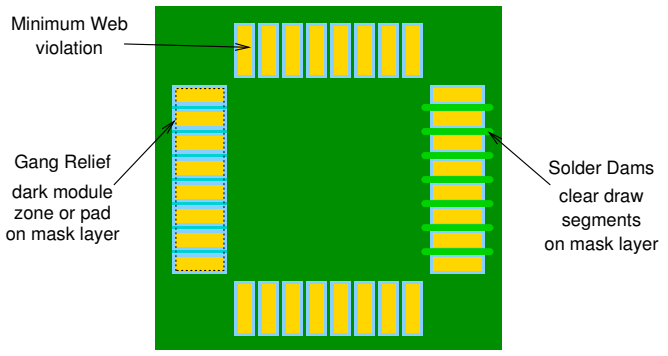

Figure 38: Solder Dams and Gang Relief

Component Keep-outs. Many high-speed or high-frequency components describe areas in which traces or other components should not be positioned. This is not directly related to component placement courtyards, in that they can affect traces and thieving on the exterior or interior layers of the board. Module keep-outs are discussed below in *Sec. 3.6.8(48)* and the new "Keep-out" layer is detailed in *Sec. 5.6(69)*.

Thermal Vias. Some QFN packages and many devices from Texas Instruments require a central grounding pad that is not only electrically connected, but is intended to be thermally connected to the board. Vias filled with a conductive (thermal) filling is required to transport heat from the package into the copper layers and ultimately the dielectric layers of the board. That is, these devices use the board as a heat sink. Thermal vias are discussed below in *Sec. 3.6.11(49)* and detailed under Vias in *Sec. 3.3.9(33)*.

Press-Fit Fixtures. Some press-fit components require a fixture that is placed in contact with the opposite side of the board from the press-fit component. This fixture protects against board distortion from stresses of the press. Press-fit fixtures are discussed below in *Sec. 3.6.12(49)* and the new "Keep-out" and "Fixture" layers are detailed in *Sec. 5.6(69)* and *Sec. 5.22(114)*.

Embedded Resistors. Embedded resistors are resistors that are formed by placing a resistance ply under the copper traces that is etched along with the traces and then have only copper removed with a secondary etching process. Embedded resistors are discussed below in *Sec. 3.6.13(49)* and the image requirements are detailed in the new "Resistance" layer in *Sec. 5.2(61)*.

3.6.1 Solder Dams

Fig. 38(42) illustrates the principles of minimum web, gang relief and solder dams. With no adjustment, the copper lands are so close together that once solder mask registration clearance is applied to the NSMD pads, there remains an insufficient web of solder mask between the pads. The two choices are: "gang relief," where a dark module-zone or pad is defined as NSMD on the mask layer that encloses the pads that need gang relief so that the violating webs are simply removed; or, "solder dams," where clear draw segments are used to enforce a minimum web between the pads. "Solder dams" results in a potential overlap

between the dam and the pads on either side. When the overlap is excessive, the DRC checks for solder mask clearance (between the clear draw segment and the NSMD pad considering allowable overlap) will generate an error when checked. This relies on the new ability to define draw segments that are "clear" instead of "dark" as is normally the case.

Solder dams and gang relief can be left to the fabricator and called out on the fabrication print as expoused by Sanmina-SCI in the following passage: [SS-2007]

> Soldermask is commonly placed between surface mount pads to reduce solder bridging during assembly. These small lines of mask between the pads are called webs. Soldermask webs have a minimum width that can be reliably exposed such that they will remain on the board during assembly. The space between the pads must allow for registration tolerances for imaging the soldermask. This leaves a small "finger" of soldermask web between the SMD pads. The web must be removed when the web width falls below the manufacturing capability. This is called "gang relief." Please specify if gang relief is allowed, or not allowed, on the fabrication print. When spacing falls below the required minimum to place a web of mask between the SMD pads and a web is required, a 0.001" overlap onto the SMD pads will allow a web to be placed if no alternative is possible. This allowance should be noted on the fabrication print. [SS-2007]

A mechanism should be provided to allow the designer to place solder mask dams between SMT pads as per above. "Gang relief" can be provided by creating a dark module-zone or pad on the solder mask layer that extends tangent to the perimeter of the group of pads that are too close together and for which no solder dams will be provided. The module-zone or pad should be described as NMSD (Non-Solder Mask Defined) so that it will be expanded by the solder mask clearance. When the pads cannot be solder mask defined when gang relief is required.

Where "solder dams" are required, a clear draw segment should be drawn between the pads. This draw segment acts as a solder dam. Segments on the solder mask layer will be given the same line width as the minimum web for the soldermask process and material. Because it is rather a guess whether a solder mask dam is even possible in this location for a given solder mask process (minimum web), DRM checks will be made to check the overlap (if any) between the solder dams and pads is acceptable, and generate an error otherwise. Where library modules are defined with solder mask dams supplied with the library module, the copper land should be extended by an equivalent area as the expected overlap on the edges that are not in conflict.

> Soldermask clearances at the fingers should extend past the bevel and the board edge. This will prevent having soldermask between the end of the bevel and the fingers, where it could flake off during manufacturing.
>
> It is preferred to have the solder mask brought down to the top of the fingers, and not to have it stop above the fingers. Stopping above the tops of the fingers will result in circuits exposed from the edge of the solder mask to the fingers. Damage to the exposed circuits may result during the gold top plating operation. A distance of 0.050" should be maintained between the edge of the mask and the edge of the nearest plated through hole or via hole clearance. If maintaining the 0.050" spacing is not feasible, then the mask should be brought down on top of the fingers to allow for the 0.050" strip of mask to be maintained. [SS-2007]

To accomplish the above, gold contacts should be SMD (Solder Mask Defined) and then a solder mask "gang relief" should be applied over the gold finger contacts past and overlapping with the bevel. The bevel should be plotted on the solder mask layer (as dark) so that solder mask will be removed from the bevel. "Gang relief" can be applied by defining a dark module-zone (or pad) on the solder mask layer and extending it to be tangent to the end of the contacts, and extending down into the bevel. The module-zone or pad should also be defined as solder-mask-defined (SMD) so that the mask opening is not adjusted by the solder mask clearance. A clear draw segment can also be used along the top of the contacts to extend the mask down onto the contact.

3.6.2 IPC Footprint Identification

IPC footprint identification an naming for SMT components is described in [PCB-2010]. This identifying name can also be used to decipher the IPC component type as listed in *Tab. 2(45)* and *Tab. 3(45)*.

Many production assemblers demand that IPC land patterns be used. When generating package names on centroid (pick-and-place) files and in BOM (parts lists) manufacturing outputs, the lack of a proper IPC land pattern name can cause the assembly to be quoted high, reviewed, placed on hold, or rejected. To indicate that a particular pcbnew library module corresponds to a given IPC land pattern, the IPC land pattern name in accordance with the IPC land pattern naming convention [PCB-2010] must be provided in the manufacturing outputs. Because the IPC land pattern naming convention is not necessarily the best way to name or organize the module libraries, an IPC name should be separately associated with the module file.

3.6.3 IPC Component Types

The basic SMT IPC component types from [PCB-2010] are listed in *Tab. 1(45)*, *Tab. 2(45)* and *Tab. 3(45)*. To identify type-dependent courtyard relationships between various component types, it is necessary to be able to associate a module with an IPC component type. Type types of relationships are discussed below in *Sec. 3.6.5(44)* and *Sec. 3.6.6(45)*.

(R) 19 (pcbnew) *Library modules will have their IPC component type derived from their IPC component name.*

3.6.4 Courtyards

Historically, pcbnew has not directly addressed placement courtyards. The silk screen outlines were the only thing that could be used by the designer as a guide to placement courtyards, with the problems and issues attendant to trying to make the silk screen outline do triple duty.

As illustrated in *Fig. 39(44)*, basic placement courtyards consist of creating an outline around the mechanical component and its lands. To accommodate component manufacturing and basic placement tolerances, an excess is added to this basic courtyard. This provides the fundamental needs of SMT component placement. Typically and in IPC footprint libraries, the component footprint is adjusted for component manufacturing tolerances (size and dimensions of the terminals and body). Added to a courtyard is an excess that needs to be accounted for that addresses the placement tolerance, including floating of the part on its lands during reflow [Hausherr, 2006]. Wave soldered components are typically glued to the board before mass soldering. Additional details are in *Sec. 3.6.6(45)*.

The original purpose of the placement courtyard was "to provide the PCB designer a guideline for placing land patterns next to each other with enough room to compensate for component tolerances." [Hausherr, 2006] They "do not compensate for assembly machine heads and manufacturing allowances." [Hausherr, 2006] *PCB Libraries, Inc.* maintains that the placement courtyard is used only "as a CAD visual graphic aid for part placement and never post-processed." [Hausherr, 2006] However, this

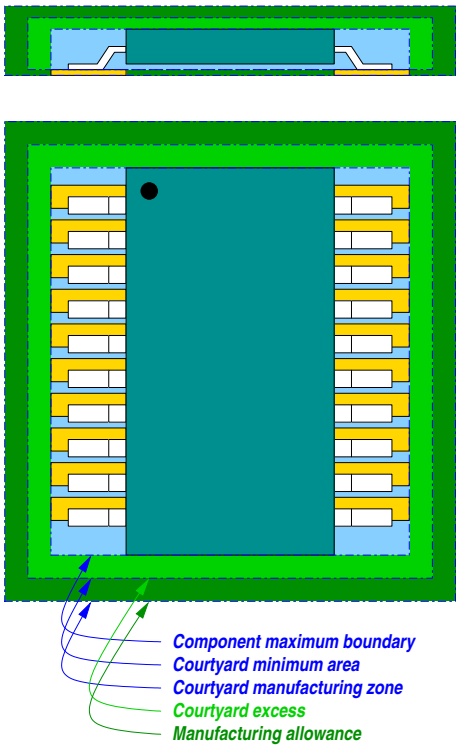

Figure 39: Courtyard Determination

ignores the obvious uses of placement courtyards and a corresponding extrusion height for the purpose of developing selective wave soldering fixtures[14] and for performing general fixture analysis enabled by IDF [Kehmeier and Makowski, 1998] import and export to mechanical CAD systems.

As illustrated in *Fig. 39(44)* [Hausherr, 2006], the IPC libraries calculate the courtyard minimum area from the courtyard maximum boundary. This courtyard includes manufacturing tolerances for the component and some consideration of placement tolerance. This addition is the courtyard excess. The courtyard minimum area is the courtyard that is draw in component libraries and this is the contour that will be used in the pcbnew module editor for the new "Courtyard" layer. IPC courtyards do not, however, include the manufacturing allowance or component manufacturing zone. This requires additional rules. For example, see *Tab. 7(101)* in *Sec. 5.16(97)*.

The courtyard manufacturing zone is critical for assembly; however, IPC libraries neither include the manufacturing allowance nor the courtyard manufacturing zone. Manufacturing allowance is difficult to include in a component library due to the need to support many assembly soldering processes such as full-wave, selective wave fixtures, selective wave and miniwave. [Hausherr, 2006] While these are not choices that need to be included in a component library, they are easy enough to calculate during, or for, placement on the board.

> It's the Courtyard Manufacturing Zone that is critical for the assembly process. This is the body-to-body clearance that you set in your Design Rules for Design Rule Checking. The size of the manufacturing tolerance must come from the assembly shop that is going to be used to populate the parts on the PC Board. Every assembly shop has different assembly tolerances, but the average is 0.1mm (4mil).
>
> The assembly process makes it very difficult to determine placement courtyards for through hole components. It's easy to determine SMT to SMT and even Through Hole to Through Hole, but SMT to Through Hole gets complex, especially when placing Through Hole parts on the top side and SMT parts on the bottom side. Since Through Hole parts require holes that go all the way through the PC Board, the Through Hole part Top Side courtyard would be different than the bottom side courtyard due to the wave solder process used to solder Through Hole component leads. If wave solder is used for the Through Hole component leads, the SMT parts mounted on the bottom side must have a 5mm (200mil) clearance between the edge of the Through Hole pad and edge of the SMT pad. If a selective wave is used there is a different tolerance between the pads depending on the assembly shop requirements.
>
> Therefore, building in placement courtyards for through hole parts is almost impossible due to too many variables. A PCB designer must use common standard rules provided by the assembly shop when performing part placement. The assembly shop should always approve the part placement prior to routing any traces on the board. This is an official check point that must not be avoided. [Hausherr, 2006]

3.6.5 Rework

The manufacturing allowance illustrated in *Fig. 39(44)* can also include rework clearances. Rework clearances are the clearances surrounding a component's courtyard that provides access for rework desoldering and soldering. It provides access for the tools required to heat the joints of the component for soldering and desoldering; and, to grasp the component for removal and reinsertion. When components are intended on being reworked, these additional clearances apply. pcbnew has historically given the designer no assistance in this category.[15] This feature provides assistance for rework.

Rework calculations in KiCad requires the following steps:

1. Calculations for rework clearances begins with the component courtyard. Component courtyards and provided by the courtyard layer. Component courtyards are mapped onto a proximity map for fast clearance calculations.
2. The component courtyard is expanded by a component-to-component dependent clearance. To address all components at once, the component courtyard is expanded by the maximum component-to-component clearance and used to index the proximity map and determine which components have courtyards that overlap with the maximum clearance zone.
3. When the close components are found, the component being checked has its courtyard contour expanded by the component-to-component dependent clearance. The tested component does not have its contour expanded.
4. The kbool engine is then used to determine whether there is an intersection between the target component expanded courtyard contour and the tested component's unexpanded courtyard contour. If there is an intersection, the centre of the intersection is the point at which a DRC marker can be placed.
5. When moving a component, the same calculation is performed and the component can resist movement to or placement at wherever a DRC violation occurs, by splitting the cursor point and component position.

14. See *Sec. 3.6.6(45)*.

15. Making it a CUD (Computer Unassisted Design) instead of a CAD (Computer Assisted Design) tool in this regard.

Table 1: Component Types

Type	Description
Chip	Chip capacitor, resistor or inductor.
SOIC	Small Outline Integrated Circuit
Tantalum	Tantalum capacitor
SOT23	Small Outline Triple lead package
DIP	Dual In-line Package
PLCC	Plastic Leadless Chip Carrier
QFN	Quad Flat No-lead Package
QFP	Quad Flat Package
CSP	Chip-Scale Package
BGA	Ball Grid Array

Table 2: Decipherable Component Types

Regex	Description
/BGA/	A BGA device.
/CAPC/	A chip capacitor.
/CAPT/	A tantalum capacitor.
/CFP/	Ceramic flat package.
/CSP/	A chip scale package.
/DIP/	A dual-in-line package.
/INDC/	A chip inductor.
/PLCC/	A plastic leadless chip carrier.
/T?S?C?QFN/	A quad flat no-lead.
/T?S?C?QFP/	A quad flat package.
/RESC/	A chip resistor.
/SOIC/	A small outline package.
/SOJ/	A small outline J-leaded package.
/P?SON/	A small outline no-lead.
/T?S?SOP/	A small outline package.
/SOT/	A small outline transistor.

6. When moving a component and displaying clearance lines, clearance lines can be drawn for all components within the proximity of the target position, at the clearance to the moved component.

To determine the component-to-component clearance requires an attribute of the module that specifies what type of component it is. A designer-editable matrix is maintained that provides rework clearance values between each pairing of component types. An example of a component-to-component clearance matrix is *Tab. 7(101)*. Note that through-hole or edge-mounted components can have a different front and back side clearance. A minimal set of component types is listed in *Tab. 1(45)*.

In general, when the footprint names of modules are consistent with IPC-7351, IPC-7351A or IPC-7351B SMT land pattern naming conventions, it is possible to discern the component/device type from the name. The regular expressions for matching the packages of *Tab. 1(45)* are listed in *Tab. 2(45)*. In the event that the IPC-7351x name is used for the footprint, or when a footprint field containing the IPC name is included in the module, it is possible to determine the package type from the name. Note also that, as of IPC-7351B, the height of the package is also codified in the land name even though the IPC land information (courtyard) does not include height.

3.6.6 Wave Soldering

Most modern design are either SMT only, have TH assembled with Intrusive Soldering (a technique where solder paste is deposited in the hole of TH components and soldered using no-clean reflow), or are assembled using selective wave. Full wave soldering is not normally encountered. Some of the placement requirements for wave soldering and selective wave soldering are

Table 3: IPC-2511A Device Types

Code	Description
RES	Two terminal resistor.
VRES	Variable resistor.
FABRES	Embedded resistor.
RPCK	Resistor pack.
DPCK	Diode pack.
LEDPCK	LED pack.
HYBRID	AtoD or DtoA.
CAP	Two terminal capacitor.
VCAP	Variable capacitor.
PCAP	Polarized capacitor.
TCAP	Tantalum capacitor.
FABCAP	Embedded capacitor.
CPCK	Capacitor pack.
IND	Inductor.
VIND	Variable inductor.
XFMR	Transformer.
DIODE	Diodes, including Schottky.
DIAC	Diac.
ZENER	Zener diode.
BRIDGE	Silicon bridge rectifier.
PNP	Transistor, unijunction or darlington.
NPN	Transistor, unijunction or darlington.
NFET	FET family.
PFET	FET family.
NJFET	FET family.
PJFET	FET family.
TRIAC	Triac.
SCR	Silicon controlled rectifier (Thyristor).
VR	Voltage regulator.
OPTO	Opto-isolators.
LED	Light Emitting Diode.
OPAMP	Operational Amplifier.
XTAL	Crystal.
RELAY	Relays.
SWITCH	Switches.
FUSE	Fuse.
JUMPER	Jumper.
CONN	Connector.
SOCKET	Socket.
LOGIC	Logic devices.
ANALOG	Analog IC.
OTHER	Used when none of the standard types is sufficient.

given in *Sec. 5.16(97)*.

In general, the requirements for determining component-to-component clearances for wave soldering are similar to the calculation of rework clearances (see *Sec. 3.6.5(44)*) in that component-to-component clearances can be maintained in matrix form for each pairing of component types. Component types can be determined by any of the methods described in *Sec. 3.6.5(44)*. There are, however, some additional requirements:

1. Certain board features, such as large NPTH, or cutouts, can cause flooding of the component side when wave soldering, requiring either that these holes be plugged when mass soldering (e.g. with peel-off mask), or protected by a selective wave fixture, or larger clearances provided from the hole to nearby pads or vias.
2. Selective

Conformal Selective Wave Solder Carriers – Design Rules: The design of an exclusion fixture is "non-trivial." As well as checking the process suitability, calculations need to be done to ensure

- carrier stability and longevity avoiding Euler Buckling of the masking due to differential heating;
- minimal thermal deflection caused by the thermal gradient across the carrier;
- optimal solder flow, maximum sculpting locally to the PTH while maintaining overall stiffness.

These calculations are only aids: successful designs evolve with experience.

Process suitability: The viability of this process route crucially relies upon being able to estimate the number and distribution of PTH components that:

- cannot be processed (with and without glue dotting): these will have to be miniwaved or hand soldered;
- are likely to be poorly soldered: these will require further inspection and manual touch-up.

The estimation can be done in three ways:

1. If a PCB is available (preferably populated), sales engineers can rapidly evaluate the board.
2. If PCB design data is available it can be processed, analyzed and remotely assessed.
3. The designer can use the design rules presented below.

Pin Land to SMT pad clearance evaluation: The TH land to SMT pad clearance, edge-to-edge, is dependent upon the direction of travel of the wave as well as the height of the SMT component (that must be covered).

Less clearance (2.2mm to 2.8mm, slightly dependent on height) is required when the perpendicular distance is perpendicular to the direction of the palet travel through the wave. An approximation of the function is (all units are millimetres):

$$\begin{array}{l|l} s = \text{undefined} & h = 0.0 \rightarrow 0.5 \\ s = 2.2 & h = 0.5 \rightarrow 2.5 \\ s = 0.4(h - 2.5) + 2.2 & h = 2.5 \rightarrow 4.0 \end{array} \quad (47)$$

More clearance (4.0mm to 8.0mm, strongly dependent on height) is required when the perpendicular distance is parallel to the direction of the pallet travel through the wave. Additional clearance is required in this case as a function of the height of the SMT component being protected by the fixture. An approximation of the function is (all units are millimetres):

$$\begin{array}{l|l} s = \text{undefined} & h = 0.0 \rightarrow 0.5 \\ s = 0.95(h - 0.5) + 3.9 & h = 0.5 \rightarrow 2.5 \\ s = \frac{3.8}{3}(h - 2.5) + 5.9 & h = 2.5 \rightarrow 4.0 \end{array} \quad (48)$$

Table 4: Data requirements for selective wave fixtures

Component Side (Top):	
Solder Paste	To assist with determining where not to place board fixings.
Solder Mask	To assist with determining where not to place board fixings.
Silk Screen	To identify what is to be soldered.
Solder Side (Bottom):	
Silk Screen	Not needed when there are no TH components install from the bottom.
Solder Paste	To determine what needs to be reflowed.
Solder Mask	To assist with determining where not to place material, and to avoid leaving large icicles on ground planes.
Drill Data	To identify where plated and non-plated through holes are located.
Placement Data:	
	For overlays.

General design rules: Additional general design rules for selective wave fixtures:

- Keep large (height) SMT components away from PTH areas.
- Leave the leading and trailing areas around PTH as clear as possible.
- Do not put any SMT components within 3mm (120mil) of any PTH components.
- Do not put all PTH components in line along on edge of the board, leave space to support the masking in the center.

Data required: The normal board data is Gerber and ideally RS-274X. Many carrier manufacturers accept a wide range of formats. Data should include:

To realize optimum soldering results, the particular process conditions have to be met. To a high degree, process reliability depends upon:

- the design of the pads (type, distance between each other);
- the distance between pad and neighbouring component pad (e.g. SMD which is not supposed to be touched);
- the pin length underneath the board;
- the pitch between pins (e.g. pitch of connectors).

These influence coefficients effect directly the peel-off of the flowing solder. To ensure no-bridging results, a defined and reproducible peel-off is required. Bridges are the main soldering failure (more than 80%).

Basically, it has to be distinguished between drag soldering process with a miniwave and dip soldering process. Each process demands a special PWB design.

The following design guidelines ensure the best possible process conditions. In case you should decide not to follow these recommendations, the all-over process window will be smaller and additional steps will be required to stabilize the process. These additional steps could demand higher maintenance requirements and an increased wear of tooling parts.

Optimum Layout for Dip Soldering Processes: The clearance between solder pads:

- round pads are preferable;
- distance between round pads (edge-to-edge) is $\geq$0.60mm;
- pin distance (centre-to-centre) is $\geq$2.54mm.

Pin length underneath the board:

- pin length extension under the board should be $\geq$2.5mm.

Clearance of solder nozzles: the distance to neighbouring pad (not to be soldered):

- on three sides of a rectangle enclosing the pins to be soldered, the clearance is ≥3.0mm;
- on the fourth side, the clearance is ≥5.0mm.

Optimum Layout for Miniwave/Drag Soldering Processes The clearance between solder pads:

- round pads are preferable;
- distance between round pads (edge-to-edge) ≥0.6mm;
- pin distance (centre-to-centre) ≥1.9mm.

Pin length underneath the board:

- pin length extension under the board should be 0.8–2.0mm.

Clearance of the miniwave: the distance to neighbouring pad (not to be soldered):

- on three sides of a rectangle enclosing the pins to be soldered, the clearance is ≥2.0mm;
- on the fourth side, the clearance is ≥5.0mm.

Smallest Dimension of Solder Nozzle: A rectangular solder nozzle has dimensions of 8.0mm × 5.0mm, with a soldering area ≤ 40 mm^2. A round solder nozzle has a diameter of 3.0mm, with a soldering area ≤ 7 mm^2.

Maximum Height of Neighbouring Components: On the bottom side (soldering side) the maximum component height is limited through the height of the solder nozzle. The standard solder nozzle height is 32mm. Therefore, the maximum component height should not exceed 25mm. Higher components demand a higher solder nozzle design.

Furthermore, it is necessary to take notice of the distance between any component and the solder joint. In case of drag soldering process the component could touch the nitrogen shroud, e.g. if the component is higher than 10mm and the process is carried out with a solder angle. Rule of thumb when the angle is employed and the component is more than 10mm high:

$$\text{height of component (mm)} \leq \text{distance to the solder joint (mm)} \tag{49}$$

Solder balling is a phenomenon in all wave soldering processes that always occurred in the past and which will occur in the future as well. It, however, appears more frequent in lead-free soldering processes and process temperatures are remarkably higher than in traditional soldering processes. The higher process temperatures can have a negative effect on the solder resist. Depending on the quality, the solder resist might soften during preheating which abets arising solder balls to stick at the solder resist. In traditional lead bearing processes or applications featuring high quality lead-free solder resists, arising solder balls would bounce off. *Therefore, if possible, solder resist close to the solder joint should be avoided.*

Particularly in multi-nozzle dip soldering processes, special nozzle designs can help to avoid solder balling as well. These nozzle tools are featured with a defined solder flow which is directed by means of a flow plate. In addition, the complete nozzle tool is covered with a second top plate. Any splashes, which might occur while the liquid solder is flowing back to the reservoir, therefore will not get a chance to touch the printed circuit board.

The phenomenon of poor hole fill is mostly based on an insufficient heat transfer rate that also can be improved with an appropriate PCB layout.

The length of leads plays an important role in this regard, particularly in multi-nozzle dip soldering processes. Multi-nozzle dip soldering processes require a lead length greater than 2.5mm. This is related with the energy transfer rate that directly affects hole penetration. Longer component leads are dipped deeper into the liquid solder that improves the heat transfer that finally results in an improved hole fill.

Another issue that should be considered in respect to hole fill is an ideal ratio between the pin diameter and the via. If this ratio should be too large, no capillary action will emerge. Should this ratio be too small, flux cannot soar through the via and therefore solder joints cannot be formed properly. *As a rule of thumb, the diameter of the via should be equal to the diameter of the pin plus 0.2 up to 0.4mm. Lead Free process even can require a plus of 0.5mm.*

Thermal energy also will be transferred better when the pad size is enlarged to a certain extent or if oval pads are used. *If possible, solder resist close to the solder joint should be avoided. This helps to keep the heat at the solder pad and in addition also helps to avoid solder balling.*

Attention should be given also to thermal decoupling. *With an appropriate thermal decoupling of the PCB, the heat will not be completely withdrawn to the strip conductor, but will be held for a longer time at the pad.*

Flowing solder waves, also in dip soldering process, should generally be preferred. This ensures that oxide-free and correctly heated solder alloy is continuously supplied to solder joints. Even during the contact phase, the solder alloy does not cool down. This improves fill remarkably, even in case of high-mass pins, at pins with connection to inner layers or pins that are located at the outer edges of an assembly.

Most frequently a selective soldering process cannot be realized because of missing clearance between the solder joint and neighbouring components, such as:

- SMT devices that might be washed off during the process; or,
- housings of other leaded components that could be touched and damaged by the solder nozzle.

In may other cases, solder bridges and poor hole fill are the main reasons for faults. In addition, solder balls can cause difficulties. The solder's pull-off behaviour, which is influenced by several factors, is what is mainly responsible for reliable soldering results in the selective soldering process.

In general, one has to distinguish between the different selective wave soldering processes.

Selective soldering as a single miniwave process can be performed in either a drag or a dip soldering mode and allows soldering with an angle. This offers high flexibility and fewer restrictions with regard to board design; however, depending on the number of joints to be soldered, single miniwave processes show a longer cycle time. Typical cycle times range between 1 and 10 minutes.

Multi-nozzle dip soldering process, on the other hand, use product-specific solder nozzle tools that results in a certain inflexibility As all solder joints of an assembly, however, are processed simultaneously, multi-nozzle dip soldering processes are featured with a short cycle time that ranges between 20 and 30 seconds. Most machine systems do not feature soldering with an angle.

3.6.7 Outlines

Historically, pcbnew has made the silk screen layer serve multiple purposes: legend, placement courtyard for layout, component outlines for assembly drawings. This is insufficient. A drawing of the component outline is necessary for detailed assembly drawings that can be used in documentation (repair manuals, rework drawings, user manuals). Viewing the component outline can also assist in layout where courtyards are being violated, or simply to assist in the visualization of the components on the board.

Assembly shops tend to use the pad master and silk where it is available to verify centroid files and IPC-D-356 net lists. However, a Gerber plot of the assembly layers is also common place amongst other CAD tools in the industry. Therefore a "Component" layer was added to contain component outlines and to allow them to be plotted to assembly drawings. See the "Component" layer in *Sec. 5.17(104)* for more detail on component outlines.

Previously in pcbnew I used to draw the component outline in the "Drawings" layer in the module editor. This provided nice grey component outlines that were not too obtrusive when laying out the board. The "Drawings" layer should actually be reserved for the main fabrication print, so when reading a module, the "Drawings" layer should be mapped to the "Component" layer when reading a version 1 module file.

(R) 20 (pcbnew) *The "Drawings" layer from version 1 modules will be mapped to the new "Component" (outline) layer when reading version 1 library modules into pcbnew.*

3.6.8 Keep-Outs

In a number of high-speed component definitions, keep-out areas where no other traces, components or vias can be placed. This is in addition to the component courtyard (which only describes the above-board courtyard). Keep-outs can be used to affect the external or inner copper layers as well. Keep-outs can be described to keep out thieving, venting, traces, vias, from external or internal layers.

The instance that I have run into is the use of SFP (Small Formfactor Pluggable) optical module cages. These are high-speed (10Ghz) component cages that must not have EMC sources beneath them or they will not function very well. Also, traces are only allowed to enter the cage on external layers through narrow windows on the front of the cage. To model this, a keep-out zone can be described that indicates that theiving and traces are to be kept out of the specified areas.

To accommodate these module keep-outs, the new "Keep-out" layer was defined. Creating zones (closed contour polycurves) on the "Keep-out" layer in the pcbnew module editor will provide the necessary ability. See the new "Keep-out" layer in *Sec. 5.6(69)* for for more detail on keep-out layers.

3.6.9 Rooms

For additional information on "rooms," see the new "Courtyard" layer in *Sec. 5.16(97)*.

3.6.10 Thermal Calculations

Components are typically designed to dissipate their heat into the board, $\theta_{JC} \gg \theta_{JB}$, or into the ambient (case), $\theta_{JC} \ll \theta_{JB}$. Some chip-scale QFNs lack the surface area for ambient (case) dissipation or proper mounting of heat sinks, and so they dissipate their heat into the board through a large thermal grounding pad in the centre of the package. Most large BGAs packages or other flip-chip technologies are designed to dissipate heat into a metal hermetic cover that dissipates to the ambient or connects a heat sink through thermal adhesive. Combinations of the two are possible. The general model is:

$$T_J = T_A + P_A\theta_{JA} \tag{50}$$

$$T_J = T_B + P_B\theta_{JB} \tag{51}$$

$$\theta_{JA} = \theta_{JC} + \theta_{CS} + \theta_{SA} \tag{52}$$

$$P = P_A + P_B \tag{53}$$

$$T_J = T_A\frac{\theta_{JB}}{\theta_{JA}+\theta_{JB}} + T_B\frac{\theta_{JA}}{\theta_{JA}+\theta_{JB}} + P\frac{\theta_{JA}\theta_{JB}}{\theta_{JA}+\theta_{JB}} \tag{54}$$

Where P is the power dissipated at the junction; T_J is the temperature of the junction; T_A is the temperature of the ambient; T_B is the temperature of the board; θ_{JA} is the thermal resistance between the junction and the ambient; θ_{JB} is the thermal resistance between the junction and the board. When θ_{JA} or θ_{JB} are controlling, *Eqn. 54(48)* reduces to *Eqn. 50(48)* or *Eqn. 51(48)*, respectively. Therefore, when one of θ_{JA} or θ_{JB} is not specified, it should be assumed to be infinite (or at least much larger than the specified value).

Now, the temperature of a locality of the board, T_{B_L}, is determined by the amount of power dissipated by devices into the locality, the amount of power dissipated by the conductors carrying current within the locality, the temperature of the ambient, T_{A_L}, on the outer surfaces of the locality, the temperature of the board, T_B, surrounding the locality, and the shape of the locality. Full 3D heat-flow analysis can of course be performing using heat-flow equations and the geometry and characteristics of all of the material involved. However, usable approximations assign a general value to T_B and T_A. T_A is usually the maximum inlet temperature for forced airflow, and the airflow is set to dissipate the maximum wattage of the PWA. So, for example, for PCI Express, the maximum wattage for the form factor can be 25W and the maximum inlet temperature is 55°C, so $T_A = 55°\text{C}$ would be used. Of course it would be advantageous to have critical junctions upwind.

The temperature of the board is generally approximated using the power dissipated by and into the board, and the thermal resistance of the board to the ambient airflow. Usually the only significant airflow is along the bottom of the board (where there are few and short components), so the thermal resistance of only one board surface to the air should be considered. In that case, the temperature of the board can be estimated as $T_B = T_A + P_B\theta_{BA}$, where P_B here is the power dissipated into and by the board, T_A is the temperature of the airflow, and θ_{BA} is the thermal resistance between the board and the ambient and depends on the material used to construct the board and the area of the board surface exposed to the airflow. Of course this does not consider localized effects where the temperature of the board may be locally higher (hot spot) due to localized effects (heat diffusion within the board). Typically for maximum junction temperatures, the temperature of the board should be considered at operating expectations and local maximums. The board itself cannot handle temperatures higher than the T_g of the dielectric materials involved.

While it is true that restricting a "room" to a given maximum power dissipation is somewhat useful for controlling hot spots within a localized area, the maximum power dissipation should be separately specified for power dissipated into the ambient and power dissipated into the board. Again, the maximum power density would be a better control of localized hot spots than just the maximum dissipation for the entire room. Therefore SPEECTRA DSN does not support thermal analysis for placement well, whereas IDFv3.0 does, because it does not restrict in any way the candidacy rules for placement groups.

IDFv3.0 [Kehmeier, 1996] provides component properties: capacitance in μF, resistance in Ω, tolerance in % deviation, operating power rating in mW, maximum power rating in mW, thermal conductivity in W/m°C, junction to board thermal resistance, θ_{JB}, in °C/W, junction to case thermal resistance, θ_{JC}, in °C/W, and other user-defined values. The interesting thing here is the specification of θ_{JB} and θ_{JC} which can be used to determine which model (into the board, into the case) the component follows.

IDFv4.0 [Kehmeier and Makowski, 1998] provides a more complex thermal model composed of a network of thermal capacitances, resistances and areas. The model can, however, be simplified to the model of IDFv3.0.

SPEECTRA DSN does not support thermal modelling, only

a single power-dissipation assignment for components and the ability to specify the maximum power dissipation of a "room".

(R) 21 (pcbnew) *Support will be provided for the "rooms" models of SPEECTRA DSN and the more general-purpose model of IDFv3.0.*

For additional information on "rooms" based on power-dissipation and thermal considerations, see the new "Courtyard" layer in *Sec. 5.16(97)*.

3.6.11 Thermal Vias

Many QFN packages have a large ground pad in the center of the package that is intended on providing an electrical ground as well as for dissipating heat from the chip into the board. These thermal grounding pads are supposed to have thermal vias (vias that are filled with a thermally conductive material) to facilitate heat transfer. pcbnew has historically not been able to describe these thermal vias adequately.

(R) 22 (pcbnew) *A new pin attribute will be assigned that identifies pads created inside this large grounding pad as thermal vias so that they can be treated separately for hole filling.*

Thermal vias are described in more detail in *Sec. 3.3.9(33)*.

3.6.12 Press-Fit Fixtures

Press-fit components that have a significant insertion force (due to the press-fit pin design or friction of surface finish) require a fixture for the back-side of the board to protect the board from distortions due to the insertion force of the press-fit component. The demands on the CAD system of such press-fit fixtures have not historically been handled by pcbnew. There are several approaches that need to be supported:

1. Mechanical definition with IDF.
2. Mechanical definitions within the CAD system.

In general the press-fit fixture can be described with mechanical and electrical keep-outs. Where the press-fit fixture makes contact with the board, it is necessary to keep out components mountings on the opposite side of the board to the press-fit component. Depending on the insertion force, it may also be necessary to keep traces or other exterior copper features away from the areas where the fixture contacts the board. These areas can be defined using the new keep-out layer as a set of component and/or copper feature keep-outs. Cutouts can be used to define areas in which the pin traverses the board and extends out the opposite side of the board through the press-fit holes. These keep-outs should already be capable of being exported to or imported from IDF. An attribute should be provided to classify the keep-outs as belonging to a press-fit fixture.

Another approach is to use the new "Fixture" layer to describe the fixtures in more detail. A set of closed contours with extrusion heights defined for the reverse side of the board is likely a better alternative that the "keep-out" layer, which does not really intended to have a height associated with it. In practise, however, either approach could be used.

For more details on the "Keep-out" layer, see *Sec. 5.6(69)*; for the "Fixture" layer, see *Sec. 5.22(114)*.

3.6.13 Embedded Resistors

Embedded resistors using resistance layers have not historically been supported by pcbnew. Embedded resistors are formed by laminating an etchable resistance layer between the copper cladding and the dielectric on either side of a core. Both the copper and resistance layers are etched using the regular copper etching artwork, and then an additional photoimaged etching process

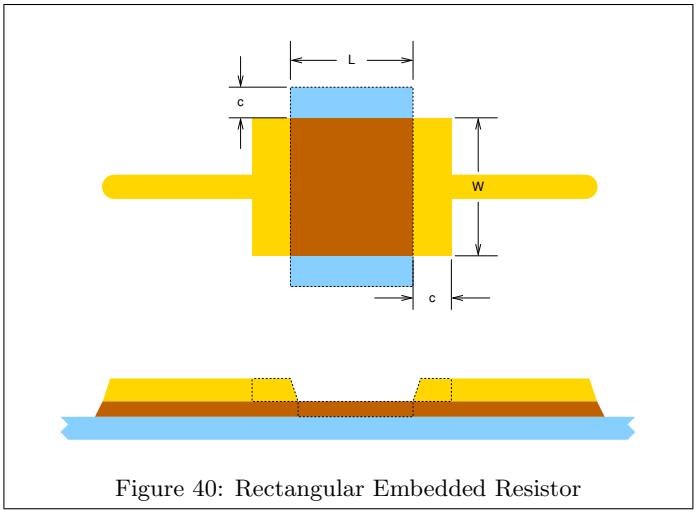

Figure 40: Rectangular Embedded Resistor

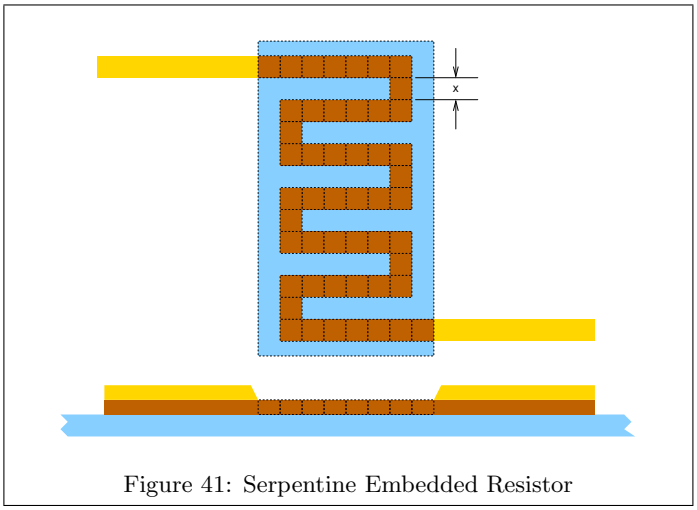

Figure 41: Serpentine Embedded Resistor

is applied to remove only the copper from areas, leaving only the resistance layer intact. There are two common ways of forming embedded resistors with this approach: a rectangular area and a serpentine.

A rectangular area embedded resistor is illustrated in *Fig. 40(49)*. The width of the resistance area is controlled with the copper etch mask by controlling parameter W, the width of the rectangular area. An overlap, C, is used to form terminals, so the rectangular copper etch mask has a length, $L_c = L + 2C$. The length of the resistance area is controlled with the resistance etch mask by controlling parameter L. An overlap on the resistance etch mask is also formed so that $W_r = W + 2C$. The resistance created by the resistance area is $R = \rho A = \rho(L \times W)$, where ρ is the resistivity of the resistance layer in $\Omega/\square$ and is dependent upon the volumetric resistivity of the material used in the resistance layer, ρ_v, and the thickness of the resistance layer, t, $\rho = \rho_v \times t$.

A serpentine embedded resistor is illustrated in *Fig. 41(49)*. A illustrated, the resistance of the serpentine can be broken into squares to determine the resistance of the overall structure. For squares of $x \times x$, the resistance for each square can be determined as $r = \rho x^2$, where ρ is as before. Each square laid out in a linear path contributes an additional r to the overall resistance (series). The squares at the corners do not contribute an complete r but contribute approximately $0.599r$ instead. So the total resistance is $R = mr + 0.599nr = (m + 0.599n)\rho x^2$, where m is the number

of linear squares, and n is the number of corners. In *Fig. 41(49)*, $m = 38$ and $n = 12$, so the resistance is $R = (38 + 0.599 \times 12)\rho x^2$. If $x = 1$ and $\rho = 1\Omega/\square$, $R = 45.188\Omega$.

As the process is a photo image mask process, other geometries are possible, subject only to the creativity of the designer.

There are several problems with embedded resistors: the heat dissipation of the resistor is trapped within the board. The thermal conductivity of the surrounding dielectric is poor, so the only path for heat to escape is along the copper traces or the sandwiched resistance layers. Copper zones, thermal vias, and PCB heat sinks are required to dissipate large amounts of heat. Therefore, this technique is really most useful for low voltage embedded resistors. However, terminating resistors for low voltage signal families are an excellent candidate for this approach for two reasons: termination resistors for low voltage signal families are small and do not dissipate much heat; and, termination resistors cause problems in layouts alongside large BGA devices with many terminations, such as FPGAs and ASIC logic chips.

Embedded resistors cause problems for traditional PCB CAD systems in general, and specifically for pcbnew. Traditional devices are only placed on the external surfaces of the board, so convincing pcbnew that a module can be placed on an internal layer of the board is a challenge. Also, it is not usual to make every internal layer of a multilayer PCB a resistance layer. The resistance layer under the copper layer can have negative skin effect ramifications for very high speed signals. For lower speed signals this is not a problem. Therefore it is advisable to restrict the layers to which the technique is applied. Also, the resistance layer adds thickness to the board, but still requires the same z-axis clearance because it is essentially a conductor. So, applying resistance layers to each layer of the board can have ramification on both high-speed signals and overall board thickness. Therefore, resistance layers must be considered in the stack-up.

So, to support embedded resistors, pcbnew must be taught the following:

1. That resistance layers can be sandwiched with copper layers in the stack-up.
2. The thickness and resistivity of the resistance layer.
3. That embedded resistors can only occur on layers that have a sandwiched resistance layer.
4. That modules that represent embedded devices can occur on any external layer in the stack-up (core construction only) or internal layer in the stack-up (foil construction), but only on layers that have a defined resistance layer.
5. To initially place embedded resistance modules on valid layers.
6. How to autoplace embedded resistors.
7. How to calculate the resulting resistance.

(R) 23 (pcbnew) *To provide support for embedded resistors, the module must be enhanced to identify embedded resistors (so pcbnew know where they can be placed). The pcbnew module editor must be capable of laying out resistors to formulate a user, project, or KiCad library of embedded resistors. Clearance rules must be added for embedded resistors. pcbnew module editing must support placement of embedded resistor modules on internal layers.*

Following are some thoughts about implementing embedded resistors in pcbnew:

- Add a module type for embedded resistors.
- Add the *Resistance* layer for photo images of the secondary etching masks necessary to form embedded resistors. Make this layer visible to the module editor.
- Possibly treat the existing "front" copper layer in the module as the true copper layer and the existing "back" copper layer as the resistance layer (only when the module is marked as an embedded resistor.)

For more information on the "Resistance" layer that is used to image the etch masks necessary for etching away the copper from the resistance ply, see *Sec. 5.2(61)*.

3.7 Nets

3.8 Layer

There are a number of copper layer operations that deserve separate treatment as follows:

Thieving: Plating thieves, or thieving is the process of adding neutral copper to areas isolated from traces and planes to better distribute plating anode current for resolution of the board's exterior layers. It is a layer operation rather than a zone operation, although zones could be used to define thieving areas. Thieving is considered in *Sec. 3.8.1(50)*.

Venting: Resin venting, or simply venting, is the process of adding neutral copper to areas of low pressure during lamination, to equalize the press-out of the prepreg layers in multilayer board construction. It is a layer operation rather than a zone operation, although zones could be used to define venting areas. Venting is considered in *Sec. 3.8.2(52)*.

Crosshatch: Crosshatch is the process of removing copper from large planes for the purpose of equalizing plating current or prepreg press-out in a manner similar to thieving or venting, but also to change the thermal expansion characteristics of a copper plane to avoid bow and twist of the board during construction or operation. It is an layer operation, but is typically applied only to copper zones. Crosshatch is considered in *Sec. 3.8.3(53)*.

3.8.1 Thieving

Plating Thieving is the process of adding neutral copper zones or grid patterns to an external layer of a single or multilayer PCB in an effort to balance the local concentrations of active plating chemistry and protect against over-plating in some areas and under-plating in others. Plating thieving might not be necessary for all plating processes.

Outer layer copper features are created by imaging a base pattern and plating copper to a specified thickness. The pattern plating process is very dependent on the uniformity of the copper features on the outer layers Areas that have a very small amount of copper (such as isolated single traces) will plate a slightly thicker amount of copper on those features. Non-functional thieving pads can be added to low-density areas to balance the copper distribution. This action is called thieving. Thieving should be incorporated during the PCB layout to ensure identical coverage from all suppliers. Board fabricator CAM systems are also typically capable of performing automatic thieving; however, there are a number of strong reasons for the CAD system to perform this function:

- The more actions a fabricator performs, the more locked-in that fabricator becomes.
- DIY shops need to do what the fabricator does, and do not wish to purchase expensive CAM systems.
- Placement of thieving on high-speed and high-density card could be disruptive to the design. Specifying thieving keep-outs and keep-ins might be more difficult that simply performing the thieving with the CAD system.
- Thieving is a simple thing to do. It is not rocket science.

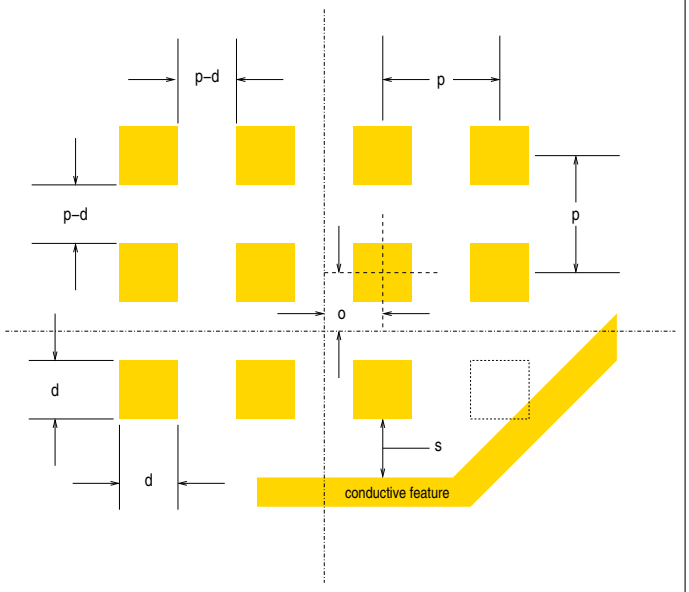

Figure 42: Thieving Pattern

Thieving consists of an orthogonal grid of squares or circles of size, d, set on a pitch, p, resulting in a pattern separation of $p - d$. The clearance between any element of the pattern and any conductive feature is typically $s = p - d$ or $s = p/2$. Some fabricators recommend $d = 30$mil, $p = 50$mil, $s = 20$mil. Others recommend a coarser pattern at $d = 80$mil, $p = 100$mil, $s = 50$mil.

For example:

> Typical thieving consists of a 50mil pitched orthogonal grid of 30mil squares. Clearance from any square to any board feature is 20mil (the same as the inter-square spacing). If thieving is applied to all open spaces without consideration for components, the fabricator will remove thieving from inside SMD arrays on request.

Also:

> Thieving patters are either square .080" non-functional pads on 0.100" centers, or round 0.080" non-functional pads on 0.100" centers. Thieving is a term used to describe the use of non-functional copper pads to balance the copper area of any particular layer. Thieving is intended to avoid isolated copper features. Isolated features tend to plate faster than other features and can even burn during the plating process. Thieving patterns will prevent the over-plating of these features and provide a more uniform copper surface. Thieving patterns are typically 0.080" round or square pads on 0.100" centers. Located closer that 0.050" from any conductive area and 0.050" from the board edge. On panelized arrays, all breakaway areas should have thieving patterns.

And:

> Thieving or non-functional copper used in sparse areas of internal layers is also beneficial to manufacturability. Designs that are not well balanced in respect to copper distribution, many times introduce warpage issues. Adding copper to internal layer also gives that layer more dimensional stability. This is more so with layers having thin dielectrics (less than 6mil). High layer count PCBs benefit greatly when signal layers are paired with a plane layer.
>
> Adding a non-functional plane around signals can prevent the funnel plating or any other over-plating of isolated features. Over-plated features that require a soldermask coating can introduce air entrapment in the soldermask. Mushrooming of over-plated traces can cause shorts.
>
> Plating thieving of low pattern density and cross hatching of external plane area should be considered.

Thieving Process The challenges in applying thieving patterns are as follows:

- *Placing the pattern in the necessary areas.* This has been accomplished by most inferior tools by simply filling all areas of the board. This, of course, causes problems with respect to the other items that follow. pcbnew will use the new keep-out/in layer to specify keep-ins for thieving. When no keep-in is specified for thieving, the entire board surface is considered eligible for thieving. Thieving keep-ins can be plotted on the fabrication print or detail for situations where the fabricator is to apply thieving.
- *Keeping thieving elements the appropriate distance from conductive features.* This can directly be accomplished by removing any elements that fail DRC checks to other elements. This will be accomplished in pcbnew by specifying the separation between the element and conductive features. This can be considered as a copper clearance value for the thieving element. An element that would violate DRC is simply removed.
- *Keeping thieving elements out from under SMT components.* This is obviously more complicated because some CAD tools have historically not performed this function. Fabricators will remove thieving from under SMT arrays when requested on the fabrication print. pcbnew will use the courtyard layer to define thieving keep-outs for components. Thieving keep-outs can be plotted on the fabrication print or detail for situations where the fabricator is to apply thieving.
- *Keeping thieving away from conductive patterns where the thieving would have a negative impedance impact.* This has not historically been handled by inferior CAD tools, leading to complex notations on the fabrication print. pcbnew will use thieving keep-outs defined on the keep-out/in layer to define areas on the board where thieving is not to be applied. Thieving keep-outs can be plotted on the fabrication print or detail for situations where the fabricator is to apply thieving.

The parameters of thieving patterns (illustrated in *Fig. 42(51)*) are as follows:

— *Shape*, the shape of the elements in the thieving pattern. Obvious choices are squares, circles, octagons, squares with cut corners (an octagon is a trivial case of a square with cut corners), squares with rounded corners.

— d, the horizontal and vertical diameter of the elements. Typical thieving patterns have elements with a diameter of 30mil to 80mil.

— p, the pitch (center-to-center spacing) of the elements. Typical center-to-center distance between elements is 50mil to 100mil.

— s, the separation or clearance between the neutral elements and conductive features. This is typically $p - d$ or $p/2$ and is typically from 20mil to 50mil.

— o, the offset of the grid (element) centers with respect to the origin. Typical choices for interlocking patterns would be 0 or $p/2$. The coverage of the pattern can be tweaked by tweaking the offset.

Thieving Material Because thieving pattern elements are simply formed by etching copper, the materials are the same as the copper cladding or foil on the external layer in question.

Thieving DFX Rules Thieving will follow the same rules as pcbnew already uses for copper clearances to non-functional copper (that is, copper not assigned a net).

Thieving Imaging The imaging of thieving should, for the most part, follow the imaging of other copper areas on external board surfaces. There are some special considerations as follows:

Display: Thieving should be able to be displayed using separate visibility and colour controls. It is a question whether to place thieving on the "plating" layer. To assist with tweaking the thieving pattern, the thieving elements that are removed due to conflict with conductive features should be able to be drawn in outline mode (that is, sketched in over other items). This might most easily be accomplished by XOR'ing in the outline of the removed thieving elements. Application of thieving is not something that is necessarily a good thing to do while drawing traces, due to the significant computing involved. Also, because thieving is normally something considered toward the completion of a design, performing a separate option equivalent to filling a zone, is not too restrictive. Therefore, population of thieving patterns will be a "filling" operation on a thieving zone.

Plotting: Whether to include thieving when plotting should be a general fabrication output option. When requested, thieving will be added to plots. When not requested, thieving will not be plotted. When plotted on RS-247X plots, thieving will be plotted using a separate information (merge) layer and will use an aperture defined only for thieving. This will permit fabricators to remove user-supplied thieving as necessary. It might be an idea to have an option to plot thieving in a separate file. Thieving keep-out areas can be plotted in RS-274X and DPF plots as clear overlays placed underneath the conductive features of the board. When thieving is later applied ahead of (underneath) the overlay, thieving will be clipped from these areas. There should be an option to remove these overlays as it may confuse some CAM systems.

3D Display: Thieving should be displayed like any other copper feature. Visibility rules for thieving should also be available to the 3D display.

Thieving Post-Processing The plotting of thieving should be a fabrication output option. The designer should always have the ability to specify thieving keep-ins and keep-outs and have these printed on the fabrication print or detail. Thieving elements should not normally be included as testable points in (IPC-D-356) netlists: they are normally covered with soldermask anyway. Nevertheless, it is possible to test that an exposed thieve pattern is not shorted to a conductive feature when performing 100% board tests.

3.8.2 Venting

Venting (also called *"Resin Thieving"*) is the process of adding neutral copper zones or grid patterns to an internal layer of a multilayer PCB in an effort to balance the distribution of resin in B-stage epoxy laminate material (prepreg) and protect against excessive bow and twist of the resulting laminated board.

When copper features are laid out on a PCB, the uneven distribution of copper may lead to difficulties when pressing layers together into a board. Uneven distributions of copper lead to resin flowing away from copper features and toward open spaces, and can lead to layers floating out of position, elongation and contraction of layers under heat and pressure as well as during cooling. To avoid these difficulties it is necessary to balance the distribution (density) of copper features across any axis emanating from the board center. Copper feature density is not normally a problem where copper features exist on the layer. This is because the layout task is always focused on cramming more features into less space. The problem is normally open spaces devoid of copper features.

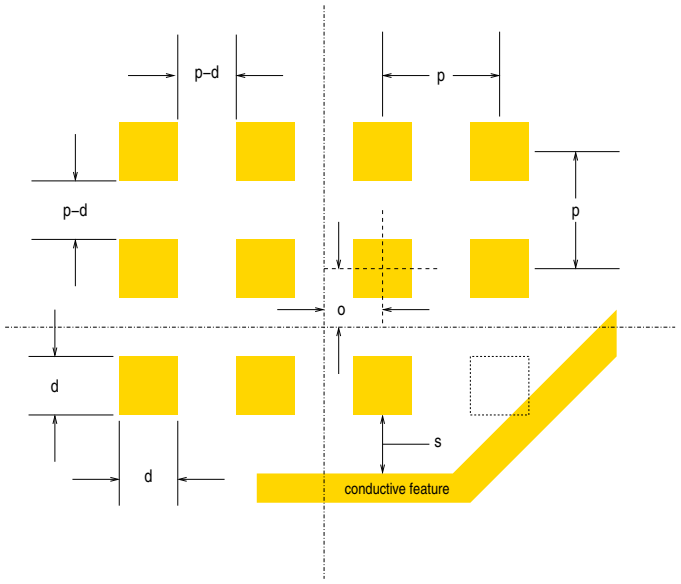

Figure 43: Venting Pattern

The easiest way of ensuring a good copper balance is to flood open spaces with neutral or grounded copper. This copper can be "neutral" in the sense that is it not connected to any net, or can be connected to ground or a voltage reference if compatible with the electrical design. Another way of providing venting to place a neutral grid pattern in the open areas.

Venting keeps the plane of material uniform and level. During lamination, the resins flow and has a tendency to become "frosted" or crystallized as they approach the outside edges of the panel. Therefore, the test coupons should be placed inward beyond the one-inch minimum border so as to avoid being in the crystallized area. It is best to remove non-functional pads as the pads without traces are less restricted and may "float" when the epoxy is flowing prior to cure.

Venting Process pcbnew has not historically supported automated venting. It was possible to place no-net pads of the appropriate shape; however, this was a tedious manual process. It is a job better suited for a computer. The parameters associated with venting are the same as those associated with thieving.[16] The only difference are as follows:

- The pattern is applied to inner layers of the board instead of exterior layers.
- There is some advantages that might be gained by offsetting the pattern on one layer with the pattern on adjacent layers to achieve an interlocking of layers.

Venting patterns are applied within a zone on an inner copper layer. The zone is marked for venting pattern and the following parameters control the venting pattern and its distribution:

— *Shape*, the shape of the elements in the venting pattern. Obvious choices are squares, circles, octagons, squares with cut corners, squares with rounded corners.

— *d*, the horizontal and vertical diameter of the elements. Typical venting patterns have elements with a diameter of 30mil to 80mil.

16. See *Sec. 3.8.1(50)*.

— p, the pitch (center-to-center spacing) of the elements. Typical center-to-center distance between elements is 50mil to 100mil.

— s, the separation or clearance between the neutral elements and conductive features. This is typically $p - d$ or $p/2$ and is typically from 20mil to 50mil.

— o, the offset of the grid (element) centers with respect to the origin. Typical choices for interlocking patterns would be 0 or $p/2$. The coverage of the pattern can be tweaked by tweaking the offset.

These parameters are illustrated in *Fig. 43(52)*.

Venting Material The material used is the same material as the inner layer copper to which venting is applied.

Venting DFX Rules The DFX rules are the same as for any other non-functional (no connection) copper, with the exception that pattern elements that violate DFX Rules are simply removed.

Venting Imaging Imaging is the same as Thieving[17] with the only difference being that they are applied to inner layers instead of outer layers.

Venting Post-Processing Post-processing of venting is, again, the same as thieving. Venting and thieving should be grouped together.

3.8.3 Crosshatch

Application of a crosshatch to copper planes is a technique to perform a negative thieving or venting; or to alter the mechanical and thermal characteristics of large copper planes. The technique derived from when one had to "paint" copper zones with line segments. The objective of crosshatching zones is to lower the density of copper in the area. This process has three objectives:

1. *Control of board bow and twist and laminate separation:* Large areas of copper on a layer of a multilayer board introduce a situation where the lamination of copper and dielectric forces two materials with a different coefficient of thermal expansion over a large area. The forces resulting from expansion and contraction as the board is heated and cooled can cause laminate separation. Crosshatching large copper areas serves to relax these forces and makes a multilayer board more resilient to temperature extremes.
2. *Negative plate thieving:* Large copper areas can thieve too much plating current from nearby features. Crosshatching of large copper areas can reduce the plating current absorbed by the area and improve the resolution of nearby features.
3. *Negative resin venting:* Large copper areas can cause a layer of the laminate in a multilayer board to float when pressed. Crosshatching, just as normal resin venting, provides for a more controlled press of the laminate.

Crosshatching is not as important for small board designs, where the forces involved are not so great. It is, however, important in large boards. Also, these objectives must be weighed, of course, against the electrical function of the copper zone. Also, where the area serves no electrical function, thieving or venting using a non-functional pad pattern is superior.

pcbnew has not historically handled crosshatching of copper zones, although it is a feature that has been requested on occasion.

(R) 24 (pcbnew) *Zones will be enhanced to provide for application of crosshatch as a general option of any zone on any technical layer.*

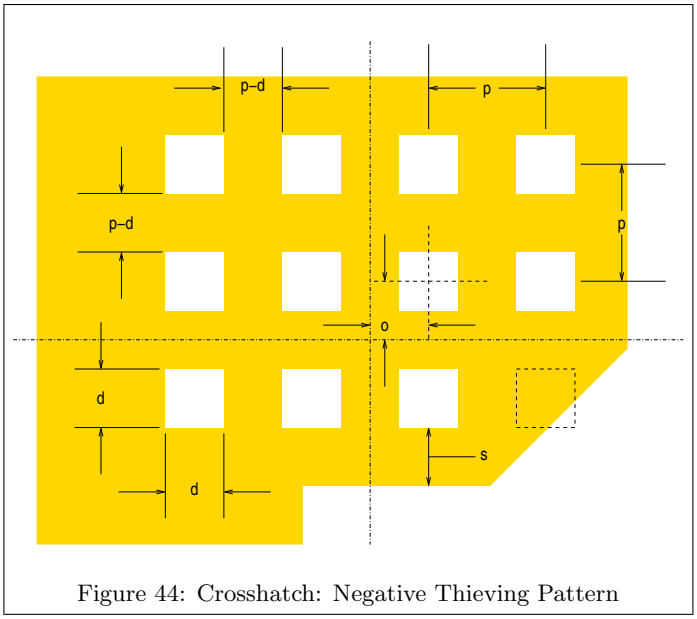

Figure 44: Crosshatch: Negative Thieving Pattern

There are a number of challenges when applying crosshatch to a copper zone. Most fabricators are concerned about the little areas (isolated corners at the periphery of the pattern) that might cause problems when etching or plating. When painting the pattern with line segments, it is difficult for some CAM systems to analyze all the little pieces. (Not to mention that painting zones in the first place has an effect on DRC check runs). To side-step these challenges, pcbnew will perform, what I call, a "negative thieving" when applying crosshatch. *Negative thieving* is what I call the process of providing a clear overlay of thieving patterns on a dark copper zone. The negative thieving pattern is illustrated in *Fig. 44(53)*. All elements that are completely within the zone, shrunk by its minimum thickness, will be applied. (See the white elements in the figure.) All elements that cross the edge of the zone, shrunk by its minimum thickness, will be discarded. (See, for example, the dashed-line element in *Fig. 44(53)*.) This approach has several advantages:

- There will be no little pieces or isolated corners in the resulting pattern.
- Because the negative thieving is applied in a clear overlay, it can be a separate information layer in a RS-274X plot and can use a unique aperture number. This will permit the easy removal of the pattern, or portions of the pattern, by the fabricator when desired. That is, it makes manual editing of the crosshatch easier.
- The pattern can use objects other than a square: e. g. they can be round, octagonal, or have blended interior corners, to avoid blossoming or mushrooming of the etch.

The technique uses the same parameters as thieving: p is the pitch for placement of objects: it corresponds to the line pitch when drawing crosshatch; d is the space between lines; $p - d$ is the line thickness; and s is the minimum separation between a crosshatch opening and the periphery of the zone; o is the offset of the pattern from the board or local origin. When octagonal shapes are used, it can correspond to the addition of 45° lines to the crosshatch. As with thieving, the absolute grid upon which these patterns are applied for a particular layer can be offset or interlocked with the patterns of opposing layers (the layer on the other side of a B-stage epoxy (prepreg) layer, or the layer on the

17. See *Sec. 3.8.1(50)*.

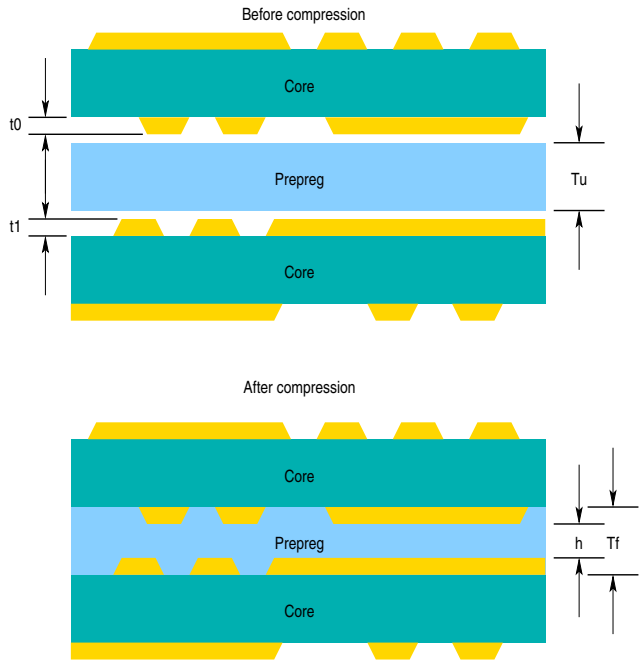

Figure 45: Press Factor

other side of the board). The parameters and calculations used for thieving and venting can therefore be reused for crosshatch.

(R) 25 (pcbnew) *The "negative theiving" approach to crosshatch will be provided in* pcbnew *to support the crosshatching of any defined* ZONE_OUTLINE *on any layer.*

3.9 Stack-Up

pcbnew has not historically handled stack-up. The furthest extent to which pcbnew went was to specify the overall board thickness and number of layers. The overall board thickness was only used to equally space layers for 3D display.

Items of particular concern are as follows:

3.9.1 Dielectric

3.9.2 Press Factor

Press Factor is the difference in thickness between a B-stage laminate (prepreg) before and after lamination (compression). The finished thickness of the laminate is dependent upon the copper coverage of the copper layers being pressed into the prepreg. This is illustrated in *Fig. 45(54)*.

Historically, pcbnew did not handle press-out: it did not handle dielectric layers or z-axis profile at all. pcbnew has been enhanced to handle stackups, dielectric layers, thicknesses, tolerances, dielectric constants, xy-axis thermal expansion, glass temperatures, etc.

To perform the calculation of press factor, etch compensation[18] may be calculated, but is often unnecessary: the volume displaced by the compensated etch is typically the same as the uncompensated (CAD) displacement.

Several parameters contribute to the press factor and the resulting finished thickness. The unfinished thickness of the prepreg, T_U, is, unfortunately, usually not accurate to better than ±10%. This means that the overall result can be no better than ±10%. The thickness of the upper copper foil, t_0, is a little more accurate. So is the thickness of the lower copper foil, t_1. If we calculate the surface area of the upper copper, A_0, and the surface area of the board, A_B, then the volume of the upper copper is

$$V_{C0} = t_0 \times A_0. \tag{55}$$

The same is true for the lower copper:

$$V_{C1} = t_1 \times A_1. \tag{56}$$

The volume of the prepreg is:

$$V_P = T_U \times A_B. \tag{57}$$

The finished thickness is, therefore:

$$T_F = \frac{V_P + V_{C0} + V_{C1}}{A_B} \tag{58}$$

$$T_F = (T_U \times A_B + t_0 \times A_0 + t_1 \times A_1)/A_B \tag{59}$$

$$T_F = T_U + t_0 \times A_0/A_B + t_1 \times A_1/A_B \tag{60}$$

The dielectric height, h, between conductors is:

$$h = T_F - t_0 - t_1 \tag{61}$$

$$h = T_U - t_0(1 - \frac{A_0}{A_B}) - t_1(1 - \frac{A_1}{A_B}) \tag{62}$$

Now, A_0/A_B and A_1/A_B can only range from 0...1, so,

$$h_{min} = T_U - t_0 - t_1, h_{max} = T_U \tag{63}$$

Guesstimating that $h = T_U - (t_0 + t_1)/2$,

the maximum thickness can be $h_{max} = 1.10T_U$, and the minimum thickness can be $h_{min} = 0.9T_U - t_0 - t_1$, so the maximum error is:

$$\pm(0.10 + \frac{t_0 + t_1}{2T_U} \tag{64}$$

For $T_U = 0.003$" and $t_0 = t_1 = 0.0002$", the error is ±16.7%, which is obviously not good enough to guesstimate. The actual copper areas must be calculated or accurately estimated to get the accuracy back down toward ±10%. Of course, the dielectric height, h, is fundamental to calculating stripline impedances.

An approximation can be made using resin-venting patterns. For a pattern of 0.080" squares on a 0.100" pitch, the coverage is about 90%. Selecting at this copper density, the error will be the deviation of functional conductive features from the 90% coverage which may or may not be so great depending on the complexity of the signal layer. The error can maybe then be trimmed to between 3% and 5%.

Estimation of copper areas: The copper areas can be estimated as follows:

- For stripline, one copper layer is normally a reference plane and the other is normally a signal plane.
- The copper density, A_0/A_B, for ground and power plane layers can be assumed to be $A_0/A_B \approx 1$.
- The copper density, A_1/A_B, for signal layers can be assumed to be $0.5 \leq A_1/A_B \leq 0.8$. This is assuming that unused areas are flooded with copper, or resin venting. Cutting the difference, say $A_1/A_B \approx 0.65$.

To refine the areas of signal layers requires a technique along the lines of the following:

1. *User input:* ask the user. A human being can likely estimate the copper density on a layer faster than (but, of course, not as accurately as) a computer.

18. See *Sec. 3.4.1(34)*.

2. *Statistical approximation:* sample the layer by selecting points using a pseudo-random number generator. Use the statistic as the expected copper coverage. Choose additional points until the expected value is significant to 4 sigmas.
3. *Track and zone summation:* For each track, add its length times width to the copper area. For each zone, calculate the filled area of the zone minus the area of its cutouts. Use this approximate number as the copper area.
4. *Bit map:* create a bit array for the entire board. For each of the copper elements on the layer, plot the element on the bit-map. Add the number of ones in the bitmap. (This can be done while plotting by counting only bits that are not already plotted.) Use the bit-map count as the copper area.
5. *Brute force:* calculate the area of every copper feature, subtracting its intersection with other features.

Likely the easiest approach is the first: a reasonable approximation can be achieved without computational or programming effort. ***Note that when calculating the copper areas, resin thieving (venting)***[19] ***must also be considered or a poor estimate will result for signal layers.***

3.10 Panels

pcbnew is only capable of providing rudimentary panel features. Its inability to do step-and-repeat is the major indication that its purpose was to provide only one-up board data; however, it can provide one-up data that is outside the board edges. Note also, that multiple boards (should be different boards with the same stack-up; e.g. a motherboard and a daughter card). The sections that follow detail some of the rudimentary panel items that can be supported.

The following are implementation considerations for fully supporting panelization:

- A new ARRAY class will be created that has placement information for boards. The class will refer initially to only one board. The class will contain a list of placements for the board. The array class will contain board items like a board and can be derived from the BOARD class. The board will be a proximity item within the array class and the array class will have a proximity map for board instances. The board will place itself in the array's proximity map. Board items belonging to the array will be placed directly in the proximity map.
- A new PANEL class will be created that has placement information for arrays. The class will refer initially to only one array. The class will contain a list of placements for the array. The panel class will contain board items like a board can can be derived from the BOARD class. The array will be a proximity item within the panel class and the panel class will have a proximity map for array instances. The array will place itself into the panel's proximity map.
- Two new modes will be provided for viewing and editing arrays and panels.
- Arrays can have board items that are owned only by the array to provide array tooling and fiducial marks.
- Panels can have board items that are owned only by the panel to provide panel tooling and fiducial marks.
- When board items are displayed, in board editing mode, they will be
- When in board editing mode, only the board will be displayed and edited. Other features belonging to the array and panel will not be displayed. This is the classical view.
- When in array editing mode, board items and boards belonging to the array will be displayed. Board items whose parent is a board will display themselves multiple times, once for each board instance and transformation in the array. In array editing mode it will not be possible to edit board items belonging to boards. Boards belonging to arrays can be moved as a unit (by changing the transformation). Board items whose parent is the array will be displayed once relative to the array's origin. It might not be necessary to display more than a board's boundary box in this mode.
- When in panel editing mode, board items belonging to board, arrays and the panel will be displayed. Board items whose parent is a board will display themselves multiple times, once for each board instance and transformation in each array instance and transformation in the panel. In panel editing mode it will not be possible to edit board items belonging to boards or arrays. Arrays belonging to panels can be moved as a unit (by changing the transformation). It might not be necessary to display more than an array's boundary box in this mode.
- To support the proper plotting of arrays and panels, two additional modes will be provided for plotting: array plot and panel plot. In these plotting modes, each board item will plot itself multiple times, once for each transformation within the array or panel.

3.10.1 Fiducial Marks

1. Local Fiducial Marks.
2. Board Fiducial Marks.
3. Array Fiducial Marks.
4. Panel Fiducial Marks.

See *Sec. 3.2.13(25)* for more information on the pad features necessary to properly support fiducial marks.

Board Fiducial Marks. Board fiducial marks are used to mark a bad board in a panel so that the assembly robot can detect whether a board is bad. These fiducial marks are typically placed outside the board edges.

3.10.2 Coupons

1. *Scoring Coupons:* Scoring coupons are used to detect the lack of a score mark on a path that needs to be scored. Detection of the missing score mark is accomplished during bare-board testing, due to the presence of the coupon.
2. *Beep Test Coupons:* Beep Test coupons are used to detect copper layer registration or drilling registration that is outside the required tolerance. For example, on a IPC-6012 Class 1, 2 or 3 board, a drilled hole that does not meet the requirements of the class can be detected using a beep test coupon.
3. *Impedance Coupons:* Impedance coupons are used to detect impedances that are out of range. They are intended to reflect the impedance requirements for traces on specific impedance controlled layers of the card by mimicry of those traces within the coupon. An impedance test set can then measure the actual impedance of the coupon versus the target impedance to determine whether the board as built is within the tolerance requirements for impedance.

19. See *Sec. 3.8.2(52)*.

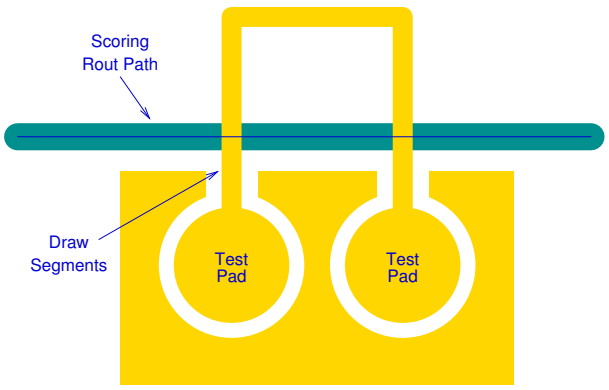

Figure 46: Scoring Coupon

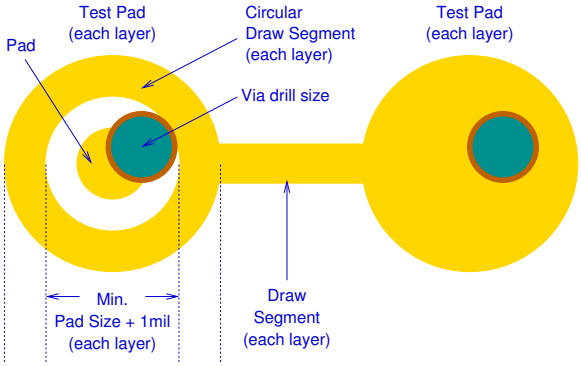

Figure 47: Beep Test Coupon

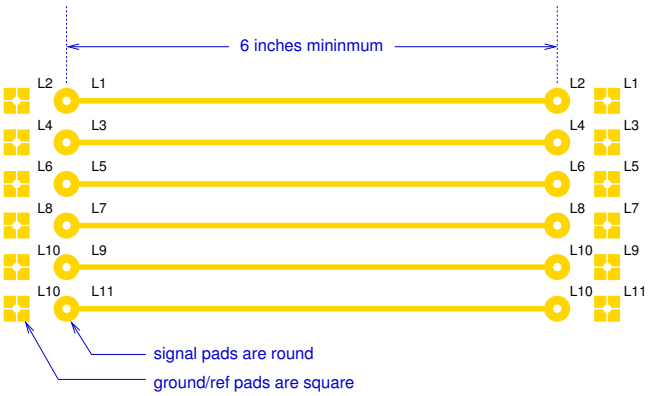

Figure 48: Impedance Coupon

Scoring Coupons. Scoring coupons are used to detect the lack of a score mark on a path that needs to be scored. Detection of the missing or misplaced score mark is accomplished during bare-board testing, due to the presence of the coupon. Scoring coupons can also be designed to determine whether the score mark deviates from the design path by a greater amount than the allowed or promised tolerance. A scoring coupon that fails a bare-board test can be cause to visually inspect the board for acceptability.

As illustrated in *Fig. 46(56)*, scoring coupons consist of two pads within the board outline that connect across the score mark at the edge of the board. This is for v-groove scored boards and not for routed boards (routed boards do not need such coupons because the tool bit passes through the entire book). Because the pads need to be considered an open circuit in the board testing data, pcbnew draw segments (DRAWSEGMENT on the top and bottom copper layers should be used to connect the pads inteaad of tracks. This way, the board testing data generated does not consider the pads to be connected to anything. Also, using draw segments on copper will not invoke the wrath of DRC when drawing traces across a score mark.

It is quite possible to create scoring coupons with the pcbnew module editor. This coupon consists of two pads that exist on front and back layers (but no internal layers): this can be accomplished with four pad definitions. Arbitrary pin numbers can be assigned (1, 2, 3, 4). A component could even be created in the eeschema library to invoke the addition of one of these coupons. The pins could even be marked as not connected. A reference designator of SC* could be assigned. However, it is tedious and unnecessary to define a component in eeschema every time one wants to place a scoring coupon. The number of scoring coupons on a board should not be a function of the schematic, but one of the board design. Because addition of any module into pcbnew is allowed, this can be accomplished solely from pcbnew. When a module is added in this fashion, the pads are not assigned a net. In this case, we want the pads to not be assigned a net (they will then be marked as N.C. (not connected) in the IPC-D-356 file).

Beep Test Coupons. Beep Test coupons are used to detect copper layer registration or drilling registration that is outside the required tolerance. For example, on a IPC-6012 Class 1, 2 or 3 board, a drilled hole that does not meet the requirements of the class can be detected using a beep test coupon. A beep test coupon that fails a bare-board test can be cause to visually inspect the board, possibly with x-ray equipment, to determine acceptability of the board.

Beep Test coupons consists of a drill hole surrounded by copper at a specified distance. The purpose of the coupon is to indicate to board testing when a drill hole is out-of-range, by plating through from the upper annulus to the lower annulus, completing a short (according to the netlist). Beep Test coupons, because they are normally open, can be formed with normal board features.

> The clearance diameter must be sized in a manner that takes into consideration etching capability based on copper weight. The clearance diameter should be determined at the foot of the etched feature. The clearance must be a minimum of 0.001” larger than the minimum annular ring diameter. This prevents beep test failure at tangency and provides allowance for etch tolerance. The optimum beep test clearance should be no less than 0.013” larger than the drill diameter used to drill the hole within the feature. Optimum drill diameter used to drill the clearance feature of the coupon should be between 0.030” and 0.070”. Specify only one beep test coupon per corner of the panel (4 total). [METRIX04]

Because the pass-state of the beep test coupon is open-circuit, the connections between the test pad and the beep test clearance area must be performed with draw segments in pcbnew so that the netlist (IPC-D-356) file will not consider the patterns to be electrically connected. As with the scoring coupon, the beep test coupon can be created using existing pcbnew pads and draw segments in a module. The only complication is that draw segments must be generated on all PCB layers, and the clearance of the circular beep-test portion must match the minimum annular ring on all layers. The existing pcbnew module editor does not allow this. A specialized module might be necessary to avoid errors in calculation and exact drawing.

Impedance Coupons. Impedance coupons can be included on the board (within the board edges); however, it is not normally possible in high density designs to do so. On the other hand, the route path might permit placing these coupons within an area cut out by edge connector contacts, or it might be possible to place these coupons on an assembly rail. These coupons might require copper or silkscreen text.

> It is common practice to perform 100% testing of controlled impedance boards, usually required as the acceptance criteria by the customer. Merix utilizes a thirteen board AQL sample plan. Any failures from this plan will generate a 100% test requirement. With any testing requirement there will be difficulty in accessing the controlled impedance trace for verification. There may be an option of adding test pads on the board, but this may affect the performance of the trace and occupy valuable routing space.
>
> Separation within the plane layers or split planes can raise issues with the accuracy of impedance testing. Separations within plane layers can introduce crosstalk and add inductance to traces running perpendicular over the separations. This increased inductance is sue to the long return current loop and can have an additional [effect] on the rise times of the signals. Actual PCB traces are typically shorter and include branches to circuitry and to vias between layers. This can affect the integrity of testing impedance-controlled traces within the actual circuitry of the board.
>
> With these factors affecting the measurement accuracy of the actual circuitry, it is common practice for the PCB manufacturer to generate impedance test coupons. These coupons represent the circuitry designed within the board, but are controlled to known trace lengths, which the impedance [formulae] are generated from. [sic] The test coupons are typically small PCBs with trace lengths at a minimum of six inches and are located at opposite ends of the panel, external to the board image. The coupons are processed under the same controls as the main PCB. The testing is then performed on the coupon and the impedance data can be correlated back to the feature size and processing parameters.
>
> Another option used by some designers on very large board images is to incorporate a test coupon or test traces internal to the board image. It is imperative that these internal coupons or test traces adhere to the test coupon description below. This is to [ensure] ease of testing and result in accurate measurements.
>
> **Test Coupon Description**
>
> The impedance test coupons are actually small PCB on the same manufacturing panel as the main board itself. They are typically one inch wide by eight inches in length and include traces of the same width and copper weight as the main PCB (depending on the number of layers within the design). The hole size and patterns are critical to [ensure] accurate testing with standard SMA probe heads used by most TDR manufacturers. See Figure 1.
>
> The coupon will be located on the manufacturing panel in a position to best represent board conditions for etching, plating, lamination, and surface finishes. The coupon is also positioned on the panel to best represent consistency across the entire manufacturing panel. When the artwork data is generated for the controlled impedance traces, the same aperture or 'D' codes used on the main board are also used on the coupon. ... [METRIX06]

1. Dielectric separation will replicate impedance structure on printed boards.
2. Test connection holes shall be plated through hole connections to access all ground/power reference planes.
3. Square pads identify plated through hole connections to access all ground/power reference planes.
4. Conductor widths will replicate critical conductors on each impedance layer.
5. Via holes to be added as required.
6. Cross hatching to be added to outer layer as required.
7. Two coupons per panel. These to be individually identified with letter A & B respectively.
8. Job No. + Date code to be added as per customer requirements.
9. All planes to be interconnected (on test coupon only).

To minimise capacitive loading during test, you should minimise the size of pads and vias on coupons, especially for high impedance traces. Although this standard coupon design shows pads at both ends, you are likely to obtain better results on high impedance traces if you only place a pad at one end.

- Suitable test lines need to be provided by the designer for each layer with impedance requirements.
- These lines need to be a minimum of 3.0 inches long (ideally 5.0 inches) without networking into another layer.
- They also need to be accessible from the external layer with a 0.030" minimum diameter hole, and be within 0.150" of another hole of the same diameter, making connection to the reference plane.
- If you do not have a test coupon, one will be provided for you.
- For the test pattern, select an appropriate hole from from the circuit board drawing.

3.10.3 Boards

3.10.4 Assembly Panels

1. Breakaway Assembly Rails and Tabs.

Breakaway Assembly Rails and Tabs. For scored boards there is no issue (except the cost of v-scoring both sides of the board, when, by comparision, a number of boards may be routed and drilled simultaneously in a book.). For route boards, breakaway tabs are normally made with perforations (drill holes betweeen route path ends) that permit the board to be snapped from the breakaway bars after assembly.

In fact, pcbnew used to provide board edges as a specific width of edge centered on the board edge. While such data is suitable for scoring, it is not suitable for routing and perforating tabs. Therefore, new attributes have been added to permit the definition of exterior and interior tool paths for routing as well as the ability to specify a segment as perforated (dotted or dashed for graphics). Also the ability to describe a panel edge for breakaway tabs is also necessary.

3.10.5 Production Panels

1. Tooling Holes.
2. Plating Bars.

Tooling Holes. Tooling holes (2, but better 3) should be created diagonally across the board as a reference for routing the board. These tooling holes are placed outside of the route tool path.

Plating Bars. Board edge connector contacts require electrical connection to a plating bar necessary for electroplating the contact material (usually hard gold). Again, because these traces need to cross a bevel and score or route path, they need to be defined as copper-layer draw segments. Sometimes it might be necessary to even provide the electropating device connector pads;

however, in many panelizations, this overlaps from one-up to one-up and is considered panel data. Placement of the plating bar might require a knowledge of the eventual panelization, but the user is in control here.

Note that the board is not routed, scored or bevelled until after the electroplating is complete.

4 Design for Excellence

Note that this section talks about DFM (Design for Manufacturability) or DFX (Design for Excellence) instead of just DRC (Design Rule Check). This is the case because I always take a holistic approach to these things. I resist the *Tower-of-Babel* effects that come from classical division of labor exercises in the electronics industry, and instead believe in *"one tool to bind them all"*.

The serious downside to division of labor and specialization in the manufacture and assembly of PWBs in the industry appears to be that once a fabricator and assembler are chosen, there is a significant amount of lock-in because the fabricator has massaged the data and added value to it and the assembler then massages the data and adds value to it. Once everything works, changing fabricators or assemblers is the last thing from a designer's mind.

Another downside is that selection of a fabricator and assembler in the first place is sort of hit-or-miss. Most recommendations seem to be word-of-mouth and of the nature of *"worked for me..."* The more that fabricators take on, the more hidden the deficiencies of their capabilities becomes.

But from the standpoint of bleeding-edge, high-speed boards with ultimate signal integrity, this division of labor has got to stop. Running expensive and time-consuming 3D FEM (Finite-Element Modelling) and full wave simulations on traces that have been laid out in precise dimensions and separation to height ratios only to be widened, built up too high, and pressed into a mess of prepreg goop that comes nowhere close to the height needed and nowhere near the separation needed (sure differential impedance might be OK, err, say, 50 Ohms, but what about the Zo and Ze: these have tolerances too!) All with the seeming goal of setting the resulting bow and twist of the board above its ability to function electronically. I say, "*Give me a bowed board, just as long as it functions from 40Gbps to 100Gbps!*" I hate to tell a fabricator, but the best way to manufacture a board that has no bow or twist is to *not put any traces on it! Darn... pesky.... traces.*

Some fabricators have gone to the point of saying, "*Hey, we'll build the whole darn board for you!*" Yah, using protel or orcad or something. To this I say, "*Nooo... because I frankly don't want you to know a darned thing about how it works: because, if you do, you're just going to steal it!*"

Soooo, for this and other rants, raves, and religious doctrines, I seek to place all of the necessary features, information and functions into pcbnew so that the designer can, once again, control the final product! So don't complain about all the fabrication and manufacturing mumbo jumbo and details that follow: if you don't want to know about them, *let 'em default!*

5 Layers

5.1 Dielectric Layers

Historically pcbnew has not handled dielectric layers. It has just assumed that dielectric is that stuff between the spaces of the copper layers. pcbnew did naively assume that the dielectric thickness was the overall board thickness divided by the number of copper layers minus one for multilayer boards, on the thickness of the board for single-sided boards. This was not sufficient for several reasons:

1. Stack-up editing was not possible.
2. Impedance calculations were not possible.
3. It was not possible to export a 2.5D or 3D-image of the board for 2.5D or full wave simulation of electric fields.
4. The 3D display cannot represent the board properly in the z-axis.
5. Press factors could not be calculated.
6. The tolerance of the overall board thickness could not be calculated.
7. Whether a given number of layers was even possible for a specific board thickness was not possible.

All of these deficiencies contribute to the ability of pcbnew to design boards that simply cannot be fabricated. This is very bad for the design process, contributing to wasted time and effort on the part of the designer.

5.1.1 Dielectric Process

The dielectric and laminate construction process is well known and can easily be modelled mathematically. Single-sided boards are made using a copper-clad dielectric that is typically a C-Stage (fully cured) fibreglass epoxy-resin composite (FR-4). Multilayer boards are made by laminating alternating layers of copper clad C-Stage fibreglass-resin cores and B-Stage (partially cured) fibreglass-resin prepreg layers. Layers of prepreg and core are balanced across the centre of the board in the stack-up on the z-axis. To maintain symmetry, an even number of copper layers are provided. Two constructions exist: one with a core at the centre and prepreg covered with copper foil as the external layers. This is called a copper foil construction. Another has prepreg at its centre and copper-clad core as the external layers. This is called a core construction. Copper foil constructions are by far the more popular constructions. The book of layers are pressed in a heated press, sometimes under vacuum, to laminate the layers into the multilayer board.

C-Stage FR cores do not press-out during lamination and retain their unfinished thickness. B-Stage prepreg, on the other hand, presses into and around the traces that form the adjacent copper layers. This press-out contributes to deviation in thickness of the prepreg layers from unfinished to finished, and affects both overall board thickness and impedance calculations for those layers. A process of placing copper in low pressure areas of the prepreg, called resin venting, can be used to control the thickness of the finished dielectric layers. Press-out is detailed in *Sec. 3.9.2(54)*. Resin venting is detailed in *Sec. 3.8.2(52)*. Calculation of the finished thickness of a prepreg layer is a simple volumetric calculation as the layer does not exhibit significant plastic elasticity.

5.1.2 Dielectric Materials

Dielectric materials are available in a range of dielectric constants, dissipation factors, loss tangents, glass expansion temperatures, and other physical properties that affect the finished thickness, impedance, structural stiffness, resistance to thermal shock, and other mechanical and electrical considerations. The primary parameters are as follows:

- The permeability, permitivity, loss tangent and conductivity of the material.
- The thermal conductivity of the material.
- The dielectric breakdown voltage for a given dielectric distance.
- Whether the dielectric layer is core or prepreg.
- The glass temperature.
- The coefficient of thermal expansion in the xy-plane.
- The coefficient of thermal expansion along the z-axis.
- The unfinished thickness (or finished thickness for cores).
- The number of fibreglass plies.
- The density of the fibreglass weave.
- Tolerances for all items.

5.1.3 Dielectric DFX Rules

Design for manufacturability rules are as follows:

1. Dielectric core and prepreg can be supplied in a range of thicknesses, however, not all thicknesses are possible.
2. Dielectric core and prepreg must be alternated throughout the stack-up with the possible exception of the outermost layers, which can be sequential laminations of prepreg.
3. There is a minimum core thickness and a minimum prepreg thickness.
4. There is a maximum core thickness and a maximum prepreg thickness.
5. For a given dielectric strength, the distance between two conductors in any spatial direction can be no less that a specified minimum distance. This can be as large as 0.0035".
6. Dielectric layers (and the intervening copper layers) must be symmetric in material and thickness about the centre of the stack-up.
7. Dielectric layers can be sequentially laminated and drilled to form blind and buried vias; however, a core must always be included in the sub-laminate; sub-laminates must also be symmetric about their centres; and there is a limit to the number of press cycles that can be performed while maintaining xy-plane registration of the copper layers within reasonable limits (usually only three press cycles are usable[20]).

The most compute intensive tasks for design and rule calculations for dielectric layers involve the calculation of press factor (see *Sec. 3.9.2(54)*) for finished thickness, and the DFM calculations for the minimum dielectric distance between conductors. Where multiple signal layers are sandwiched between reference planes, it is possible that the dielectric thickness of the intermediate layers need to be as large as the required minimum dielectric distance for the construction. It is possible to interleave signal layers in such as way that conductors do not cross. Therefore, it can be difficult to calculate the minimum dielectric distance without considering the geometry of the signal layers. This is a good task for a computer.

5.1.4 Dielectric Design

Designing dielectric layers, with the exception of embedded capacitors, are a simple matter of designing the cross-section of the lamination of the multilayer board, in accordance with the rules. Dielectric layers do not require photo images and, therefore, do not require plotted artwork. Most of the design parameters associated with dielectric layers have to do with the electrical and

20. UL limits the maximum number thermal press cycles to three.

mechanical properties of each dielectric layer. The dimensions of concern are the cross-sectional thicknesses of the layers. Also of design consideration is the sequence of lamination of dielectric and copper layers to form sub-laminate sequences for the construction of blind or buried vias. For additional considerations for stack-up design, see *Sec. 3.9(54)*.

5.1.5 Dielectric Imaging

Display Dielectric layers are not normally imaged, and therefore they are not displayed in the pcbnew canvas of the board editor. The are assume to occupy all of the space between copper layers that are not cut out with board edges or rout paths, which are imaged separately.

3D-Display Dielectric layers, on the other hand, can be displayed in the 3D view. They are either represented as the empty space between copper layers, or can be displayed by fogging the space between the copper layers. The extremities of the dielectric are the board edges, or other rout paths forming cutouts, slots or wells. Drilled, punched or ablated holes pass through one or more dielectric layers.

5.1.6 Dielectric Post-Processing

Plotting. Dielectric layers are not normally imaged. Therefore, the dielectric layer is not plotted in pcbnew.

Nevertheless, it is typically to provide a stack-up drawing on the main fabrication print, or detail sheet, that provides the cross-sectional view of the multilayer board, identifies the dielectric material, finished thicknesses, and impedance requirements.

5.2 Resistance Layers

Resistance layers are a layer of etchable resistance that is sandwiched between a copper layer and a dielectric core that are used to form embedded resistors. Copper clad dielectric core material can be ordered that includes this resistance layer under the copper cladding. The resistance ply is etched along with the copper etch mask, and then a secondary mask is applied to remove copper from the area of the resistor. Therefore, this layer is photo image layer and is post-processed and plotted.

5.2.1 Resistance Process

5.2.2 Resistance Materials

Resistance materials are preconstructed copper clad dielectric cores that have the resistance layer sandwiched between the copper cladding and the dielectric core. The resistance materials are chosen for their resistivity and their ability to be etched along with the copper cladding, or to have the copper cladding etched from above them without etching the resistance layer itself.

Copper clad dielectrics with embedded resistance plys can be ordered in a combination of thicknesses for varying $\Omega/\square$ resistivity. Characteristics of the laminated core consist of:

- The volumetric resistivity (conductivity) of the resistance ply material.
- The thickness and tolerance of the resistance ply material.
- The normal copper cladding parameters (roughness, thickness, conductivity, tolerance).
- The normal dielectric parameters (dielectric constant, thickness, tolerance).

5.2.3 Resistance DFX Rules

Stack-up handling. There are several ways to handle stack-up rules for embedded resistors. First, one could require the user to define embedded resistance layers prior to allowing embedded resistor modules to be placed on the corresponding layer. This would buck the ability of the designer to experiment with placement of embedded resistors without having to suspend DFX rules, or switch back and forth between the stack-up definition dialogue and the main board editor canvas. I don't think that is the best approach.

Second, one could allow the designer to place embedded resistors in any layer that they choose. Then, when editing the stack-up, or controlling the visibility of layers, resistance plys would only appear for layers on which embedded resistors are populated. This approach would allow the designer to experiment with embedded resistor placement and is the approach that will be taken.

5.2.4 Resistance Design

The design approach to creating embedded resistors using resistance plys is detailed in *Sec. 3.6.13(49)*.

5.2.5 Resistance Imaging

Display To facilitate the display and editing of embedded resistors, resistance layers (or plys) will only appear on the visible layers list for resistance layers that actually contain an embedded resistor. Any other resistance layers that do not contain embedded resistors will not be displayed in the visibility list.

The resistance portion of embedded resistors will form a new visibility item, or could share visibility with other modules.

3D-Display Resistance plys can (and should) be rendered in the 3D display.

5.2.6 Resistance Post-Processing

Plotting. Forming the resistance layer into shapes requires the etching of the corresponding copper layer along with the resistance ply to form the outline of the resistor. A secondary etch mask is then used to etch away the copper from the resistor areas. Therefore, two plots are required for every layer that contains a resistance ply: one (normal) plot of the normal copper traces and resistor outlines; and another (special) plot of the copper that is to be removed from the resistor. The "resistance layers" are provided to plot this latter image. The numbering of resistance layers is the same as the numbering of copper layers. The plots, however, must be identified as resistor etch plots. The abbreviations, EtchRes or ResEtch, have been used elsewhere. Copper layers that do not have embedded resistors do not need to have a corresponding ResEtch layer plotted.

5.3 Copper Layers

Copper layers are layers for etch masked copper (and possibly resistance plys under the copper). It addresses all of the specifications and rules that are common to all copper layers. Historically, pcbnew only supported a maximum of 16 copper layers. All copper layers on the board. Plots associated with these layers are primary photo-resists and plating-resist patterns. These layers are board fabrication layers.

Number of Copper Layers: The number of copper layers is either 1, or a multiple of 2. To avoid bow and twist (warpage) in multilayer boards, the weight of copper should be symmetric about the center of the board. An exception to this rule is when HDI layers are added to one side of the board.

Copper Weight: Copper sheets come in a number of thicknesses. Fabricators do not usually roll their own copper but order it in standard sizes. pcbnew should be programmed with these standard sizes. Thicker copper sheets are not as workable on inner layers as they are on outer layers. This is because the height of copper traces is taller and pressing the board together becomes more difficult. Thinner copper fois are not as workable on outer layers as they are on inner layers. This is because handling the foil becomes more difficult. Typical inner signal layers use 0.5 ounce copper; inner power and ground planes 1.0 ounce. Typical outer layers are 0.5 ounce copper.[21] Deviations are normaly only performed due to high-current handling requirements.

Plating Factor: As a result of etching and plating, the finished thickness and unfinished thickness of the copper layer are different. Because the process for inner layers is often different from that of outer layers, there is a difference between finished and unfinished thickness on inner and outer layers. The difference in thickness, from unfinished to finished, is predominately the plating factor.

Etch Factor: The etch factor is the ratio of the difference in width of a finished trace at the upper copper surface to the height (copper thickness) of the trace at the inner copper suffice (against the laminate). A factor of zero would be a perfect process, resulting in a rectangular trace cross section. Typically the etch factor is 0.30 to 0.40, resulting in a roughly trapezoidal trace cross-section. The etch factor is necessary for more accurately calculating (or estimating) the theoretical impedance of a microstrip or stripline construction. Therefore, although it is traditional to use a sort of finished width for traces, and that is what we do on design file plots anyway, ask your fabricator exactly what they are going to do, and what their etch factor will be, and tell them precisely how to compensate for it. Trace width deviations are one of the most significant influencing factors affecting impedance: get this right. Etch factor (or compensation) is discussed in detail in *Sec. 3.4.1(34)*.

Minimum Trace and Space: Because the characteristics of etching and plating processes often differ for the inner and outer layers of a multilayer board, the minimum trace width and feature spacing (so called, "trace and space") depend on the layer.

Plating Thieves: To increase the predictability and accuracy of the plating process, it is typical to place plating thieves on outer layers. Plating thieves are little unconnected squares of copper that are laid in a grid pattern on board surfaces that do not contain any copper. These plating thieves rob plating ions from the plating process so that the main copper on the board is not over-plated in localized areas. These plating thieves can also minimize bow and twist resulting from heating and cooling while pressing the laminate. The new pcbnew keep-out/keep-in layer can be used to describe areas in which plating thieves are not permitted; or areas where they are permitted. These areas can also be used during plotting to apply a thieving pattern. The need for plating thieves can be avoided by flooding open spaces with copper. A good way to keep the fabricator from thieving, is to thieve it first! Plating thieves are discussed in detail in *Sec. 3.8.1(50)*.

Press Factor: The press factor is expressed as the difference between the unfinished and finished thicknesses of a prepreg layer, where the thickness is measured from the outer copper surface of the upper core layer to the mating outer copper surface of the inner copper layer. Knowing the traces on the board, their etch factor, their finished copper height, and the placement and size of resin thieves on the mating layers, it is possible to calculate how much of the prepreg will "press" into the spaces between the copper, and, thus, the resulting thickness of dielectric between the copper traces. When the laminate is pressed together in a vacuum, this theoretical value will be very close to the actual. The most significant deviation will be that the tolerance of prepreg is poor (worse than $\pm 10\%$). One of the reasons for going to all this trouble to calculate the copper-to-copper spacing is for determining the impedance of microstrip, embedded microstrip and stripline traces. Another is to check the minimum dielectric strength between traces on the board. Press factor is discussed in detail in *Sec. 3.9.2(54)*.

Resin Venting: To reduce bow and twist as well as to provide a more predictable Z-axis profile for a multilayer board, it is typical to place resin thieves in open copper spaces on inner layers. Resin thieves fill space and keep resin from flowing away from higher-density features. This is, of course, more important on signal layers than on reference plane layers. The new pcbnew keep-out/keep-in layer can be used to describe areas in which resin thieves are not permitted; or areas where they are permitted. The need for resin thieves can be avoided by flooding open spaces with copper. A good way to keep the fabricator from thieving, is to thieve it first! Resin venting is discussed in detail in *Sec. 3.8.2(52)*.

Registration and Annular Ring: Because inner layers are more difficult to align in the stack-up, XY-plane registration of inner layers can be worse than outer layers.

Annular Ring: Registration of outer and inner layers aggravates annular ring calculations for plate-through holes for inner layers. Often, inner layer registration will be worse than outer layer registration simply because the outer layers are visible in the layup. X-ray machines; however, have made registration of inner layers better in recent years, and most fabricators now specify the same inner and outer annular rings.

Clearances: Characteristics of the plating process (e.g. to what extent copper wicks down glass fibers in the laminate) place constraints on the clearance of copper traces or zones and PTH or vias. This results in several clearances:

- clearance of traces or zones to PTH;
- clearance of traces to NPTH;
- clearance of traces to board edges or routed slots;
- clearance of traces to vias or pads.

5.3.1 Copper Layer Process

The primary process considerations for copper layers are as follows:

21. Note that in imperial units, copper foil is measured by weight density (ounces per square foot), whereas in metric units copper fois is measured by thickness (microns).

- The etch factor for etching of copper traces and associated tolerances.
- The plating factor for plated copper foil on external layers and the associated tolerances.
- The registration tolerance of features etched into the copper layer, particularly with respect to drilled PTH.
- The ability of the etching process to resolve copper features, normally expressed in minimum trace and space capabilities.
- The ability of the plating process to resolve copper features, also normally expressed in minimum trace and space capabilities. Also, the plating factor or thickness and tolerance of plated copper on copper foil for external layers is also provided.

5.3.2 Copper Layer Materials

There are a number of copper layer material parameters that affect the electrical properties of the material:

- The conductivity of the copper material.
- The microscopic construction of the copper material. That is, whether it is electrodeposited, rolled, or annealed.
- The surface roughness of the copper material. This can significantly affect the resistance of copper traces at high frequencies for which skin-effect is significant. There can be a difference between the finished surface roughness and the unfinished surface roughness because an abrasion pattern is often applied to inner copper layers to increase the adhesion to the dielectric epoxy resin.
- The finished thickness and tolerance of the copper. The thickness is not normally specified, but the copper weight ($oz/\square$) is normally specified when imperial units are used, but micron thickness when metric units are used.

5.3.3 Copper Layer DFX Rules

Fabricators would have you ignore most of these things during design. They want artwork for copper traces that are design width (whatever that means) and they will adjust the thickness of those traces to match their etch factor and process parameters. They also want to apply their own plating thieving and laminate thieving. They also want to simply state the minimum PTH finished hole size and the minimum annular ring that goes with it (following IPC requirements for a 2mil minimum annular ring at all times).IPC (Institute for Interconnecting and Packaging Electronic Circuits)

Historically pcbnew has only handled minimum trace and space. The remaining parameters were ignored.

(R) 26 (pcbnew) *Characters on copper will be checked for the design rules for minimum space and trace. The design parameters will include the default stroke width, default character height, default character spacing and default character font.*

Clearances for plating wicking along fibreglass fibres:

Annular ring calculations:

5.3.4 Copper Layer Design

Copper layer design is the primary focus of all CAD tools for PCB design. It is afterall the copper traces that determine most the electrical characteristics of the board and must form the electrical circuit. Therefore, the discussion here does not discuss all copper design considerations, because many of these are adequately accommodated by KiCad in the first place. The consideration here is of specialized copper features or rules that have not historically been handled by pcbnew and that need consideration.

Teardrops: Designing a copper layer that includes fine-pitch BGAs and other high-density components requires the use of vias with minimal pad sizes to provide additional room for chip escape and routing. When minimal pad sizes are provided, and cards are designed to IPC-2010 Class 1 or 2, it becomes necessary to add teardrops to minimal pads to avoid open circuits where the trace meets the pad, and to provide additional exposed NSMD copper for BGA lands. Otherwise, IPC-2010 Class 3 boards would need to be built, increasing the annular ring requirement and shrinking the available room for chip escape and routing. When board thicknesses are constrained by product standards, simply adding another layer is oftimes not an option. Resorting to BBV (Blind/Buried Via) constructions is too expensive an alternative. Teardrops are a feature in copper that consist of simple artwork and, in fact, reduces the costs due to an increase in yeilds.

The mechanisms and design parameters for teardrops are described in more detail in *Sec. 3.4.2(35)*.

Thieving: Designing a copper layer for fabrication requires that plating thieves be applied to the board in areas isolated from traces so that the current of the plating anode is better distributed, resulting in higher yields for a given resolution (trace and space) demanded by the design. Thieving design considerations are detailed in *Sec. 3.8.1(50)*.

Venting: Designing a copper layer for fabrication requires that resin venting be applied to the board in areas isolated from traces so that a low pressure area is not created during press-out of the B-Stage laminate. Venting design considerations are detailed in *Sec. 3.8.2(52)*.

Characters in copper. pcbnew does not currently consider the design parameters associated with characters etched into copper. However, the etching process is limited in its ability to form characters, just as the screen printing process is limited in its ability to form characters in ink. The primary design parameters associated with character resolution are: minimum line width, minimum character height and minimum character spacing. The minimum line width should be the same or greater than the minimum trace. The minimum character height should be set so that the minimum space rules are not violated by the character font faces. The minimum character spacing should be set to the minimum space. Also, the character font should be selected so as to best form the characters. See *Sec. 7.12(142)*.

(R) 27 (pcbnew) *Placement of characters onto copper layers will observe the design rules of minimum trace and space as applied to characters.*

5.3.5 Copper Layer Imaging

Display The display order is similar to the plotting order described below. When displaying a dark information layer, it is OR'ed with the existing canvas. When displaying a clear information layer it is XOR'd with the existing canvas. When erasing a dark or clear information layer, they are XOR'ed with the existing canvas. One difference to the plotting of copper layers for display is that vias a treated differently for display versus plotting, and so are pads for TH, SMT, or contacts. Also, for diplay of embedded resistors, copper is removed from the display where the resistance ply is exposed, whereas for plots the copper areas to be etched away are included in the plot. Also, holes are displayed differently on the canvas from the plots.

1.

3D-Display

Plotting Both display and plotting of copper layers requires a number of overlays (information layers within the plot). These information layers, from back to front are:

1. *Dark.* At the bottom is copper areas representing copper zones without cutouts.
2. *Clear.* Next is the cutouts for copper zones. Cutouts provide both clearance areas around tracess of a different net, as well as antipads around vias and thermals around same-net through holes.
3. *Clear.* Next is cross-hatch for copper zones. These should not line up with any other dark or clear feature other than the zones themsleves, so it really does not matter in which order this layer is plotted, so long as it is plotted after the copper zone dark layer.
4. *Dark.* Next is all of the traces for the copper layer.
5. *Dark.* Next is all of the (functional) pads for the copper layer. Pads are printed after traces because the apertures could actually have a clear knockout in the center (drill mark). That is, they can be donuts instead of discs. For the same reason, pads are plotted before vias in case the via is a VTP and has a knockout (clear) centre.
6. *Dark.* Next is plating thieves (external layers) or resin vents (internal layers). These should not line up with any other dark or clear feature, so it really does not matter in which order this layer is plotted.
7. *Clear.* Last is all of the clearance holes required to avoid direct drill on ground plane on inner layers, or drill marks for outer layers when requested by the user.

5.3.6 Copper Layer Post-Processing

Teardrops: For annular rings that do not permit a sufficient (IPC 2mil) annular ring to prevent breakout, teardrops (also called “fillets”, “keyholes” or “snowmen”) can be used to permit reduced annular ring and still prevent tangency or breakout. Teardrops comes in several shapes: two shapes of fillet, teardrop, and snowman.

Thieving Thieving is required for high-yield plating of external copper layers when the construction has copper foil on its external layers. Where thieving is required, it can normally be applied by the board fabricator. This causes several problems, however, for high-technology cards: there are areas where the application of thieving could degrade the design electrical properties. Therefore, it is necessary that the fabricator that is to apply thieving be aware of keep-out areas for thieving. pcbnew will therefore provide the ability to plot thieving keep-out and keep-in areas to the fabrication print or detail. See *Sec. 5.24(116)* for more information on plotting keep-out areas to fabrication or assembly drawings.

For plot-and-go services, thieving must be provided by the design. In this case, thieving must be plotted on copper layer plots. The internal representation of thieving and where it is to be plotted is discussed in more detail in *Sec. 3.8.1(50)*. However, the thieving patterns themselves should be plotted to a separate information layer within an RS-274X plot, using dedicated apertures, to permit easy removal or modification of the patterns.

Venting Venting is required for low pressure areas to provide high-yield press-out of internal prepreg layers, and controlling the overall board thickness and the impedance of internal layers. Where venting is required, it can normally be applied by the board fabricator. This causes several problems, however, for high-technology cards: there are areas where the application of venting could degrade the design electrical properties. Therefore, it is necessary that the fabricator that is to apply venting be aware of the keep-out areas for venting. pcbnew will therefore provide the ability to plot venting keep-out and keep-in areas to the fabrication print or detail. See *Sec. 5.24(116)* for more information on plotting keep-out areas to fabrication and assembly drawings.

For plot-and-go services, venting must be provided by the design. In this case, venting must be plotted on copper layer plots. The internal representation of venting and where it is to be plotted is discussed in more detail in *Sec. 3.8.2(52)*. However, the venting patterns themselves should be plotted to a separate information layer within the RS-274X plot, using dedicated apertures, to permit easy removal or modification of the patterns.

5.4 Plating Layers

When flash and button, or button plating buried vias, whether filled or unfilled; prior to planarization: a button plating resist image is required. This image includes only capture pads for PTH on the current layers. Plating images could be generated form information such as NC Drill files and copper layer Gerbers for etch mask.

Historically, pcbnew did not support plating layers. A photoimagable plating resist layer for plating of buried vias during buildup. Buried vias can be button-plated or panel-plated. When panel-plating, photoresist images generated for this layer is not required. When button-plating, resist is applied, imaged with this layer, developed and plated. Buttons are normally removed but can be removed through planarization, potentially after via filling. This is a separate imagable conductive coating layer that is potentially associated with each copper layer (but is only used for outer plated layers). This layer information is for board fabrication.

This technical layer is for imaging plating resist layer for plating of buried vias during buildup. Buried vias can be panel-plated, button-plated or flash-panel and button-plated. When panel-plating, this information layer is not required unless there is selective plating (such as for nickel-gold finish on traces). When button-plating, resist is applied, imaged with this layer, developed, plated and stripped. Buttons are typically removed but can be removed through planarization, potentially after filling buried vias. This is a separate imagable conductive coating layer that is potentially associated with a number of copper layers (but is only used for outer plated layers about a core).

This layer is not a data capture layer (there are no direct inputs to the layer by the designer), but is a post-processing layer used for button-plate photo-resist images.

5.4.1 Plating Process

5.4.2 Plating Materials

5.4.3 Plating DFX Rules

5.4.4 Plating Design

5.4.5 Plating Imaging

5.4.6 Plating Post-Processing

5.5 Hole Filling Layers

These layers are not data capture layers (there are no direct inputs to the layers by the designer and they are not editable, but is a post-processing layer used for hole-filling stencil generation when selectively filling vias.

5.5.1 Hole Filling Fabrication

5.5.2 Hole Filling Process

Hole Filling is a process used to close holes by paritally or completely filling them with a conductive or non-conductive material. There are several reasons for filling holes, and especially via holes:

Thermal Vias: Thermal vias are vias of special construction that are intended to transfer heat from a component into the copper planes in a multilayer printed circuit board. They are used for specialized components that have thermal grounding pads under the component. Texas Instruments, Linear Devices, and Analog Devices have a number of components with this requirement. The construction of thermal vias inexorably include filling the hole of the via with a conductive material (electrically conductive materials also have better heat-transfer characteristics than the usual non-conductive fillings). Thermal vias require a mechanism for filling the via with a conductive filling.

High-Current Vias: Vias designed to carry a high DC current can be constructed without reducing the space available for layout by filling the via with conductive material. Of course the current handling capacity of a conductively-filled via is much higher than that of the same unfilled via. Therefore, conductively-filled vias are far more volumetrically efficient.

Buried Vias: Buried vias are vias that are completed embedded within the a multilayer printed circuit board's inner layers. When a void is at the center of the via barrel, this void can become filled with chemistries, contaminants, or debris, that are trapped within the laminate, where it later causes problems during assembly or, worse, after a period of service.. (Gas could be trapped, however, it is a simply matter to laminate the material together in a hard vacuum thus both avoiding escaping heat and trapped gas.)

Immersion Silver or Tin: IAg (Immersion Silver) or ISn (Immersion Tin) finishes are susceptible to long term reliability issues due to corrosion of the finish material from contaminants or residue chemistries. Because of this characteristic, it is not advisable to use via capping to form a vacuum seal for ICT with these finishes. However, it is possible to plug the via holes altogether using a hole filling process. Either conductive or non-conductive filling can be used. This is really only applicable to through- and blind-vias: that is, where at least one end of the via is on the exterior of the board.[22]

Via Through Pad: Via-through-pad is a construction for vias that permits the via to be placed inside of or near a pad, particularly lands for BGA devices in high-density designs. Because of the need for a highly planar pad surface, it is typical to both fill the via and plate over the filling with 0.0005" of copper. This permits the via construction to go effectively unnoticed to the assembly processes.

ICT Test Points: Via capping with a secondary mask can provide the vacuum seal required by ICT (In Circuit Test). However, some fabricators warn that they cannot guarantee that secondary mask will not bleed or wick through to the uncapped side of the board. Where a via is used on the uncapped side as an ICT test point, this bleeding can impair the function of the test point, causing false negatives. The purpose of vacuum retention of the ICT test fixture is to reduce false negatives. Therefore, the possibility of bleeding must be weighed against the cost of achieving a vacuum seal. ICT test fixtures that a retained by vacuum instead of mechanical means provide reduced false negatives during testing, and thus the objective of capping or filling vias to obtain a vacuum seal. The best way of achieving both a vacuum seal and a reliable test point for ICT is to either fill the via with a conductive material, or to fill the via with a non-conductive material and then plate over one or both ends as for VTP.

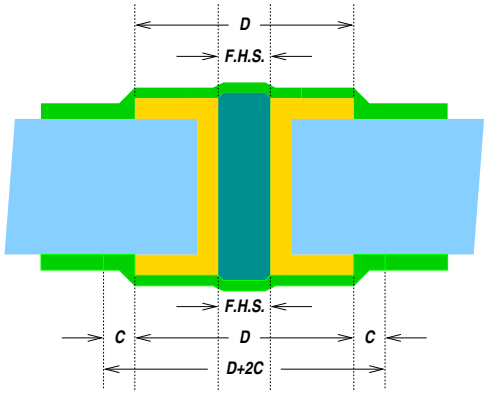

Figure 49: Conductive or Non-conductive Via Filling

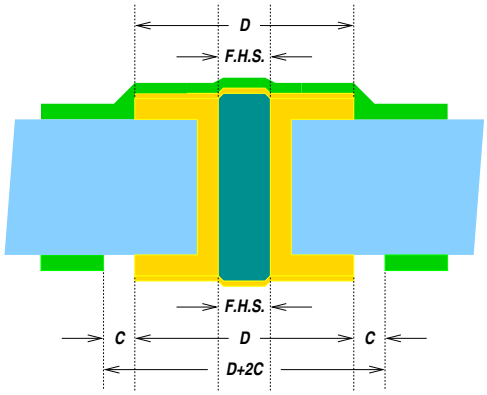

Figure 50: Conductive or Non-conductive Via Filling and Plating

Fig. 49(66) illustrates a through-via that is filled with conductive or non-conductive filling.

Fig. 50(66) illustrates a through-via that is filled with conductive or non-conductive filling and then plated over with copper. Either side (or both sides) can be covered over with primary mask or have primary mask removed. Primary mask is normally removed only when the via is to be used as a test pad.

Fig. 51(67) illustrates a through-via that is plated closed. Plated closed vias have the advantage that they do not require plating over. Normally only blind vias a plated closed, and only vias smaller than a given hole size and aspect ratio. For example, laser-drilled microvias can easily be plated closed. As there are no artwork requirements for plating closed a via, it is sufficient to indicate the vias that need to be plated closed on the fabrication

22. Another option, of course, for IAg (Immersion Silver) or ISn (Immersion Tin) surface finishes is to leave the holes open.

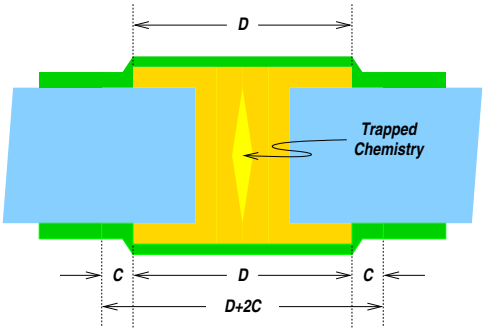

Figure 51: Via that is Plated Closed

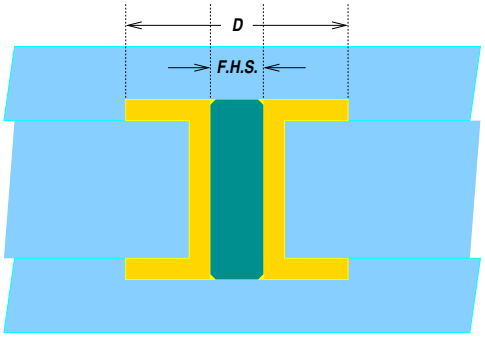

Figure 52: Buried Via: Conductive or Non-conductive Filling

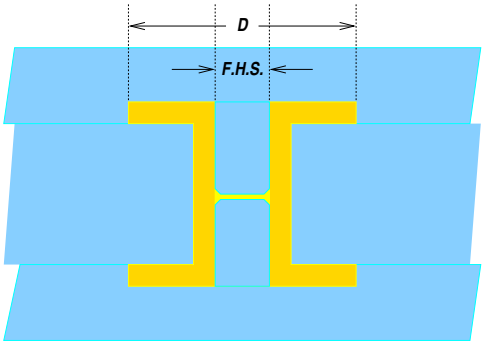

Figure 53: Buried Via: Resin Filled

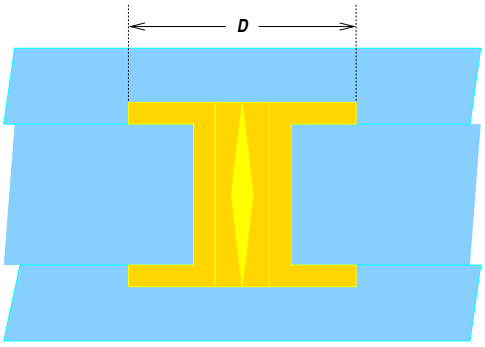

Figure 54: Buried Via: Plated Closed

print or detail.

Fig. 52(67) illustrates a buried via filled with conductive or non-conductive material. Non-conductive filled buried vias are normally filled using button plating followed by stencil printing of an epoxy-based non-conductive filling. Once the filling is cured (typically by heat), both the protruding filling and the button plate are planarized in a brushing or milling machine until the surface of the copper is again flat. Once flat, the normal resist and plasma etching process begins. This, therefore, adds four or five additional process steps as follows:

1. Drill through holes in sub-laminate.
2. Clean, desmear, etchback holes.
3. Electroless copper to seed holes.
4. Flash plate 0.0006" of copper to create a knee.
5. Apply etch resist, photoimage with buttons, develop.
6. Plate 0.001" copper on buttons and in barrel of holes.
7. Stencil print non-conductive epoxy-based fill material into holes.
8. Cure hole filling with heat.
9. Planarize copper buttons and cured hole filling.
10. Photoimage and develop trace patterns.
11. Plasma etch outer layers.
12. Set aside for use in higher order sub-laminates.

Fig. 53(67) illustrates a buried via filled with resin. Resin filling is normally only an option for buried vias, where resin is applied during the lamination process. Note that resin filling is not normally complete, but partial. Use of vacuum during the lamination sequence assists with fill. Percentage fill is never usually 100%. Filling with resin is the usual filling for buried vias.

Resin filling removes several steps from the process of filling via holes. Because it avoids process steps, filling buried vias with resin is less expensive than the other alternatives (with the exception of leaving them empty altogether). Because no imaging is necessary to create the filling (just about any hole in the board at this intermediate stage is only a fabricator tooling hole or a buried via), a note on the fabrication print is sufficient, such as:

```
6. ALL BURIED VIAS ARE TO BE FILLED WITH RESIN AS PART OF
   THE LAMINATION PROCESS. MINIMUM FILL IS 40% IN TOTAL,
   AND 20% FROM EACH SIDE.
```

Fig. 54(67) illustrated a plated-closed buried via. Although this construction is possible, it is unlikely. There is little purpose to intentionally plating closed a buried via. Because no imaging is necessary for plating closed vias, a note on the fabrication print is sufficient, such as:

```
7. ALL BURIED VIAS ARE TO BE PLATED CLOSED.
```

Note that the dielectric constant of the material used to fill the hole is not important. This is because, at high frequencies, skin effect pushes currents to the exterior surface of the barrel of the via. It is far more important what material is around the via, than what material is inside it.

Filling a buried via with a conductive material should only be necessary when the low frequency or DC currents require the volumetric efficiency of a conductively filled via. Other reasons for conductively filled buried vias include vias designed for the transfer of heat between copper layers of the board.

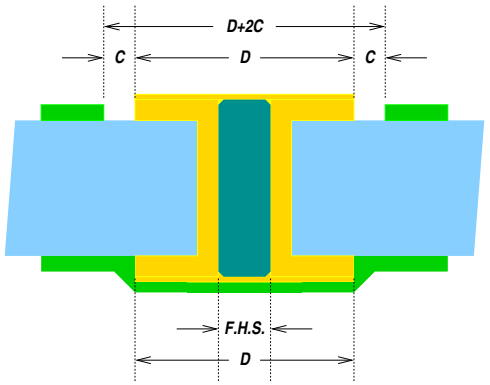

Figure 55: Via Through Pad

Fig. 55(68) illustrates a via-through-pad (VTP) construction. In this case, the via is also a through-via, but in most cases the via is a blind or microvia.

5.5.3 Hole Filling Materials

Choices for hole filling materials include:

Plated Closed. The barrel of the via hole is plated closed with copper. Plate closed holes should be indicated on the fabrication print with a note such as:

```
8. VIA HOLES ARE TO BE PLATED CLOSED WITH COPPER FOR THE
   F.H.S. HOLES PROVIDED IN THE BOARD_HOLEFILL_XXX.GBR
   FILES AND INDICATED ON THE DRILL MAP DETAIL SHEET XXX.
```

Non-conductive. The barrel of the via hole is filled with non-conductive hole filling material, and possibly plated over with copper. Non-conductive filling of via holes should be indicated on the fabrication print with a note such as:

```
9. PROVIDE VIA FILL WITH PETERS PP2795-SD OR EQUIVALENT
   NON-CONDUCTIVE MATERIAL FOR THE F.H.S. HOLES PROVIDED
   IN THE BOARD_HOLEFILL_XXX.GBR FILES AND INDICATED ON
   THE DRILL MAP DETAIL SHEET XXX.
```

Conductive. The barrel of the via hole is filled with a conductive material. Conductive filling of via holes should be indicated on the fabrication print with a note such as:

```
10. PROVIDE VIA FILL WITH DUPONT CB100 OR EQUIVALENT
    CONDUCTIVE MATERIAL FOR THE F.H.S. HOLES PROVIDED
    IN THE BOARD_HOLEFILL_XXX.GBR FILES AND INDICATED ON
    THE DRILL MAP DETAIL SHEET XXX.
```

Resin. The barrel of the via hole is filled with laminate resin. This is indicated on the fabrication print with a note such as:

```
11. ALL BURIED VIAS ARE TO BE FILLED WITH RESIN TO A
    MINIMUM OF 80% FILL.
```

5.5.4 Hole Filling DFX Rules

In general, DFX checks are unaffected by hole filling and hole filling materials.

Nevertheless, there are DFX rules and modifications to existing rules for VTP (Via Through Pad) and VBP (Via Beside Pad) constructions, as follows:

- VTP vias larger than a given hole size must be filled and plated over with copper.
- VBP vias larger than a given hole size must be filled.
- VTP vias of greater than a given aspect ratio must be filled and plated over with copper.
- VBP vias larger than a given aspect ratio must be filled.
- VTP and VBP vias are not subjected to the normal clearances from pads of the same net.

To perform these rule checks, the following process parameters must be known:

Maximum unfilled via F.H.S. This is the maximum finished hole size that a via may have before it must be filled with a conductive filling, or filled with a non-conductive filling and plated over, or plated closed.

Maximum unfilled via aspect ratio. This is the maximum aspect ratio that a via hole may have before it must be filled with a conductive filling, or filled with a non-conductive filling and plated over, or plated closed.

5.5.5 Hole Filling Design

5.5.6 Hole Filling Imaging

Plotting: Where hole filling is performed by stencil, a stencil mask may be necessary. When all PTHs are filled and are filled with the same material, the filled holes correspond to the corresponding drill map. Therefore, some fabricators simply want the holes identified on the drill map, and the fabricator can create the stencil from the NC drill files. Where not all holes are filled, or they are not filled with the same material, the holes and which material they are filled with need to be identified. The easiest way to identify the holes is to provide Gerber or Bacro plots of apertures that correspond to the F.H.S. of the holes to be filled, and a separate plot for each filling material. An exception is where buried vias are filled with resin as part of the lamination process. A stencil is not necessarily required. Also, it is a simple thing to identify on the fabrication print that all buried vias are to be filled with resin. Nevertheless, if a hole-filling plot is generated, the designer does is not forced to include it in the package.

Therefore, the best way to handle hole-filling is likely as follows:

- Use a separate tool number in NC drill files and identify tool numbers that represent filled holes and their filling on NC drill maps and reports. Whether filling is selective or non-selective should be obvious from the drill report.
- Plot apertures the same size as F.H.S. on a plot for each drilling layer-pair identifying the filled holes. Provide a separate plot for each material and each drilled layer-pair.
- Identify the data files and filling material on the fabrication print and detail sheets (drill maps).
- Identify the Gerber or Barco files and their contents in the README.

Display: To assist the designer with assigning filled holes, hole filling and material should be identified on displays as a separate visible item. Normally, holes on the pcbnew display have been displayed as a black hole set to the F.H.S. of the hole. When holes are filled, they can be displayed, instead, as a coloured disc. The appearance of the disk should also be different depending on the filling material (non-conductive, conductive, resin).

Editing: It should be possible to change hole-filling attributes of vias using right-click or properties dialog.

5.5.7 Hole Filling Post-Processing

5.6 Keep-out Layers

Historically, pcbnew did not directly support keep-out layers. Silkscreen outlines for components provided some guide as to component-to-component placement (courtyards); however, keep-ins and keep-outs were not described.

Keep-out layers are non-imagable technical layers. That is, they have some meaning to the layout process, but, although necessary for DFX rule checks, are not necessary for fabrication output. That is: they are not post-processed (unless simply to diagram and annotate them on assembly drawings).

There are two types of entities in this layer: keep-ins and keep-outs. Keep-ins and keep-outs can pertain to components (placement keep-ins and keep-outs), tracks (routing keep-ins and keep-outs), vias and TH pins (hole keep-ins and keep-outs), zones (copper keep-ins and keep-outs) and venting or thieving (thieving keep-ins and keep-outs). Keep-ins and keep-outs related to the placement of copper features on the card have no height.

Coper feature (routing, thieving) keep-ins and keep-outs pertain to one or more copper layers. Because copper feature keep-ins and keep-outs can be present in a library module, definition without knowledge of the number of copper layers, their specification of layers can select board-side(s0 and inner layers.

Historically for autorouting within pcbnew, the board edges were used as the keep-in boundary for routing. DRC rules were applied to determine whcih regions within the board edges routing cannot occur.

Component keep-ins and keep-outs have a height associated with them. A keep-in height describes the height beneath which a component will fit within the keep-in region. A keep-out height describes the height above which a component must be excluded from the region. A keep-out height of zero specifies that all components must be excluded from the region for the board side.

Historically for autoplacement within pcbnew, the board edges were used as the keep-in boundary for placement. DRC rules were applied to determine which regions within the board edges routing cannot occur. The placement boundary box of the component was the bounding box of the items on board fabrication layers.

These layers are data capture layers: the designer directly edits keep-out and keep-in contours on the layer. These layers are also post-processed layers (but only in the generation of electrical and mechanical CAD models for input/output formats such as IDF. [Kehmeier and Makowski, 1998]

5.6.1 Keep-out Process

The keep-out process is primarily concerned with board fabrication and assembly (including design, placement, autoplacement, routing and autorouting), board test, board assembly, rework and test. It defines keep-in and keep-out areas that describe the mechanical, and somewhat the electrical, rules for the placement of components and board features. The layer is intended on capturing a simple set of rules that restrict the placement of components and board features. This simple set of rules is then applied as part of a DRC (Design Rule Check) that is performed either off-line, when checking a design for its adherence to the rules, or on-line, while editing the design. This set of design rules can also be imported or exported in various file formats so that the rules can be exchanged with other PCB design systems, whether electrically focused (ECAD), mechanically focused (MCAD), or manufacturing focused (CAM). Therefore,

(R) 28 (pcbnew) *The internal representation of keep-out and keep-in zones in pcbnew must support the models provided by the intended exchange file formats: GenCAM [GenCAM], GenX [GenX], 258X [258X], IDF [Kehmeier and Makowski, 1998] and SPEECTRA DSN [SPEECTRA].*

IDVv3.0 [Kehmeier, 1996] models components using a single component outline (closed contour) and a height, where the height is determined at a zero mounting offset. Each placement of a component specifies a mounting offset. The format does not support specifying component-attached keep-outs. IDFv4.0 [Kehmeier and Makowski, 1998] models components using one or more extrusions. Footprints can contain zero or more keep-outs. GenCAM and GenX [GenCAM, GenX] models packages using body artwork on a technical layer and provides: a package height, the maximum height of the package as measured from the finished surface of the board or panel on which it is mounted; a package standoff, the clearance height to the bottom of the package body as measured from the finished surface of the board or panel on which it is mounted; and, a package centroid, the centroid of the package as used by manufacturing equipment during assembly, inspection and test, measured from the origin of the package. GenCAM components can contain zero or more component-attached keep-outs; however, keep-outs for components do not specify any restrictions (all components or no components). GenX associates 258X [258X] models packages with a height from the mounting surface and then a placement describes the maximum protrusion the upper surface (mounted height) and a standoff (at mounted position). It does not refer to the position. This CAM format does not model keep-outs or keep-ins at all.

IDFv4.0 [Kehmeier and Makowski, 1998] defines keep-in areas as follows: The keep-in entity is used to represent an area on the board or panel in which component instances or board features must be located. There are four predefined keep-in types, that are described as follows:

All components. An area in which all component instances must be placed. The keep-in may affect any valid component placement layers. The keep-in has an optional height to specify a maximum placement height value. All components with a molded height less than this value must be placed within the area. If a height is not specified, all component instances must be placed within the area.

Specific components. A component group area within which a specific set of component instances must be placed. The keep-in may affect any valid cmponent placement layer or layers. The keep-in uses the group property to specify a list of component references that are to be placed within the area.

Component rotation. A component rotation area within which component instances must be placed according to a set of predefined orientations. The keep-in may affect any valid component placement layer or layers. It specifies a set of acceptable rotation values for the component placement direction (0 and 180 degress, for example).

All routing. An are in which all routing must occur. The keep-in may affect any valid routing layer or layers.

IDFv3.0 [Kehmeier, 1996] defines three keep-in types as follows:

Routing outline. This section of the file defines a routing outline for the board or panel. Each routing outline specifies a region within which routing must be confined, and consists of a simple closed curve made up of arcs and lines. Portions of routing outlines on a panel that lie on a board in the panel are inherited by that board. Multiple routing outlines may be defined.

Placement outline. This section of the file defines a placement outline for the board or panel. Each placement outline specifies a region within which components must be placed, and consists of a simply closed curve made up of arcs and lines plus a height restriction. Portions of placement outlines on

a panel that lie on a board in the panel are inherited by that board. Multiple placement outlines may be defined. The outline height is used to exclude components from the outline that, when mounted, exceed this height. If this field is missing, there is no height restriction on the outline.

Placement group area. This section of the IDF version 3.0 file specifies an area where a group of related components is to be placed. For example, it may be desirable to place all analog components in a particular area for thermal considerations. Each placement group area consists of a simple closed curve made up of arcs and lines along with a name designating the group of components to be placed in that area. Multiple areas are allowed.

IDFv4.0 [Kehmeier and Makowski, 1998] defines keep-out areas as follows: The keep-out entity is used to represent an area on the board or panel in which component instances or board features may not be located. There are seven predefined keep-out types that are described as follows:

Components by height. An area in which no component instance having a mounted height greater than the specified keep-out is allowed. The component mounted height is equal to the maximum height of the shape of the component part plus the mounting offset of the component instance. The keep-out may affect any valid cmponent placement layer or layers. It specifies a maximum placement height value, or when not specified (or zero), the restriction applies to all component instances, regardless of height.

All routing. An area in which no routing is permitted. The keep-out may affect any valid routing layer or layers.

All traces. An area in which no traces are permitted. The keep-out may affect any valid routing layer or layers.

All vias. An area in which no vias are permitted. The keep-out may affect any valid routing layer or layers.

All testpoints. An area in which no testpoints are permitted. The keep-out may affect top, bottom or both layers.

All silkscreen. An area in which no legend text or graphics are permitted. The keep-out may affect top, bottom, or both layers.

All holes. An area in which no holes are permitted. The keep-out affects all layers.

IDFv3.0 [Kehmeier, 1996] defines three keep-out types as follows:

Routing keep-out. This section of the file defines a routing keep-out for the board or panel. Routing keep-outs specify regions where routing is not allowed. Routing keep-outs can exist on top, bottom, both top and bottom, or all routing layers. Each keep-out consists of a simple closed curve made up of arcs and lines. Portions of routing keep-outs on a panel that lie on a board in the panel are inherited by that board. Multiple keep-outs are allowed.

Via keep-out. This section of the IDF version 3.0 file defines a via keep-out for the board or panel. Via keep-outs specify regions where vias are not allowed (although routing is still allowed). Each keep-out consists of a simple closed curve made up of arcs and lines. Portions of via keep-outs on a panel that lie on a board in the panel are inherited by that board. Multiple via keep-outs are allowed. Only through vias (vias that go all the way through the board) are supported.

Placement keep-out. This section of the IDF version 3.0 file defines a placement keep-out for the board or panel. Placement keep-outs specify the regions of the board where components cannot be placed. A keep-out can apply to all components, or to only those components above a specified height. Placement keep-outs can exist on the top, bottom, or both top and bottom of the board or panel. Each keep-out consists of a simple closed curve made up of arcs and lines along with a height restriction. Portions of placement keep-outs on a panel that lie on a board in the panel are inehrited by that board. Multiple keep-outs are allowed. A keep-out height is used to exclude components from the keep-out that, when mounted, exceed this height. A value of 0.0 indicates that all components are to be excluded.

GenCAM [GenCAM] defines keep-out areas as follows:

Components A component placement keep-out is one that is associated with a specific component instance (placement, mounting location). It defines an area that is not accessible to test or insertion machines. The keep-out is described as a single extrusion closed contour; however, zero or more keep-outs can be defined for each component (placement, mounting location).

This defintion is more closely related to the concept of courtyards or component-level keep-outs, where a component-level keep-out is one that moves with the component rather than being fixed to the board; however, the GenCAM keep-out is attached to a *specific placement*. Nevertheless, when writing a GenCAM or GenCAM-XML file is is possible to replicate a component-level keep-out to all of the placements, so it can be modelled with a component-level keep-out.

This component-attached keep-out can be modelled in pcbnew with a zone on the keep-out layer that is defined as part of a module using the module editor.

Boards A board keep-out is one that is associated directly with the board (does not require a placement). It defines an area that is inaccessible.

IDFv3.0 does not really distinguish between component-attached keep-out and board level keep-outs. Under pcbnew this keep-out can be modelled as a zone on a keep-out layer that belongs to a board.

Panels It defines an area that is not accessible. A panel keep-out is one that is associated directly with the panel (neither requires placement of a component nor placement of a board). It is defined relative to the panel coordinates.

This can be modelled in pcbnew, once arrays and panels are supported, as a zone belonging to an array or panel that is defined on a keep-out layer.

Fixtures A fixture keep-out is one that is associated directly with a fixture. It define an inaccessible area of the fixture. It is defined relative to the origin of the fixture.

This can be modelled in pcbnew, once fixtures are supported, as a zone belonging to a fixture that is defined on a keep-out layer. For now, these zones are defined for all fixtures associated with the one-up data for the board, array or panel. Fixtures could possibly be handled by defining an fixture identifier that is associated with a given keep-out or keep-in zone.

First and foremost, the keep-out/in provides a component-specific, or board level, panel level, or fixture keep-out or keep-in that can be one of several kinds:

Component

Via

Route

Test Pin

Test Probe

Board

The keep-out can also apply to one or more layers: top, bottom, inner. The sense of the keep-out can be reversed so that is describes a keep-in.

The assembly process depends on the placement and orientation of components. In its simplest form, a table containing the x and y coordinates of the components and fiducial marks are used as input to the assembly automatons. This simple 2-dimensional view causes problems when considering heat-sinks and other devices mounted beneath and over other components. A mounting heigh is required as well, and either the placement must be properly ordered so that components underneath other components are populated first, or a way of describing *what* is underneath *what else*.

Some keep-outs should be relatively automatic, as follows:

Solderpaste fixture keep-outs. On the secondary side of the PCB, the long edge of the PCB muse have a clearance of 0.300" on both sides. No SMT components should be placed there. Through-hole components that are to be hand-soldered or waved are okay. The reason is that the solder screening machine has hold-downs which attach to the secondary side of the board to prevent the board from moving during the screening process. These hold-downs may damage any component in this area as the past is screened onto the PCB.

Assembly fixture keep-outs. A perimeter keep-out is required for assembly fixtures. The standard assembly fixture keep-out is 0.125" on two sides (long side) of the board. This is primarily for the conveyor that transports the card through the assembly line. The leading edge of the conveyor may also need a 0.125" keep-out for drawing the card through the conveyor.

SPEECTRA DSN [SPEECTRA] defines keep-outs a little different. Keep-outs describe an areas where certain autorouter rules are to be applied or not applied. It is not so focused on placement, but more on routing refinements. The types of keep-outs are:

Place keep-out. This is a simple component keep-out with an optional height (to keep out all components that are greater than the specified height). When the height is specified as zero, it means that all components are kept out of the area. This is the standard component keep-out supported by other formats.

Via keep-out. This is a keep-out for vias but does not keep routing off of the layers to which the keep-out applies. This is the standard via keep-out supported by other formats.

Wire keep-out. I think they mean traces here. Yes, this is the standard routing keep-out supported by other formats.

Bend keep-out. Hmmm. Can't bend. What does that mean: must be a straight line? This is similar to Elongate keep-outs, below, in that this is something for an autorouter alone and has little to do with anything that can be checked by applying simple zone rules. Therefore, it will be added as a zone attribute, but there will be no DRC checks nor enforcement for such a keep-out. It is just something that will be passed through to an autorouter in a SPEECTRA DSN file.

Elongate keep-out. I assume an elongate keep-out is one where the autorouter is not permitted to place an accordian (serpentine) structure for the purpose of lengthening a trace to match it with another trace. I don't know when one could predict the need for this, but if you look at some autorouted boards, the traces appear worse than the ratsnest that they connected. There seem to be no rhyme or reason to their paths and they go every which way and another. I don't know what this was intended to address, perhaps appearance, but there must have been a reason at one time. I can't see how a designer could possibly know where or when to specify such a keep-out, so I don't see a need to support it. Nevertheless, it is just one more attribute on a zone, so we can do it; however, there will be no DRC checks nor enforcement for such a keep-out because the program cannot detect elongations. It is just something that will be passed untouched to an autorouter in the SPEECTRA DSN file.

SPEECTRA DSN also defines a place boundary that defines an area of the PCB that permits component placement and a signal boundary that defines an area of the PCB in which signal routing can be performed. These are equivalent to component and routing keep-ins described by other formats. The placement keep-ins can also specify the component types (pin, SMD or area) where pin components are TH components, SMD is SMT devices, and areas are general keep-outs or placement keep-outs. A session structure descriptor contains a place boundary descriptor and a number of keep-out descriptors for the session. Keep-out areas can be modified with window descriptors. Window descriptors subtract from an area and are equivalent to cutouts. All keep-out areas can have cutouts. Keep-outs can specify the layers or layer types or layer aliases to which they apply. It appears that all keep-outs are board-level areas: that is, they are fixed to the reference coordinates of the board and are not movable. SPEECTRA DSN also has a concept of rooms which is a general placement keep-in concept. Room can have rules that restrict the minimum and maximum height, power dissipation restrictions, or power net. Specific components or component groupings can be included in a room or excluded from a room. Inclusion and exclusion from room can be hard or soft, where hard means that the entire component or cluster must be included or excluded from the room, where soft means that a portion of the component or cluster can remain included or excluded from the room. Rooms can describe placement restrictions on large, small, discrete or capacitor components, also by pin type (TH or SMD). Also permitted orientation, placement side and opposite side component restrictions can be applied.

5.6.2 Keep-out Materials

The keep-out layers are design layers and, for the most part, are not post-processed.[23] Therefore, keep-out layers do not have any directly associated materials.

5.6.3 Keep-out DFX Rules

For keep-outs there are no specific DFM or DFX rules. The rules that are to be followed *are* the design parameters for the keep-out zone.

5.6.4 Keep-out Design

5.6.5 Keep-out Imaging

Keep-out and keep-in layers sometimes need post-processing when the board fabricator is relied upon to modify the design before fabrication. In this case, the fabricator should know how not to modify the design.

For example, where plate thieving and resin venting is provided by the board fabricator, and there are critical areas on the board that must not contain thieving or venting, the keep-out areas need to be defined and communicated to the board fabricator

23. Exchanging keep-out zones with IDF (Intermediate Data Format) considered.

in the fabrication print or detail sheets. Therefore, it should be possible to map specific keep-out areas onto specific fabrication or assembly prints or detail sheets. This is considered in more detail in *Sec. 5.24(116)*. An example of a keep-out that should be placed on assembly drawings is the standard SMEMA transport clearance [SMEMA 1.2].

5.6.6 Keep-out Post-Processing

The only real post-processing preformed with keep-out and keep-in layers is a final DFX check. As keep-outs and keep-ins are related more to thermodynamics and physical assembly that anything else, a board fabricator that is not also an assembler will likely never catch a mechanical problem until component assembly or even final assembly. Thermal problems will likely not be caught until the product undergoes field or pilot testing, or may not show themselves until the product has been in service for a number of years.

Although keep-out and keep-in areas are not normally plotted to design files, they can be placed on fabrication or assembly drawings to support callouts, special instructions, or to provide documentation of the design. Because of this, the keep-out or keep-in may be plotted to a drawing layer as detailed in *Sec. 5.24(116)*.

5.7 Via Plugging Layers

These layers are not data capture layers: they are not directly edited by the designer, but are post-processed for use in selectively capping vias.

Via Plugging. The standard LPI solder mask process cannot tent or fill vias without the risk of exposed copper inside the hole barrel. Typically, a secondary screen print operation is used that deposits UV (Ultraviolet) or thermally curable epoxy solder mask into the holes to plug them. This is called via plugging. Via plugging is used to plug via holes with a solder resist material to prevent air leakage during ICT (In Circuit Test) or to prevent shorting from components that are close to the board surface. Typically only one side of the via can be plugged. There is a high probability that one side of the plug will be blown off during assembly if both sides of the same via are plugged due to entrapped gasses. [Fabricators prefer] to plug the solder (bottom) side of the board. The solder mask opening must be made larger than the hole to allow for reliable via plug deposition into the hole.

When via plugging is required, the [fabrication] drawing should indicate this and from which side. If select vias are required to be plugged, a separate design file should be generated by the designer, which designates which vias are to be plugged. This separate design file should be labeled in such a way as to indicate that it is a via plug file and from which side it should be used (i.e. "Component Via Plug" for component side via plugging).

If select vias are required to be plugged on one side while different vias are required to be plugged from the opposite side, then two separate design files should be generated to accomplish this. It is highly recommended that no via be plugged from both sides. This process should be indicated on the fabrication with a note stating to *"Provide Standard Via Plug for the 0.0xx" holes from the [solder/component] side using SR1000 or equivalent material."*.

100% Via Plugging. This requirement uses a similar process as described above but the entire hole barrel is filed with a non-conductive material using SR1000 or equivalent. The exact hole fill percentage is dependent upon the hole size and board thickness. It will generally fill 80-90% of the volume. This should be indicated on the fabrication print with a note stating *"Via Plug > 80% of the hole barrel volume for the 0.0xx" holes with SR1000 or equivalent material"*.

BGA Device Via Plugging. BGA (Ball-Grid Array) devices require special consideration in via plugging due to their tight spacing. The figure below describes the different dimensions required for plugging vias associated with BGA devices. The surface finish for most BGA device applications is ENIG (Electroless Nickel/Immersion Gold), IAg (Immersion Silver), ISn (Immersion Tin) or OSP.[24]

Via plugging should not be used with the IAg (Immersion Silver) surface finish. The via is added after the surface finish. Via plug that uses a thermal cure process should not be used because it tarnishes the surface finish and may affect solderability.[25] Only a UV cured via plug should be used. Via plugged vias should not be used with the OSP, IAg or ISn surface finishes because entrapped chemistry may cause via hole corrosion.[26] Vias should be 100% plugged or left open (left open is preferred).

The fabrication print should have a note specifying the via coverage rules. An example note is:

```
12. THE XXX-VIAS (SPECIFY VIA SIZE, ESPECIALLY IF MORE
    THAN ONE IS USED ON THE BOARD) WERE DESIGNED WITH
    SOLDERMASK ENCROACHED ON THE VIA PADS, MANUFACTURER
    CAN ADJUST THE SOLDERMAKS-CLEARANCE-SIZE PER
    THEIR PROCESS TO ASSURE VIA-HOLES ARE OPEN FOR
    IMMERSION-SILVER FINISH, WHILE KEEPING THE PADS
    ENCROACHED WITH SOLDERMASK AS MUCH AS POSSIBLE.
```

...

Via Plug Soldermask Guidelines.

- A separate design file should be provided when select vias are to be plugged. This file should be labeled appropriately to designate form which side the vias are to be plugged.
- Plugging vias with LPI soldermask is not encouraged. [Fabricators] only [have] this capability at a few facilities. This can be replaced with a secondary via plug.
- Holes may only be plugged from one side of the via. Both sides should ont be plugged because air will get entrapped in the middle of the via. When the PCB is heated during reflow, or rework, the gas may expand enough to blow one sid of the plug out and knock off components or create tomb stoning.
-

More requirements follow:

1. Holes to be plugged must be specified on the drill chart to prevent errors.
2. Component, Solder, or both sides can be plugged. Plugging is not recommended for covering both sides of the via on a board; use filled holes instead. Air or moisture could become trapped inside causing blow-out during assembly [reflow].
3. With holes $\leq$15mil, simply screening LPI over the via, the via will "tent." However, after assembly (thermal stress, washing, vibration, etc.), there is no guarantee that the tent will remain intact. PCB manufacturers recommend either plugging or filling the vias. Delamination of the solder mask may result. Also, PCB vendors indicated that chemicals could become trapped in the tented via causing reliability problems and corrosion later.
4. There as some differences in terminology. Some vendors say "plugged" to mean one or both sides of the hole are plugged with epoxy. Some make the distinction that "plugged" is where only one side of the via is covered with epoxy; whereas "filled" is where both sides of the via are covered with epoxy.
5. Plugged holes:
   (a) Minimum hole diameter is 8mil. The smaller the hole, the more difficult it is for epoxy to flow into the hole.
   (b) Maximum hole diameter is 20mil. The larger the hole becomes, the more difficult it is to maintain 100% coverage.
   (c) PCB vendors cannot guarantee that solder resist will not seep onto the other side of the board through the via. Therefore, this should not be used for covering the component side vias where access to the secondary side is required (as in ICT). Use filled vias with conductive epoxy for such applications.
6. Filled holes (an even more expensive option):
   (a) Minimum hole diameter is 10mil.
   (b) Maximum hole diameter is 18mil.
   (c) Conductive or non-conductive epoxy can be used. Conductive epoxy can be used for areas requiring the component side to be covered but the solder side vias accessible for test points.
7. Possible alternatives to plugging is to use dry film to tent the vias on one side. Very few PCB manufacturers use dry film so it may be difficult to find such a manufacturer.

24. This is due to the thinness and planarity of these finishes.

25. Actually, studies of tarnish on IAg finishes have shown that the tarnish does not affect solderability: it only affects the asthetic appearance of the finish (the tarnished finish looks dirty).

26. Huh? How can a solution of Silver cause corrosion?

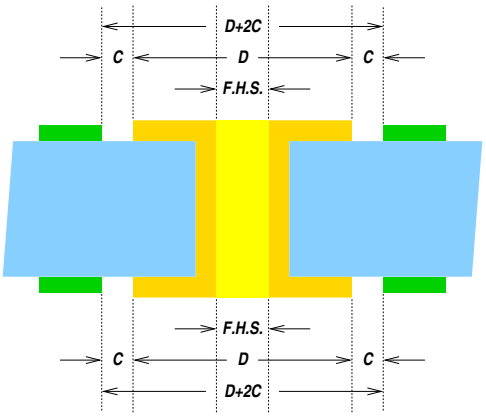

Figure 56: Via Capping: Clearance

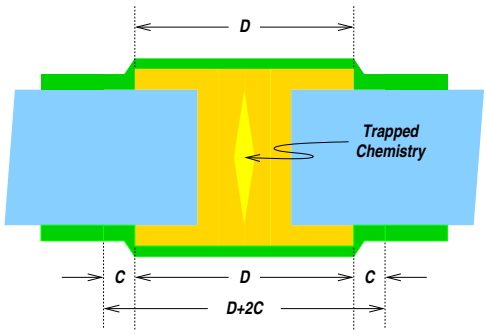

Figure 57: Via Capping: Plate Closed

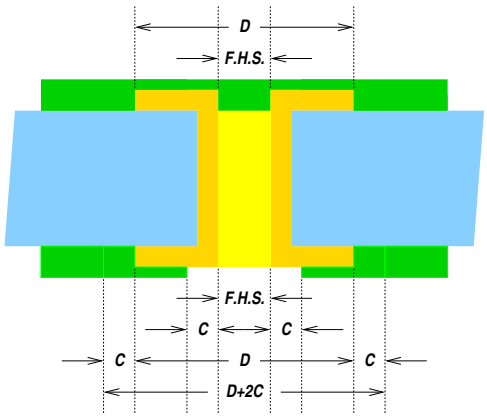

Figure 58: Via Capping: Dry-Film Tenting

5.7.1 Via Plugging Process

Via Plugging (or capping) is a process used to close one side of a via hole. The most common form of via plugging is the application of a secondary mask to form via caps that tent or plug the holes. There are two ways of applying

Fig. 56(74) illustrates a via where primary mask has full clearance, C, from the via pad size, D, on both sides of the board. Full clearances are necessary for vias that are test points on the tested side of the board. When surface finish is HASL (Hot Air Solder Leveling), the via may become plugged with solder from the HASL process. Flux can be applied to the holes to encourage plugging; however, it cannot be guaranteed that holes will be plugged by HASL. The primary mask opening is either plotted 1:1 as the pad side with instructions on the fabrication print to adjust for primary mask clearances; or is plotted as the via pad expanded by the primary mask clearance.

Fig. 57(74) illustrates a via plated closed. The via is depicted with both ends covered with primary mask. When either side of the via is defined as a test point, the primary mask opening is the pad size, D, plus twice the primary mask registration clearance, C, and the pad will be completely exposed as an ICT test point.

For non-HASL finishes, it is possible to alter the plating process to plate vias closed, however, due to void in the center of the via that can trap the plating chemistry, this is not advised in most cases. Primary mask can either cover the resulting via, or will be left with a full clearance for any vias that are used as test points.

Plate-closed vias can be fabricated by button plating until closed and then removing the buttons with brushing or other planarization. The approach can be used successfully for buried or through vias. When used on through vias, it adds processing steps. When used on buried vias, the approach can have less steps than conductive or non-condutive via filling.

Fig. 58(74) illustrates a via is tented over using dry-film solder mask. Dry-film solder mask has the ability to provide 100% tenting of via holes of up to a maximum size. Again, to avoid explosive out-gassing during reflow cycles, it is not advisable to tent both sides of the via. Therefore, the opposite side of the via either has a primary mask hole clearance, or a primary mask pad clearance (especially, where the pad also serves as a test point).

Via capping should not be performed on IAg (Immersion Silver) or ISn (Immersion Tin) finishes, because chemistries trapped against these finishes in the hole barrel can lead to corrosion and the attendant reliability issues. For IAg or ISn, leaving the hole open and using encroachment is a better solution to BGA fields and rework, whereas, for vacuum seal, filling vias is better than capping them because filling does not leave a void in the barrel.

Via capping should not be performed to achieve a vacuum seal for ICT where the fabricator cannot guarantee that the secondary mask will not bleed through to the other side of the board. When this guarantee cannot be made, holes should be plated shut, filled with a conductive material, or filled with a non-conductive material and plated over.

Fig. 59(75) illustrates a via cap applied to a via hole where the secondary mask is applied after the primary mask. Performing this sequence requires that the capping side have a primary mask opening equal to the finished hole size (F.H.S.) plus the primary mask clearance to avoid tenting the hole. The primary mask on the non-capped side must have either a clearance to the F.H.S. when encroached, or clearance to the pad when used as a test point. The secondary mask must define a button that is the pad size plus the secondary mask clearance.[27]

One difficulty with applying secondary mask after primary mask is the height and width of the resulting cap. The width of the cap is a mask clearance larger than the via pad and the height of the cap is the thickness of the primary and secondary masks. This excess width can cause interference with BGA solder balls and lands when the cap is placed under a BGA field. Therefore, when vias are capped and placed under a BGA field, the secondary mask is typically applied *before* the primary mask as shown in *Fig. 60(75)* and described below. Note that so much overlap is necessary when the secondary mask is applied second because the registration of the secondary mask to the primary mask is not so good. An additional clearance is required from the edge of the cap to the edge of any primary mask opening to

27. When LPI mask is used for both primary and secondary, the clearance is typically the same for both masks.

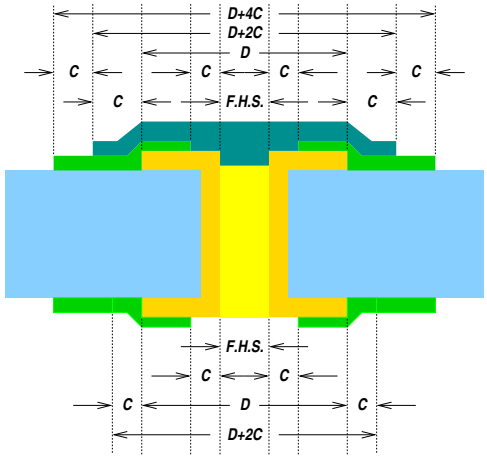

Figure 59: Via Capping: Secondary Mask (Second)

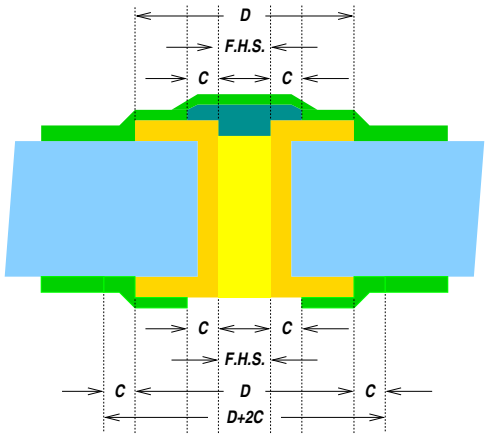

Figure 60: Via Capping: Secondary Mask (First)

ensure that the cap seals against the primary mask. This means that the total size of the capping structure, $D'' = D + 4C$. For a pad size, D, of 0.020" and a primary and secondary mask clearance, C, of 0.003", this arrange is 0.032" in size! This is far too large to fit between BGA pads in typical BGA pad fields. Therefore, this approach is normally only used with HASL finishes (ENIG and ENIPIG finishes are compatible, but these are only normally used with CBGA or other components requiring a high planarity). However, when used with BGA, this approach is typically only used to cap the bottom (solder) side of the board.

Fig. 60(75) illustrates a via cap applied to a via hole where the secondary mask is applied before the primary mask. The size of this arrangement is D, the size of the via pad, which is significantly less than the resulting cap when secondary mask is applied *after* the primary mask as above. With initial secondary mask application, clearances can be calculated just from the pad dimensions. Also note that the bump created by the cap is far narrower and therefore will not interfere with the solder balls on the BGA package.

This capping approach is compatible with HASL, ENIG, and ENIPIG finishes. Due to the planarity requirements of CBGA, this capping approach is most typically used with ENIG or ENIPIG finishes.

Fig. 61(75) illustrates a via where primary mask is encroached on both sides of the board. Note that simple encroachment the

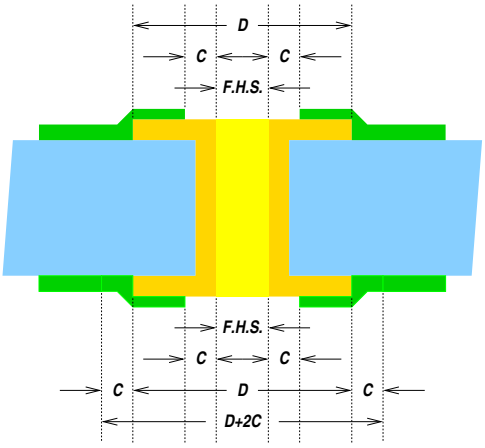

Figure 61: Via Capping: Encroachment

narrowest pattern underneath BGA fields (but secondary mask before primary mask is just as narrow). Clearances and minimum web between the via and BGA land can be calculated based simply on the copper patterns. Also, simple encroachment has no issues with height. Either side of the via (usually the bottom or solder side) can be opened up as a test point (even though the vacuum seal will not be present to support vacuum-held ICT fixtures).

Because of the long-term reliability issues associated with capping vias with IAg (Immersion Silver) or ISn (Immersion Tin) finishes, encroachment is the only option for these finishes unless 100% via filling is performed.

Note that back-drilling is compatible with via capping. This is because the back-drilling is performed before the application of primary or secondary mask. *Fig. 62(76)* illustrates a dry-film tented via hole that is also back-drilled. Also, the tenting, capping or plugging can be applied to either side of the back-drilled hole as illustrated in the diagrams. The only problem with via back-drilling is that at most one side of the via is available as an ICT test point; and, when the capture pad side of the back-drilled via is primary mask encroached, or capped, the via is unavailable as an ICT test point (just as is a buried or blind via of similar construction).[28]

5.7.2 Via Plugging Materials

5.7.3 Via Plugging DFX Rules

5.7.4 Via Plugging Design

5.7.5 Via Plugging Imaging

When via are tented with primary mask (e. g. dry-film), or when encroachment only is used, no additional information is necessary. That is, primary mask image are sufficient. However, when vias are capped or plugged with secondary mask, either before or after primary mask, the secondary mask is also photoimagable and a via capping Gerber or Barco layer is required to specify the via caps.

Via capping Gerber or Barco data must specify the size of the via cap. When primary solder mask is not fully specified by the designer (that is, it is plotted 1:1 so the fabricator can adjust), then the apertures plotted on the Gerber or Barco files is the pad size for the plugged side (secondary mask applied *after* primary mask), or the F.H.S. (secondary mask is applied *before* primary mask), and only on the side that is plugged.

28. This might not be a significant issue because most BGA breakout vias are for ground or power pins.

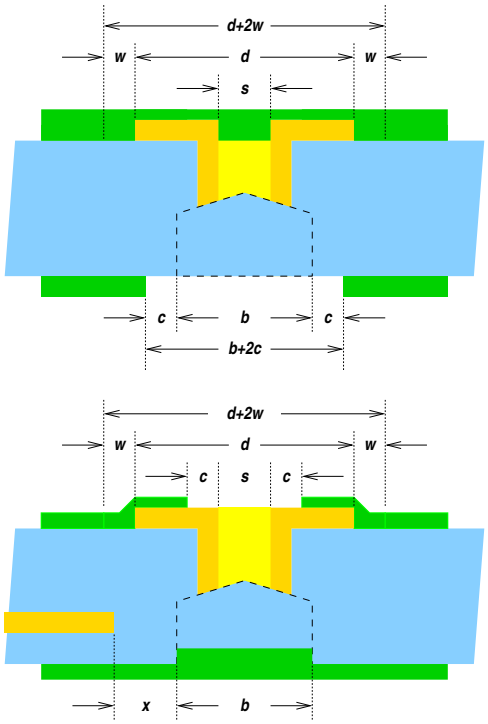

Figure 62: Via Capping: Back-drilled Via

Some fabricators want plugged or capped via holes specified on the drill chart. Others specify that a via plugging image layer be generated for each plugging layer.

5.7.6 Via Plugging Post-Processing

5.8 Primary Mask Layers

The primary mask layer was historically handled by pcbnew but primarily for plotting. Primary (solder) mask was not displayed even though a layer existed complete with visibility and color controls. Also, mask layer rules were not historically adequately provided by pcbnew. Mask layer violations are normally caught by board fabricators; however, there are a number of reasons why this function should not be performed by a board fabricator that is not also responsible for assembly:

- Errors introduced effects such as mask-on-pad might not be caught until during assembly. Bare-board testers (particularly flying probe testers) might miss this altogether.
- Debris from torn or peeling mask that does not meet the minimum web or clearance from board edges or score lines might not be caught until during assembly.
- The legend layer cannot be fully addressed without first fully addressing the primary mask layer.

As with most things, therefore, it behooves the CAD designer to ensure that a usable and manufacturable mask is applied to the board. When the correction of mask is performed by fabricators, there are several things that are done:

1. The pad master or 1:1 mask layer artwork is expanded by the xy-plane registration tolerance of the mask application process.
2. Mask is cleared from board edges, NPTH, score lines, bevels, edge-fingers, slots, cutouts, wells.
3. Where minimum web requirements are not met (typically web between high-pitch SMT pads) eight gang-busting or solder-dams are used, depending on the distance.
4. Where minimum separation is not met between tracks and edge-fingers, vias and edge-fingers, tracks and pads, vias and pads; either the mask is adjusted to risk mask-on-pad over an exposed adjacent copper feature, or the copper feature must be moved or resized.

To perform these corrections, many manufacturers look to specific instructions or established and communicated standardized policy, otherwise orders must be placed on hold.

> **Soldermask.** Soldermask is protective coating that shields selected areas of a PCB from oxidation, handling and unwanted solder during assembly. Soldermask may be applied in two separate operations call primary and secondary. Primary soldermask is applied to cover traces and planes so adjacent features won't short during assembly. The soldermask type that is most commonly used for the primary application is LPI (Liquid Photo-Imageable) due to its superior electrical properties, finer resolution allows for mask strips (webbing) down to 0.003" between SMT pads, and its thickness characteristics, which will allow complete coverage at a thickness of 0.0004" making LPI the mask of choice for Surface Mount technology.
>
> The secondary coating operation is typically used to plug vias. Via plugging is achieved by coating soldermask material into one side of the via. This is generally required when a vacuum assisted In-circuit test is utilized after assembly. The soldermask type that is generally used for plugging vias is an epoxy based thermal curable or UV curable "WetMask."
>
> - The [fabrication] drawing should specify what type of soldermask is required. It is recommended that the soldermask be specified to comply with IPC-SM-840 in lieu of specifying the actual material supplier as this may limit the avialable manufacturing sites. A list of supplier is provided in the guideline section under Sample Suppliers if specifying the supplier is mandatory.
> - It is recommended that reduced clearances be utilized for vias in place of via plugging when possible.
> - The [fabrication] drawing should indicate this and from which side if via plugging is required. See section on via plugging for more information.
> - Finish: LPI is available with a matte, semi-matte, or semi-gloss finish.[29]
> - Color: LPI is available in green, clear, blue, black, orange, purple or red.[30]
> - The soldermask opening should be created in the CAD design at least 0.006" larger than the pad size. THe soldermask clearance should be one half of the pad to trace spacing for clearances less than 0.006". The minimum soldermask to pad clearance should be 2mil on each side (0.004" overall). An optional method is to provide the soldermask openings as 1:1 with the outerlayer pads and allow [the fabricator] to ajust the clearance to meet manufacturing criteria at the tooling stage. It is recommended that high-density designs utilize the 1:1 method.
> - The soldermask opening for non-plated holes should be 0.010" larger in diameter than the maximum hole diameter to allow for registration tolerances.
> - The recomended spacing from the edge of the routed board to the edge of the soldermask is 0.020", minimum spacing is 0.015".
> - When spacing between pads and traces fell below the required minimum to place a web of mask, the result will be either mask on pads or exposed traces. If this situation cannot be avoided, it is recommended that the mask always be allowed to encroach onto the pad in lieu of allowing exposed traces.
> - When spacing between SMT pads falls below the minimum required to place a web of mask, it is recommended that the webs between the pads be removed and replaced with "Gang-Relief" or block clearances.
>
> **Webbing.** Soldermask is commonly placed between surface mount pads to reduce solder bridging during assembly. These small lines of mask between the pads are called webs. Soldermask webs have a minimum width that can be reliably exposed such that they will remain on the board during assembly. The sapce between the pads must allow for registration tolerances for imaging the soldermask. This leaves a small "finger" of soldermask web between the SMT pads. The web must be removed when the web width falls below the manufacturing capability. This is called gang relief. Please specify if gang relief is allowed, or not allowed, on the fabrication print. When spacing falls below the required minimum to place a web of mask between the SMT pads and a web is required, a 0.001" overlap onto the SMT pads will allow a web to be placed if no alternative is possible. This allowance should be noted on the fabrication print.
>
> **Soldermask at the Bottom of Gold Edge Fingers.** Soldermask clearance at the fingers should extend past the bevel and the board edge. This will prevent having soldermask between the end of the bevel and the fingers, where it could flake off during manufacturing.
>
> **Soldermask on the Top of Gold Edge Fingers.** It is preferred to have the solder mask brough down to the top of the fingers, and not to have it stop above the fingers. Stopping above the tops of the fingers will result in circuits exposed from the edge of the soldermask to the fingers. Damage to the exposed circuits may result during the gold tip plating operation. A distance of 0.050" should be maintained between the edge of the mask and the edge of the

29. For heaven's sake: what is the freaking difference between "matte", "semi-matte" and "semi-gloss"?

30. I have seen red mask boards: why would someone do this to their precious design?

nearest plated through hole or via hole clearance. If maintaining the 0.050" spacing is not feasible, then the mask should be brought down on top of the fingers to allow for the 0.050" strip of mask to be maintained.

...

General Soldermask Guidelines

- Green is the preferred soldermask color.
- Green soldermask allows a 0.003" web to be placed between pads in an SMT array, provided the minimum spacing between these pads is 0.008" (by design).
- To assure no soldermask on any ad in an SMT array, the minimum soldermask clearance for a surface mount pad in 0.003" per side. As space permits, a clearance of 0.0025" per side is available.
- Mask features below 0.004" (not including webs between SMT pads) are not allowed (as measured on the CAD data).
- Allow 0.010" per-side soldermask clearance from the edge of a score line.
- Allow a minimum of 0.005" per side clearance for nonplated holes to [ensure] no mask material inside the hole barrel.
- If spacing between a pad and a trace falls below the required minimum, the result will be either mask on the pad or an exposed trace. If this situation is expected, [fabricators recommend] adding a note on the [fabrication] drawing stating that mask on pads is acceptable, exposed traces are not acceptable. This will [ensure] what direction will be taken when the situation occurs.
- To prevent soldermask from going into and/or plugging a hole, soldermask clearance should be 0.006" (0.003" per side) larger than the finished hole size on both sides of the board.

5.8.1 Primary Mask Fabrication

5.8.2 Primary Mask Process

5.8.3 Primary Mask Materials

- Dry-Film
- LPI (Liquid Photo-Imageable)
- LDI

5.8.4 Primary Mask DFX Rules

5.8.5 Primary Mask Design

5.8.6 Primary Mask Imaging

Plotting: How the primary mask is plotted depends on a major fabrication choice: either the mask can be plotted 1:1 (without clearances), with a request on the fabrication print for the fabricator to adjust clearances to the specifications of their process; or plotted complete with clearances, solder dams, gang relief and other necessary adjustments. The trivial case of 1:1 plotting is simply a pad master.

When plotting a 1:1 solder mask, the following must be considered:

- All flashed apertures will be expanded by the process clearance value by the fabricator: therefore, apertures for complex constructions (such as via capping) must be plotted without clearance compensation.
- Zones and other contours will not be expanded and therefore are plotted without consideration for clearance compensation.
- Solder dams (clear overlaid line segments) must be inserted by the fabricator and are therefore not plotted when plotting 1:1.
- SMT and TH pads will be plotted 1:1 to their copper values.
- SMD (Solder Mask Defined) pads should be plotted 1:1 with the copper pad opening, but should be identified as requiring SMD treatment (otherwise the fabricator might expand the opening).
- Openings for via encroachment or capping (secondary mask after primary mask), will be plotted as the finished-hole-size of the hole being encroached upon or capped.

In general, it is not possible to communicate all of the requirements for solder mask design to the fabricator when traditional fabrication outputs are provided. Even though many fabricators recommend 1:1 solder mask data, for plot-and-go, or to reduce the dependence on a single fabricator, the designer should provide proper full solder mask definition.

When plotting a full primary mask definition, the xy-plane registration of the solder mask must be known. Also, the minimum web (the thinnest webbing of solder mask that can be reliably manufactured) needs to be known for DFX checks. Another factor that needs to be known for DFX checks is the maximum amount of SMT or TH pad overlap that is permitted. Under pcbnew, the following items are considered:

- SMD (Solder Mask Defined) pads will be plotted as-is.
- NSMD and TH pads will be be plotted expanded by a clearance equal to the xy-plane registration of the mask.
- Vias that are encroached or capped (where primary mask is applied before secondary), will be plotted with an opening of the finished-hole-size expanded by the xy-plane registration.
- Draw segments on the mask layer (representing solder dams) will be plotted with a thickness equal to the minimum web.
- Zones will be plotted as is. Positive and negative overlays are possible. Dark zones represent closed contours from which solder mask will be removed. Clear zones represent closed contours to which solder mask will be added. (Dark zones are typically used for gang relief, or to remove mask from areas of the board such as gold fingers, or areas of selective gold finish.)

Display:

3D Display:

Editing:

5.8.7 Primary Mask Post-Processing

5.9 Contacts Layers

There are two primary contact finishes that are used on rigid board: electroplated nickel/gold and carbon contacts. Of these, electroplated nickel/gold can be performed using two distinct processes with widely differing cost: gold finger (edge tip) plating, and selective electroplated nickel/gold.

Gold Finger. Edge plated fingers (commonly called Tips): A nickel coating is deposited on bare copper with a cobalt hardened gold surface finish electrolytically deposited on it. This is commonly called TAB plating. This is typically applied in a vertical plating line and should only be used near the edge of the board. This is a relatively inexpensive fabrication process. This is the preferred method for creating edge card fingers. There are restrictions on how far the fingers can extend into the PCB and how long the PCB may be for processing.

Selective Electroplated Gold. This process applies a nickel coating onto the bare coper and coats it with an electroplated gold surface. It differs because the entire board must have additional fabrication steps, which remove the primary surface finish, such as HASL, to apply the gold and only to the required locations. This is a relatively expensive process. However, selective non-bussed areas can be plated with nickel and gold before etching, by the use of a gold mask. Simple gold masks can be created by the fabricator when detailed information is supplied to what features will be plated.

Carbon Contacts: Carbon elements or patterns are printed with a conductive carbon ink and can be used for keyboard contacts, LCD contacts, jumpers, etc. Always clearly indicate on which side of the board carbon is to be applied. This can be one side or on both sides. Indicate the proper position by file naming and clear buildup or layer sequence description in the mechanical layer. The overall carbon design rules are:

A = minimum carbon to carbon spacing: typ. 0.400mm (16mil). (65)

= minimum carbon line width: typ. 0.300mm (12mil). (66)

B = minimum carbon on copper overlap: typ. 0.200mm (8mil). (67)

C = minimum carbon on solder mask overlap: typ. 0.100mm (4mil). (68)

= minimum isolation or spacing between copper elements of a typical carbon finger contact: typ. 0.800mm (32mil). (69)

Carbon filled PTF (Polymer Thick Film) paste is now widely used on printed circuit boards for use as an alternative to electroplated nickel/gold contacts. The paste, once cured, is highly conductive and provides a corrosion resistance comparable to gold and a hard durable surface that is able to withstand may insertions or contacts and fluxes used during soldering operations. The resistance of the cured carbon is $\rho \approx 40\Omega/\text{cm}^2$ over full contacts at parallel edges at 30μm coating thickness. The requirements for two approaches to carbon fingers is illustrated in *Fig. 63(79)* and *Fig. 64(79)*.

Approach #1. In the first approach, illustrated in *Fig. 63(79)*, solder mask is fully cleared from the copper contacts (typ. 0.1mm/4mil), and the carbon ink is overlapped on the primary mask by a full overlap (typ. 0.1mm/4mil). The parameters are:

P = minimum copper-to-copper spacing (70)

W = minimum solder mask web (71)

S = minimum carbon-to-carbon spacing (72)

O = minimum carbon overlap to mask (73)

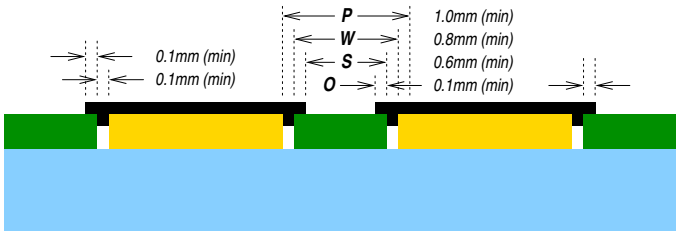

Figure 63: Carbon Ink: Approach #1

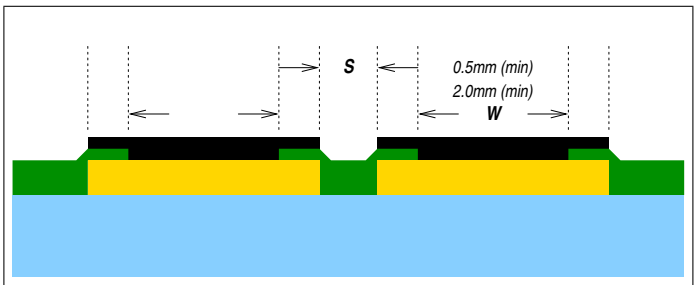

Figure 64: Carbon Ink: Approach #2

For this approach, the copper pads must be defined as NSMD pads. The carbon mask must be defined as the copper pad outline expanded by the primary mask registration clearance, and again expanded by the minimum carbon overlap, O.

Approach #2. In the second approach, illustrated in *Fig. 64(79)*, solder mask is overlapped on the copper contact and then carbon ink is overlapped over the combination of contact and mask. This approach is preferred over *Approach #1*. The parameters are:

S = minimum copper-to-copper and carbon-to-carbon spacing (74)

W = minimum width of the carbon-copper contact (75)

For this approach, the copper pads must be defined as SMD pads. For this approach, the carbon mask is defined simply as the copper pad outline (which itself is expanded by the minimum solder mask overlap for the SMD pad). In this approach, the minimum carbon-to-carbon distance, S, is typically a little narrower (typ. 0.5mm/2mil) than can be obtained with *Approach #1* (typ. 0.6mm/2.4mil). For anything but key contacts, this is likely not an issue. Nevertheless, *Approach #2* is the preferred approach.

5.9.1 Contacts Process

5.9.2 Contacts Materials

Contact materials can, in general, include any of the finish materials. However, board finish is normally more suitable to forming solder joints and less suitable to mechanical contact surfaces. Therefore, contact processes and materials are typically distinct from surface finish materials. Typical contact materials include:

Selective Electroplated Nickel/Gold: Selective electroplated nickel/gold is a process step used in the fabrication of a foil composite board. Once the traces and pads are copper plated on the outer copper foil, an separate gold mask is used to image additional plating resist onto the surface. An anode electroplated nickel/gold process is then applied to plate nickel/gold on to the copper features and then tin plating for etch resist. This results in an electroplated nickel/gold surface being applied to copper areas before solder mask and primary finish is applied.

Peel-off masks can be used to protect the contacts from primary surface finish as well as wave soldering during assembly. Although the process is distinct, the result is a thin layer of nickel covered with electroplated hard gold.

Gold Finger: Gold finger is a process step whereby tie bars connected between the edge tips and production panel electroplating contacts are used to plate nickel/gold on edge tips in a vertical plating line. Gold finger processing typically requires that the production panel be laid out in such a way that the panel may be split and only the edge fingers themselves dipped into a vertical plating line. Although the process is distinct, the result is a thin layer of nickel covered with electroplated hard gold.

Carbon Ink: Carbon ink is a polymer composite that can be screen printed onto the board and cured. It is normally applied in the final stages of the board fabrication (after application of the primary mask and before or after the application of primary and secondary surface finishes). The result is a carbon ink coating with specific electrical and mechanical (hardness) properties.

Others: Other, more esoteric, materials can also be applied to contacts. The material and process specification should not assume a limited set of available materials or processes.

5.9.3 Contacts DFX Rules

The three primary contact processes, gold finger, selective electroplated nickel/gold, and carbon ink, each have their own set of DFM rules due to the unique processes involved.

Selective Electroplated Gold. Selective electroplating of gold before etching requires a gold mask. The following DFM rules need to be followed:

1. Non-gold plated circuits and pads should be at least 0.020" away from selective gold plating features.
2. The gold mask should be designed so that the mask is 0.010" larger than the features that will be gold plated.
3. If via holes, or other plated holes, fall into the area that is to be gold plated, then both sides of the gold mask should have these included in the gold mask. Tenting a hole on one side, and gold plating the pad and hole on the other will not work with the process as it will produce holes with plating voids.
4. This process adds several additional processing steps for PCB manufacturing. Selective gold should not be specified to reduce the gold plating area to save cost. The additional processing costs will be much higher than the savings.

Carbon Ink. DFM rules for carbon ink have already been presented above. The rules are distinguished by whether the contact surfaces are marked as SMD (Solder Mask Defined) or NSMD (Non-Solder Mask Defined) pads:

1. For SMD pads, there is a minimum copper-to-copper separation, minimum defined copper width, and minimum carbon to carbon separation. This is illustrated in *Fig. 64(79)*.
2. For NSMD pads, there is a minimum copper-to-copper spacing, minimum carbon overlap to mask, and minimum carbon-to-carbon spacing. The minimum solder mask web must also be considered as a solder mask web is required between and surrounding contacts. This is illustrated in *Fig. 63(79)*.

These parameters have been added as process parameters to the *contacts* process dialogue.

5.9.4 Contacts Design

5.9.5 Contacts Imaging

Display: Contacts should be displayed as-plotted on the pcbnew canvas. A separate visibility item for contacts should be provided to control the colour and visibility of contacts on the pcbnew canvas. Contacts are displayed in their own layer, where the layer information represents the surface finish of the contact.

3D Display:

Plotting: Separate plot (Gerber, DPF) artwork is required for each surface finish of contacts and for each side of the board. Drawings made on the "contacts" layer can be used to create tie bars when required for gold electroplating.

Gold Finger: For gold fingers, artwork for the gold fingers can be provided as a print of the "contacts" layer, for both sides of the board. Contact pads are plotted at the same size as—that is, 1:1 with—the copper pad. Gold fingers can also be plotted on the peel-off mask layer where the pads are expanded by the minimum overlap for peel-off mask, typically 0.025" (0.6mm). Because gold finger contacts are typically protected with Corfin or Kapton tape instead of peel-off mask, the contacts or removable mask process should have an option for selecting the removable mask (if any) to the applied to the contacts.

Selective Electroplated Gold: For selective electroplated gold used for contact surfaces, the "gold mask" can be provided as a print of the "finish" layer, for both sides of the board. In this case, the "gold mask" is plotted as the copper pads expanded by the gold mask clearance, typically 0.010" (0.25mm) or half the clearance to non-plated features. The "gold mask" can also be plotted on the peel-off mask layer to define a protective mask for the contacts. Contact pads protected with peel-off mask should be defined as SMD pads (the mask can leave a blue residue on exposed laminate). The contacts or removable mask process should have an option for selecting the removable mask (if any) to the applied to the contacts.

Carbon Ink: For carbon ink, artwork for the screen-printed carbon ink patterns must be provided as a separate artwork file for each side of the board. Carbon ink patterns are expanded from the defined pad (whether SMD or NMSD) and are expanded once or twice the primary mask clearance depending on whether the underlying contact pad is SMD or NSMD. See *Fig. 63(79)* and *Fig. 64(79)*. Carbon ink contacts can also be plotted on the peel-off mask layer where the pads are expanded again by the minimum overlap for peel-off mask, typically 0.025" (0.6mm). The contacts or removable mask process should have an option for selecting the removable mask (if any) to the applied to the contacts.

Editing:

Selective Electroplated Gold:

Gold Finger:

Carbon Contacts: Carbon contacts require a mask image layer that defines the carbon ink print. Using the either approach, this is simply the expanded defined copper outline. Imaging of the primary mask layer must follow the normal SMD or NSMD definition of the copper pads used for the contacts. The pcbnew module editor now supports a "contact" layer. Different requirements for the definition of the size of contact images can be accommodated by adding a pseudo-pad that exists only on the contact layer.

5.9.6 Contacts Post-Processing

5.10 Peelable Mask Layers

Historically, pcbnew did not handle peelable mask or other removable mask layers. It now supports two removable mask layers. These are photo-imagable board fabrication layers.

Peelable soldermask is generally used to prevent gold plated contacts, carbon elements, and select plated through holes from being coated or plugged with solder during the HASL (Hot Air Solder Leveling) process, or during wave assembly followed by reflow. It is a temporary soldermask and can either be removed after the HASL process or left in place as a protective barrier and removed by the end customer prior to assembly.

Usually, the peel-off solder mask is only applied to one side of the board: the solder side, as this is the side that is subjected to wave (mass) soldering. When used on both sides of the board, or on the top side, this must be clearly indicated in the data package. The following rules apply to peel-off design:

- P = minimum width of any peel-off element: typ. 0.5mm (20mil).
- H = maximum coverable finished hole size: typ. 6.0mm (236mil).
- V = minimum overlap on copper pattern: typ. 0.6mm (24mil).
- W = minimum clearance to free copper: typ. 0.6mm (25mil).
- Minimum distance from board edges: typ. 0.5mm (20mil).
- Positional tolerance: typ. ±0.300mm (12mil).
- Avoid using many different small peel-off areas randomly placed on the board. Make the peel-off are as large as possible by connecting the separate peel-off areas wherever practical. This will make it easier to remove the peel-off after soldering.

5.10.1 Peelable Mask Process

The typical application process consists of screen-printing. Some considerations for peelable soldermask are the distance from the protected contact to the nearest unprotected feature susch as a plated via hole or solder pad, the amount of overlap required, the material thickness, and the minimum allowable width of a peelable strip. Additional concerns include the predicted number of thermal cycles expected prior to removal as the peelable soldermask can become difficult to remove after three thermal cycles.

- If peelable soldermask overlaps onto areas of bare laminate not covered by the primary application of soldermask, a light blue residue may remain on the baord after removal. Due to this condition, fabricators recommend that all gold plated contact pads or carbon patterns that are required to be convered with peelable soldermask, utilize the minimum soldermask clearance allowable on the primary (permanent) soldermask file. For example, if a gold plated contact pad is required to be covered with peelable soldermask, the soldermask clearance for that pad on the primary soldermask file should be no more than 0.003" per side over the gold plated contact pad.
- Boards utilizing peelable soldermask should not be subjected to more than three thermal cycles, preferred is no more than two thermal cycles.

Corfin is a protective coating that is applied over connector contacts to provide protection during assembly. It is removed from the PCB after component assembly.

5.10.2 Peelable Mask Materials

5.10.3 Peelable Mask DFX Rules

5.10.4 Peelable Mask Design

5.10.5 Peelable Mask Imaging

5.10.6 Peelable Mask Post-Processing

5.11 Finish Layers

There are a number of finishes that are applied to rigid boards. Each has its advantages and disadvantages. I prefer IAg because of its benefits for high-frequency boards and excellent planarity for high-density boards and chip-scale packages. Typically when there is a single (primary) board finish, the finish is simply specified in the fabrication print.

Where multiple surface finishes are required, it is necessary to provide images to support the secondary finish. A typical secondary finish is selective electroplated gold. On high-speed boards, it is common to remove solder mask from critical microstrip lines and finish those lines with selective nickel/gold so that the solder mask irregularities do not affect the controlled impedance of the critical lines. Typically the primary finish is identified on the fabrication print, with a reference to the data provided for the secondary finish. This is detailed below.

Historically pcbnew could only support a single finish via fabrication print notes: that is, pcbnew could not historically support secondary finishes. The new "finish" layer provides pcbnew with the ability to support multiple finishes.

Selective Electroplated Gold: This process applies a nickel coating onto the bare copper and coats it with an electroplated gold surface. It differs because the entire board must have additional fabrication steps, which remove the primary surface finish, such as HASL (Hot Air Solder Leveling), to apply the gold only to the required locations. This is a relatively expensive process.

Selective non-bussed areas can be plated with nickel and gold before etching, by the use of a gold mask. These gold masks should be supplied in the data package. Simple gold masks can be created by the fabricator when detailedinformation is supplied to what features will be plated.

- Non-gold plated circuits and pad should be at least 0.020" away from selective gold plating features.
- The gold mask should be designed so that the mask is 0.010" larger than the features that will be gold plated.
- If via holes, or other plated holes fall into the area that is to be gold plated, then both sides of the gold mask should have these included in the gold mask. Tenting a hole on one side, and gold plating the pad and hole on the other will not work with the process as it will produce holes with plating voids.
- This process adds several additional processing steps for PCB manufacturing. Selective gold should not be specified to reduce the gold plating area to save cost. The additional processing costs will be much higher than the savings.

pcbnew has not historically modelled the finish layers for board fabrication. The finish layer represents the board finish that is applied to the exposed copper areas of the board. Board finish can consist of a single finish (such as HASL), or multiple finishes. Multiple finishes is expensive because it requires additional board handling steps.

5.11.1 Finish Process

5.11.2 Finish Materials

5.11.3 Finish DFX Rules

5.11.4 Finish Design

5.11.5 Finish Imaging

Selective Nickel/Gold: Selective nickel/gold requires a "gold mask." The "gold mask" is one or two separate image layers (one each for top and bottom contacts) that defines the zone in which selective nickel/gold is to be applied. These masks are used to control the selective plating of nickel/gold after the copper plating process for foil composites. In addition to the "gold mask" the selective nickel/gold zone must likely be expanded by the minimum peel-off mask overlap and plotted on the peel-off mask layer for the corresponding side.

5.11.6 Finish Post-Processing

5.12 Legend Layers

PCB Libraries had this to say about silk screen: [Hausherr, 2006]

> Silkscreen outlines are used for cosmetic purposes only and are really not required for board fabrication.
>
> Most CAD land pattern silkscreen outlines are not representative of the true component outline. Silkscreen outlines must avoid the exposed copper pad by 0.4mm for maximum clearance 0.3mm nominal or 0.25mm minimum.[31]
>
> The standard acceptable line width for silkscreen outlines is 0.2mm.[32]
>
> The silkscreen can be drawn by the PCB designer very complex to illustrate their creative talent or very simple. In the end, it really does not matter because you can only see it when the physical PC board passes between fabrication facilities to the assembly shop. Once the parts are assembled, all the silkscreen outlines are covered up and cannot be seen.
>
> The silkscreen originated from the hand tape days. Back in the 1970's PC boards did not have silkscreen outlines. Assembly drawings were created using rectangles and circles. Some PCB designers decided the assembly drawing would look good if it appeared on the PC board and the silkscreen was born.
>
> Now, most boards have part placements that are so tight that there is no room for silkscreen outside the part. It is one opinion that with a one world standard library and full machine automation that the silkscreen will be a thing of the past. It is interesting to note that PC boards that go into outer space purposely do not have silkscreen to reduce the weight of the product.
>
> When PCB designers start to use all the principles discussed in this outline, the manufacturing assembly process can be fully automated. The Cad Library of the Future will not require the use of silkscreen outlines. [Hausherr, 2006]

These statements are incorrect in almost all respects: the numbers are exaggerated and the uses of silkscreen are misunderstood by the author. On the contrary, the benefits of the Legend or Silk Screen are as follows:

1. Reference designators printed on the legend identify components. Identifying components is necessary to assist the designer when performing board layout and design, to locate components on the board when preparing assembly tooling as well as when visually inspecting assembly, or during any point in time that the board needs to be reworked or repaired. The end-user of the product needs the reference designators to located and identify jumper and connector locations.
2. Component outlines also assist in the purpose of forming a connection between reference designators and the components to which they refer. This helps all tasks that need to identify components.
3. Component outlines identify the orientation of components that are not square or circular.
4. Component outlines identify not only the orientation, but the location for placement of components.
5. Component outlines identify which pads belong to the component and which do not.
6. Component outlines or other legend features provide markings that, in relation to the component outline identifies the precise orientation of non-symmetric parts. That is, pin #1 identifiers and polarity markings.
7. Legend markings may also identify other things such as the default setting of switches or jumpers, normally open or normally closed connections, and other information necessary to the end-user of the card.
8. Legend markings contain certification marks, copyright and patent information, help (800) numbers, manual section numbers for repair, test, or user configuration, and other markings necessary to the end-use of the product.

I believe that board fabricators (that are not also performing assembly) only care about silk screen on pads because it can impair bare-board testing.

> Some assembly shops or test engineers who are debugging a prototype require silk screen Polarity Markings. Silkscreen Polarity Markings are primarily used to illustrate the "Positive" terminal of a two-pin component. Polarity Markings are also used when there is a potential for inverting part placement in the assembly process that would result in a malfunction of the component.
>
> The term Polarity Marking came from its use to identify the positive pin on a polarized capacitor. But polarity marking is also used on diodes to indicate the Cathode Pin. It is also used on connections to illustrate the Pin 1 location. [Hausherr, 2006]

Ramification of boards with no silk screen. Sanmina-SCI, a less "polarized" board fabricator has this to say about silk screen: [SS-2007]

> Legend (silkscreen) provides graphical component outlines or alphanumeric characters to designate component locations. It is applied using a standard screen-printing process.
>
> Legend should not be used on high-density designs. Generally, a non-readable character remains when minimum character sizes are used and then modified to removed legend from all exposed conductive surfaces. Components are located by using assembly drawings or trouble shooting documentation. [SS-2007]

This is the second line of argumentation by a board fabricator against legend. Again I am convinced that board-only fabricators only care about silk screen on pads because it can impair bare-board testing. The ramifications of this throw-it-over-the-wall approach are as follows:

- Many full production shops will tool their pick-and-place machines by having an operator teach the machine by walking through the population of a pilot card. Without silkscreen, the operator has fewer points of reference to avoid errors in this process. Consulting an off-line assembly drawing does not suit the same purpose because the silk screen is normally viewed through the machine's vision system during the training of the machine.
- Turn-key and short production run assemblers develop pick-and-place programs with off-line development tools that use the silk screen information layers for programming and verification of component position, orientation and polarity. When no silk screen is provided in the data package, another suitable source must be provided. This can take the form of assembly drawings, but then assembly drawings need to be plotted just as the legend anyway.
- During manufacturing, through-hole (TH) parts are often inserted manually. That is, a person inserts the part into the board. Without a legend, this can result in misalignment of parts in the final stages of assembly. Referring to a separate assembly drawing is not quite the same as identification on the board.

31. These numbers are greatly exaggerated: the clearance between mask and pad for SMT is as little as 1 to 3 mils (0.0762mm) and then the clearance of the silk screen to the mask edge is dependent on the silk screen registration (so ink does not land on the laminate where it can chip or flake off) which is typically 3mils (0.0762mm), so the nominal clearance is 6 mils (0.15mm) and this is what a board fabricator uses for legend clipping.

32. The minimum silk screen line width is typically 6mil (0.15mm), not 8mil (0.2mm).

- Visual inspection of the assembly becomes cumbersome for the same reasons: using a separate assembly drawing does not allow the same frame of reference.
- Assemblers charge more for no silk.

5.12.1 Legend Process

There are three processes for applying legend to a printed circuit board:

1. Traditional silk screen printing process.
2. Continuous ink-jet printing.
3. DOD (Drop On Demand) ink-jet printing process.

Traditional silk screen printing processes are limited in the minimum line widths and character sizes that can be achieved. Continuous ink-jet printing is not much better and is comparable to silk screening. DOD (Drop On Demand) ink-jet printing can achieve better performance than the other methods. At 300 DPI (Dots Per Inch) the dot size is 5mil (0.13mm) nominal and can resolve a character size as small as 4pt (56mil (1.42mm)); at 600 DPI, dot size 3mil (0.08mm) nominal, character size, 2pt (28mil (0.71mm)). For 600 DPI this is smaller than silk screening.

Resolution of characters is largely influence by the printing process and the characteristics of the ink used. Parameters of the printing process that are of concern during design are:

Minimum Line Width: The minimum line width that can be resolved. This is dependent upon the smallest attainable line width for screen printing, or the dot size for ink-jet printing. Also included in fabricator recommendations for minimum line width it the susceptibility to cracking and flaking after thermal stress or due to adhesion parameters.

Minimum Character Height: The minimum character height that can be resolved without bleeding inside the character making the character illegible. Some character fonts are more susceptible to this than others (see *Sec. 7.12(142)*).

Minimum Character Separation: The minimum character separation that can be resolved without bleeding between characters. Minimum character separation is also a function of the character font (see *Sec. 7.12(142)*).

Registration: The registration of the resulting image to the surface of the board.

Clearance: The minimum clearance to solder mask openings to avoid bleeding into the opening. This is related to minimum character separation.

Height: The finished profile (or height) of the resulting ink pattern. This is related to the stencil thickness for screen printing, and to the dot size for ink-jet printing.

For legibility of characters that are silk screen or ink jet printed, see also the considerations of single-stroke font design discussed in *Sec. 7.12(142)*.

5.12.2 Legend Materials

The material for legend is typically a white, non-conductive epoxy ink. The following material characteristics are of interest to the design process:

Dielectric Properties: The dielectric properties of the cured ink. The dielectric properties of the ink can be of concern for large areas of ink applied over critical traces in high-speed circuitry. Where the dielectric properties of the ink are comparable to solder resist or expoxy resin and the height of the ink is not significant, all but large areas can be ignored.

Thermal Stress: The resistance of the ink to cracking and chipping resulting from, for example, the thermal shock of reflow or wave soldering. These material properties are typically reflected in the minimum recommended line width and character height.

Adhesion: The adhesive strength of the resulting ink to solder mask bond. This material property is also typically reflected in the minimum recommended line width and character height.

Colour: The colour. Legend can also be provided in colours other than white. One of the objectives in choosing a colour is to achieve a high contrast between the legend and the solder mask behind it. Because solder mask is available in a number of colours, legend ink is available in a number of colours too.

In general, however, there is not a lot of variance in these material parameters in available and standardized inks. Some ink-jet inks might exhibit different characteristics due to differing trade-offs in the printing process itself. For example, DOD Ink-Jet printing uses UV curable white inks that meet 4600 series specifications.

5.12.3 Legend DFX Rules

Legend layer rules were not historically provided by pcbnew. Legend layer rule violations are normally caught by board fabricators; however, there are a number of reasons why this function should not be performed by a board fabricator that is not also responsible for assembly:

- Errors introduced by effects such as silkscreen on pad might not be caught until during assembly. Bare-board testers (particularly flying probe testers) might miss this altogether.
- Debris from cracking silkscreen lettering that does not meet the minimum web might also not be caught until during assembly.
- Legibility of the legend is a matter for assembly and particularly rework or inspection. Board fabricators often express their disdain for the need for a legend at all. Again, problems with the legibility of the legend might not be caught until someone has to look at it, and it is unlikely to be the board fabricator.
- The adjustment of silk screen by the fabricator is just another lock-in item.

As with most thing, therefore, it behooves the designer to ensure that a usable legend than can be manufactured is applied to the board. When the correction of legend is performed by fabricators, there are several things that are done:

1. The thickness of lines used in legend text is set to the best minimum thickness of the legend process. This is accomplished by simply increasing the size of the Gerber/Barco aperture used to draw the lines.
2. The height of the text is altered (upward) to the smallest height acceptable by the legend process. This requires sophisticated software to isolate the draw segments making up the text in Gerber/Barco plots and scaling them. It is also a manual process with an associated cost.
3. Legend text that is adjacent and not within clearance rules to pads or solder mask openings can be treated in one of two ways:
   (a) If simply movable, the text is moved. This is a manual process with the attendant costs.
   (b) When text cannot be moved, it is clipped. This can be performed automatically from the solder mask artwork and does not normally have an additional cost.

4. Legend lines and graphics not associated with text can represent courtyards. When the courtyards are not sufficient in size to lie outside the necessary clearances, they can be expanded in size. This is another manual process.
5. Lines (not associated with moved text or outlines) will simply be clipped. This is automatic.

The following are a collection of rules regarding the positioning of legend:

- The minimum line width for legend is 8mil (0.2mm) without review, or 6mil (0.15mm) with review.
- The distance from any copper defining a solder mask opening is 12mil (0.3mm) without review, or 6mil (0.15mm) with review.
- "If there is contention between legend and other mask features, the fabricator will crop the legend to the edge of the soldermask opening. It is preferred that legend not need such correction, be being placed a sufficient distance from the primary mask opening."
- "One also has to bear in mind that screen printing is associated with fairly loose tolerances, both in positioning the print and because the ink tends to flow. Features such as lines surrounding pads to indicate a component position should therefore be kept well clear of any solderable pads."
- "Also, bearing in mind that legend prints have a third dimension, and are not totally flat, designers should not position legend under chip components, to avoid the possibility of drawbridging."
- "Legend should not be used on high-density designs. Generally, a non-readable character remains when minimum character sizes are used and then modified to remove legend from all exposed conductive surfaces. Components are located by using assembly drawings or trouble-shooting documentation." [SS-2007]
- "To [ensure] legend characters do not encroach on solderable pads the following process should be used. Modify the soldermask artwork to [ensure] no mask on pads by providing a minimum of 0.0025" per side clearance over solderable pads. Create a separate file using the soldermask artwork and increase all features by 0.005" per side. Overlays this new file with the legend artwork and relocate any legend feature that falls inside or tangent to any soldermask clearance. The result should provided a minimum of 0.0075" clearance from the edge of a legend character to the edge of any copper or solder coated feature. This distance includes the registration tolerance for the soldermask and legend features. If this procedure is followed, no editing of the legend artwork is required. If the outer layers contain large areas of ground plane and legend is required in these areas, [fabricators recommend] covering these areas with soldermask, as legend ink will not adhere well to soldered areas." [SS-2007]
- "To ensure all letters, numbers and figures are legible on the finished board, character line widths should be a minimum of 0.006" and at least 0.035" high." [SS-2007]
- "Space letters at least 0.007" apart so they don't bleed together." [SS-2007]
- "No legend nomenclature should overlap a copper pad or plane area. This is especially important for surface mount pads and fiducial marks." [SS-2007]
- "Use the fabrication print notes to specify special features to be screened onto the board." [SS-2007]
- Minimum line width is 0.006" (0.15mm) for standard technology; 0.005" (0.125mm) for high technology and 0.004" (0.1mm) for advanced. [SS-2007]
- Registration tolerance is +/-0.005" for standard; +/-0.004", high; and +/-0.003", advanced. [SS-2007]
- Minimum character height is 0.045", standard; 0.035", high; and, 0.030", advanced. [SS-2007]

There are several parameters of the board fabrication processes (capabilities) that must be specified to properly address plot-and-go legends. In general, because legend must be applied over the primary mask[33] it will be necessary to know where the primary mask may be present and where it may be missing. Therefore most of the parameters associated with the primary mask must also be known to properly assess the legend, including soldermask registration and minimum web. Fortunately, these parameters are already defined for the primary mask process. These parameters are:

Regisration of Mask. This is the positional (xy-plane) registration of the soldermask. This is also the amount by which primary mask openings must be expanded to avoid mask on pad.

Hole Tenting Ability. Before checking or adjusting legend profiles, it is necessary to understand whether the primary mask process is capable of fully tenting holes (such as is the case, for example, for dry-film soldermask), and whether particular holes are filled or capped even when tenting by primary mask is not possible.

There are additional primary mask parameters; however, those are used primarily for the calculation of mask openings and dams. This calculation must be performed before addressing legend anyway, and is no different than the calculation made for the primary mask layer itself.

Furthermore, to determine whether legend can print over any given hole (usually via holes), it must be known whether the soldermask can fully tent holes or whether the holes are filled with a material. Normally holes that are filled with a material are then covered by soldermask when they are not otherwise functional. When a soldermask opening is created around the hole, assessing the legend is simply a matter of determining whether the legend is proper considering the soldermask opening and nothing else. However, when the primary mask has no opening for a hole, whether the legend can touch or cross the hole is a question of whether the hole is tented or not (by the primary mask process). When the hole is fully tented, the legend can cross or butt up agains the hole. However, when the primary mask is not fully capable of tenting the hole, and the hole is not otherwise capped or filled, then legend that butts up against the hole or crosses the hole, is in danger of cracking at the edges of the hole. The debris associated with the little peices that crack of can cause problems during assembly or otherwise reduce the long-term reliability of the board assembly.

These parameters are:

Registration of Legend. This is the registration of the silkscreen to the board laminate. Of particular interest is the relationship of this tolerance to the positional tolerance of a given hole (whether PTH or NPTH) and the positional tolerance of primary mask openings.

Clearance Legend to Hole. This is the clearance between the hole and the legend that is necessary to avoid cracking and peeling around the edge of the hole. Note that this value depends on the registration of the legend, and the registration of the hole (PTH or NTPH).

Before applying the clearance to hole rule, it must first be determine that the hole is not fully tented or covered by primary mask, or that there is actually a hole or significant depression in the soldermask at the position of the hole (the hole could be filled or capped).

Clearance Legend to Pad (SMT). The clearance between legend and an SMT pad is dependent upon the primary mask process, enlargement of soldermask openings, and the registration of the legend with repect to the pad.

33. The most available material for legend, white expoxy ink, des not stick to laminate.

Legend Minimum Web. The minimum thickness that any legend feature can exhibit without risking cracking and separation of the legend from the primary mask. Note that this value could easily depend on the primary mask type, whether the legend passes over tracks, the thickness of the outer copper layer under the legend, and other factors.

Legend Minimum Space. The minimum space between lines to avoid bleeding of one line into the other (thus reducing legibility). This is particularly important for character separation to avoid one character bleeding into the other and making both characters illegible.

Legend Line Width. The minimum line width must be known. This is a little mores subjective than minimum web. This applies also to the readability of the resulting characters. On modern high-density boards, the minimum is the one most commonly used. This is fabrication process specific; therefore, all lines smaller than this width will be adjusted upward to that width and kept in the same position. This will only be done for display and fabrication outputs: the board database will still have the design width rather than the fabrication width. A legend character line narrower that this width will fail DFX checks unless the designer requests that minimum line be adjusted for both display and fabrication outputs.[34]

Legend Character Height. The minimum legend character height must be known. This is somewhat asthetic in that characters beneath this height are not subjectively considered legible. On modern high-density boards, the minimum character height is the predominate one. This is fabrication process specific; therefore, all characters shorter than this height will be adjusted upward to that height and kept in the same position. This will only be done for display and fabrication outputs: the board database will still have the design width rather than the fabrication width. A character shorter than this height will fail DFX checks unless the designer requests that minimum character height be adjusted for both display and fabrication outputs.

(R) 29 (pcbnew) *Process design parameters will be added to the silk screen (legend) layer class as follows:*

1. *Minimum line width: the narrowest permissible width of a line making up a single-stroke character font.*
2. *Minimum character height: the minimum permissible height of a character.*
3. *Minimum character spacing: the minimum distance between features of adjacent characters.*
4. *Minimum web: the minimum web width for logos and marks.*
5. *Registration: the xy-plane registration of the silk to the board absolute origin.*
6. *Tolerance: the tolerance in the size of printed features.*

DRC Errors and Markers. An overlap between a clearance line drawn around legend features and clearance lines drawn for features generating primary and secondary mask openings will be considered a conflict. When moving legend items, the tangency of clearance lines will be considered an intrusion and movement will resist on these boundaries when DRC is turned on and the active layer is the legend layer of the item being moved.

- Violations during DRC check will be reported as *"Legend too close to mask opening"*, or *"Legend too close to hole."*. Violation markers will be placed on the first overlapping position found during calculations. These violations can be reported as a warning because fabricators normally adjust legend as required.
- Additional checks include whether a line width is too narrow or a character height is too short, or whether sufficient space is not present between characters. These conditions will be reported as *"Line too narrow"*, *"Character too short."*, *"Character too narrow"*, or *"Character spacing too tight"*.
- *"Character spacing too tight"* will be reported as an error. It will only be reported once for each text string. A marked will be placed on the anchor point of the text string.
- *"Line too narrow"*, *"Character too short"*, and *"Character too narrow"* will be reported as warnings. This will only be reported once per feature. That is, only once for each text item. A marker will be placed at the anchor point of the item.

When editing legend items, only one violation (the first discovered) need be reported for each legend feature. When performing DRC/DFX checks, all violations will be reported. Most of these violation can be reported only as a *"warning"* because fabricators normally adjust legend as required. This is, therefore, only a concern for plot-and-go.

(R) 30 (pcbnew) *pcbnew will be enhanced to perform DRC checks against design parameters for silk screen (legend) drawings and text.*

5.12.4 Legend Design

Legend contains the following items:

- *Component outlines.* This information is contained in the module library and is designed using the module editor.
- *Reference designator text for components.* The reference designator is supplied automatically by KiCad; however, the designer must ensure that the placement of the reference designator on the board is proper. The same applies to any other text fields that are to be printed on silk screen.
- *Additional board text.* This may include copyright statements, applicable patent numbers, product name, board revision, etc.
- *Product Specific Bar codes.* A product-specific bar code used by bar-code readers to identify the line of the product. (See *Sec. 7.14(144)*.)
- *Company logos.* A decorative company logo. (See *Sec. 7.14(144)*.)
- *Certification marks.* An association or certification mark or logo. (See *Sec. 7.14(144)*.)
- *Certification serial numbers.* A certification serial number.
- *Item Specific Bar codes.* An item-specific bar-code that can include the line of the product, the board's serial number, and the data of manufacture.
- *Manufacturing date.* The manufacturing date of the specific board.
- *Serial numbers.* A serial number for the specific board.

Some of this information is traditionally applied with stickers, such as bar codes, UL certification and serial number, manufacturing dates, and card serial numbers. The application of stickers is pretty much a manual process and is therefore labour intensive and costly. Ink-Jet printing of legend can provide bar codes because it has sufficient resolution (600 DPI) to represent a bar code properly in legend ink. One of the more difficult things to do is to print serial numbers and manufacturing dates on cards. This is because the information is different for every card.

34. These are what I call *plot-and-go options*.

Character fonts and faces. In general, a CAD system needs to provide single-stroke character fonts that provide maximum legibility and renderability on silk screened legend or copper planes. A character font should keep all areas local to the character a open as possible. For example, a '4' with an open top is better than a '4' with a closed top. Also, it should avoid parallel lines at the edges of two characters. For example, the vertical bar of a '4' should be slightly more to the left than to the right. Selection of character fonts is detailed further in *Sec. 7.12(142)*. The selection of font face is a design parameter not specified by the board fabricator.

(R) 31 (pcbnew) *A selection of font for the silk-screen layer will be provided along with the layer design parameters.*

Logos and marks. In general, specialized items that need to appear in silk screened legends or on stickers should be identified in the fabrication print. A typical approach is to provide a fabrication print detail identifying the zone in which to apply a UL (Underwriter's Laboratory) sticker with serial number, and a fabrication note on the master print directing attention to the fabrication detail sheet. pcbnew will have the ability to place logos on the legend using the techniques detailed in *Sec. 7.14(144)*.

Module editor issues. There are some issues with this approach when the module editor is in used from outside the board environment (i.e. when used for creating and editing library modules). Some parameters of the design are unknown when generating footprints, such as minimum line width, minimum character height, minimum character spacing, primary mask registration (clearance), legend registration, and to some degree the tenting of holes. Therefore, for the module to be largely usable, it must meet the most restrictive of these minimums.

To mitigate this issue, a number of procedures can be employed. To mitigate some of the unknown parameters that are used in the design of the board that are unavailable to the module writer, minimums can be applied to the module at the time of import into the board. Line widths beneath the minimum can be increased to the minimum. Character height, and separation can be increased to the board minimums. The character font can be overridden with the font selected for the board. The silk screen can be clipped using the board separation from mask openings, registration and expanded solder mask openings.

(R) 32 (pcbnew) *pcbnew will be enhanced to apply the essential process design minimums to modules when they are imported to the board.*

5.12.5 Legend Imaging

I am a proponent of WYSIWYG for CAD. When things are displayed the in the same way as the fabrication outputs will create the artwork, the designer has better control over the end result. Along those lines, an option will be provided controlling the display of legend drawings and text, with the following three settings:

Filled display. The icon can be a filled 'T'. In this mode, drawings and text are always displayed filled regardless of whether they span a mask opening or untented hole. This is simply the useful filled mode for silkscreen segments.

This mode is useful for seeing the result when no clipping is performed.

Sketch display. The icon can be an unfilled 'T'. In this mode, drawings and text are always displayed unfilled regardless of whether they span a mask opening or untented hole. This is simply the useful sketch mode for silkscreen segments.

This mode is useful for seeing the segments that make up a drawing or text.

Clipped display. The icon can be a 'T' that is part filled and part unfilled. In this mode, drawings and text are displayed filled when they do not invade a mask opening or untented hole. When segments of drawings or the little segments making up a text character invade a mask opening or untented hole, the segment will be displayed unfilled (sketched).

This mode is useful for when silkscreen will be clipped and indicating to the designer which segments whill be clipped. In this way, the designer can see whether the legend is still decypherable, while meeting the DFX rules.

(W) 33 (pcbnew) *pcbnew can be enhanced to provide three display modes for displaying silk screen (legend) text: filled display, sketched display and clipped display.*

Along the same lines for editing as display, the three modes can also affect the editing of drawings or text.

Filled editing. Under filled editing, with DRC on, drawings and text that is moved or created will resist being placed where portions would overlap with mask openings or untented holes.

Sketch editing. Under sketch editing, with DRC on, drawings will not resist being placed at any location.

Clipped editing. Under clipped editing, with DRC on, drawings will not resist being placed at any location; however, as the object is being moved or created portions begin clipped will be displayed in sketch mode wherease portions not being clipped will be displayed in filled mode. In this way, the designer can see while moving or creating an segment or text whether portions of that segment or text will be clipped.

These editing modes will be tied to the equivalent display modes. Note that with DRC off, drawings and text will not resist begin placed at any location in any mode. However, the same display rules apply regardless of the DRC setting.

(W) 34 (pcbnew) *pcbnew can be enhanced to provide three modes for editing equivalent to the three new display modes for silk screen text: filled editing, sketched editing and clipped editing. The editing modes will cooperate with the DRC items added for silk screen to assist the designer in the placement of legend.*

Drawing of Legend Components. When drawing Legend components, there are two approaches to addressing DFX rules for legend minimum line width, character height, and character spacing:

As designed. The first approach is to display legend components as designed (that is, as specified by the CAD system). In this case, when a line width is too narrow, or a character is too short, or characters are too closely spaced, it will still be displayed with the as-designed dimensions. The same is true for the width of regular line segments (drawings). The only exception to this rule for display is when the resulting line thickness is too narrow for the display. In such a case, the absolute minimums for display will be used instead of the size that cannot be displayed.[35]

As fabricated. The second approach is to display legend components as fabricated (that is, as required by the manufacturing process rather than as specified by the CAD system). In this case, when the line width is too narrow, it is increased to the minimum for display. When a character is too tall, it

35. Note that pcbnew has historically allowed a line to be set so small (e.g., 1 mil) that it cannot be selected reliability with the mouse even when the canvas is at its maximum zoom.

is increased to the minimum for display. When the character spacing is too tight, it will be increased to the minimum character spacing. The same is true for the width of regular line segments (drawings).

(R) 35 (pcbnew) *A general option will be provided to allow the designer to switch between "as designed" and "as fabricated" display approaches. The general options should not only apply to the legend layer, but should apply to all layers. That is, the display should be "as designed", or "as fabricated".*

Drawing of Legend Clearance Lines. The drawing of clearance lines is another useful tool for the designer to visualize the clipping that would result from moving legend text or drawings. pcbnew has not historically drawn any clearance lines for legend, so this is new territory. The drawing of Legend clearance lines will be in accordance with the following rules:

1. Legend clearance lines are not drawn unless the designer has selected the display of clearance lines on the toolbar.
2. Legend clearance lines are only drawn when a legend layer is the active layer.
3. Legend clearance lines are only drawn for the legend layer that is active.
4. A clearance line equal to the positional (xy-plane) registration of the legend is drawn around either:
   (a) the bounding box of the text;
   (b) a clearance around a silkscreen layer graphic line; or,
   (c) possibly the individual segments making up each letter of the text.
5. A clearance line equal to the finished hole size is drawn around every hole that is not fully tented. This means that hole clearance lines are not drawn for holes that are:
   (a) fully tented by primary mask;
   (b) fully tented (capped) by secondary mask; or,
   (c) filled and covered by primary mask.

   Holes that have both a primary and secondary mask opening are not considered.
6. A clearance line equal to the primary mask opening is drawn for every pad with a primary mask opening. Solder dams are not considered as usable for legend, so only the primary mask opening is provided.
7. A clearance line equal to the primary mask opening is drawn for every route path, score, bevel, perforation, slot, cutout and well or other feature (other than holes, which are handled above) for which a primary mask opening exists.
8. Dark zones and draw segments defined on the soldermask layer will have a clearance line equal to the primary mask registration drawn about them. Clear zones and draw segments will not be considered.

(N) 36 (pcbnew) *Clearance lines should be drawn around primary mask openings on the silk screen layers for assistance in developing legend.*

5.12.6 Legend Post-Processing

Plotting. At the time that fabrication outputs are generated, violations of DFX rules for legend can still be corrected. This can be accomplished as described in the following paragraphs:

Clipping:— At the time that fabrication outputs are generated, violations of the DFX rules for solder mask opening clearance can still be corrected. This can be accomplished by plotting a clear layer over the dark legend features equal to the soldermask opening increased in size by the silkscreen registration. This will result in clipping of the legend in any place where a violation occurs.

In general, however, it is very difficult for this clipping to meet the requirement for minimum line widths. A clear overlay created in this fashion can, for example, cut a line in half along its width rather than along its length, causing a violation of the minimum line width rule. For this and other reasons (e.g. wanting the fabricator to do the cutting), clipping of legend should be an option.

Another approach to clipping that does not stand to violate the minimum line width rules is to instead remove the offending components. Line and arc segments can either be truncated or split and truncated or removed entirely. Zones (dark zones on legend layers) can have cutouts created and the zone filling rules used for copper applied.

(R) 37 (pcbnew) *A post-processing option will be provided that selects whether clipping will be applied to silk screen (legend) layers when plotting. This can be a general "as designed" or "as built" option.*

Line Width and Character Height Adjustment:— At the time that fabrication outputs are generated, violations of the DFX rules for line thickness and character height can still be corrected.

Assembly Drawings:— One of the purposes of legend is to provide an image of the assembly drawings for use in ARE (Automatic Reverse Engineering) of lands and centroids for assembly. For this purpose an unclipped silk screen layer is better than a clipped one. The clipping is unnecessary for the ARE process. A clipped version of the silk screen is only useful to the assembler for review of the board fabrication parameters and for determining the legibility of the resulting image on the board surface.

Historically, pcbnew has only provided the silk screen layers as assembly drawings, but it will be enhanced to provide assembly drawings for the front and back of the board (as viewed through the board from the top).

5.13 Heat Sink Layers

The design of a number of power chips include a QFN package that contains a thermal grounding pad for dissipating heat into the board. One of the problems with this approach is that the multi-layer board is a sandwich of good thermal conducting copper layers and very poor thermal conducting dielectric layers. Therefore, to dissipate the heat transfered to the interior copper areas of the board, external copper areas must be provided to dissipate heat to the ambient. Solder resist also acts as a good insulator, so basic SMOBC designs are limited by the ability of solder mask to dissipate heat. One alternative is to place large planes of exposed copper, however, some finishes are not intended on being used in this fashion, or the soldering process (wave soldering mostly) can have significant issues with large exposed solderable areas. Then there is also the problem that exposed copper or finish can cause electrical shorts.

One way of providing the heat-sinking of inner copper layers is to use a electrically non-conductive, but heat-conducting, paste to fill thermal vias and provide a surface for dissipating heat to the ambient. The curable paste is applied using a basic stencil process.

5.13.1 Heat Sink Process

Here, heat-sink refers to a special type of heat-sink that is created using non-conductive (electrical) paste that is highly thermally conductive. The paste is stencil printed onto the board in large areas and presses through open vias and other holes in the aperture area. The result is a plate attached to the board that acts as a surface heat-sink. This is a very specialized process and is not supported by all board fabricators.

Where heat sinks printed in this fashion are too thick, they can interfere with the stencil printing process for solder paste. In some cases the heat sinks might need to be created after solder paste has been stencil printed.

5.13.2 Heat Sink Materials

5.13.3 Heat Sink DFX Rules

5.13.4 Heat Sink Design

5.13.5 Heat Sink Imaging

5.13.6 Heat Sink Post-Processing

5.14 Solder Paste Layers

The solder paste layers are assembly layers that consist of conductive paste applied to the exposed board surfaces in preparation for the reflow of surface mount devices, and, with special techniques, through-hole devices. pcbnew has historically modelled solder paste layers, but has not displayed the layers during editing, only generated plots during post-processing of the solder paste layers. Also, pcbnew has historically only generated pad masters for solder paste until recently (April 2010). In recent versions of pcbnew, the module editor and module library format has permitted altering the solder paste apertures. This feature, however, only permits adjusting the aperture by a fixed clearance value that is specific to the pad or to the module. Unfortunately, this alone is not sufficient to meet industry standard requirements, such as the IPC-7525 Stencil Design Guidelines [IPC-7525], which requires that some leaded SMT components (e. g. J-lead and Gull-wing) be adjusted by differing amounts in the length and width of the aperture. Also, FBGA and CSP require that a square aperture be used even though the pad is circular. Also, chip components should have bow-tie, home-plate, d-shaped or obround apertures to avoid the formation of solder balls and the occurrence of solder under the chip. Also, MELF components require a C-pattern to keep the cylindrical package from rolling during reflow or on the conveyor into the reflow oven.

In general the deficiencies in the single clearance parameter can be circumvented by unchecking the solder paste layer for the pad and adding a new pad, that only has a solder paste layer, with the proper dimensions. Even the change of circular FBGA pads to square solder paste apertures can be accommodated in this fashion. However, it does not accommodate use of bow-tie, home-plate, or obround (d-shaped) apertures for chip components because the module editor does not offer these as shapes. Changes to the module editor to permit the description of complex apertures can alleviate the problem somewhat. With diligent library design, these can be accomplished. However, it does not permit the changing of bow ties to home plates, or any such approach. Also, providing for self-cleaning stencils (with rounded corners) is not an option with this approach either.

Another thing that pcbnew could not accommodate, historically, is the design rule checks necessary to ensure that a stencil can be fabricated. Given a single-step stencil of a given thickness, the aspect ratio and area ratio of a given aperture must be checked to ensure proper transfer printing of solder paste. Some tricks (e. g. overprinting) can be taken to keep to a given stencil thickness. Common single-step stencils are 6mil thick. Thinner stencils are needed for finer components. Step-stencils, that have different thicknesses require specialized keep-outs across the transition from one thickness to another that are extremely difficult to represent in a CAD system and for the designer to fathom. Two-stage printing can handle a wide variation in stencil thickness without keep-outs, but additional information must be used to separate the solder paste bricks into two levels. Two-stage printing is normally used for intrusion printing, where TH are filled with solder paste in the second stage, while etched cutouts in the bottom stencil protect the solder paste bricks formed with a thinner stencil in the first stage.

5.14.1 Solder Paste Process

The application is typically performed using a stencil and a process that applies paste to the board via the openings in the stencil. For example, a stencil made by laser-cutting a thin steel sheet is placed in a jig in contact with the board's surface. Solder paste is dispensed into the jig and a squeegee is used to wipe the paste across the surface of the stencil. This results in the deposit of solder paste in the openings of the stencil. The amount of paste applied to a given opening is proportional to the area of the opening and the thickness of the stencil.

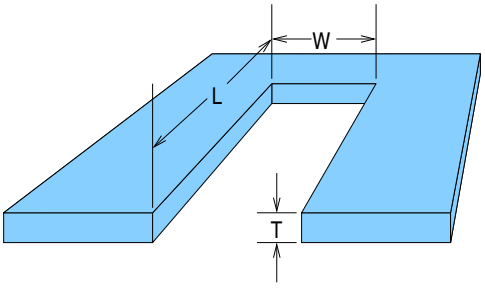

Figure 65: Stencil Opening

5.14.2 Solder Paste Materials

Solder paste materials are normally a combination of flux and solder solids forming a paste. Solder paste is available with or without lead content. It is available in various constituencies characterized by the variation and mean size of the conductive solids within the paste. Stencil design must typically consider the constituency of the material used.

5.14.3 Solder Paste DFX Rules

In general, given the target stencil thickness (or thicknesses for step-stencils), it is possible to calculate aspect ratio and area ratio for any given aperture and determine whether it falls withing the process window. This is a design check of the soler paste apertures to the requirements of IPC-7525 [IPC-7525].
The stencil opening is modelled in *Fig. 65(90)*, where,

- T, is the thickness of the stencil.
- W, is the width of the opening.
- L, is the length of the opening.

Volume. The volume of solder paste that fills the opening is simply

$$V_S = T \times A_S = T \times (L \times W). \tag{76}$$

Area Ratio. The *area ratio*, α_A, is defined as the area of the land exposed to the aperture, $A_L = W \times L$, over the area of the walls, $A_W = 2T(L + W)$, or

$$\alpha_A = \frac{A_L}{A_W} = \frac{L \times W}{2T(L + W)} = \frac{1}{2T(\frac{1}{W} + \frac{1}{L})}. \tag{77}$$

$$\text{When } L \gg W,\ \alpha_A \approx \frac{W}{2T}. \tag{78}$$

For shapes other than rectangles, the aspect ratio is the area of the shape divided by the thickness of the stencil times the perimeter of the shape, or

$$\alpha_A = \frac{A_S}{T \times P_S}, \tag{79}$$

where, A_S, is the area of the shape, and, P_S, is the perimeter or circumference of the shape. So, for a circle with radius, r, and diameter, $d = 2r$, the area ratio is

$$\alpha_A = \frac{A_S}{T \times P_S} = \frac{4\pi r^2}{T \times 2\pi r} = \frac{d}{T}. \tag{80}$$

Other shapes can easily be calculated using:

$$\alpha_A = \frac{A_L}{A_W} = \frac{A_S}{T \times P_S}, \tag{81}$$

where, A_S is the area of the shape, and P_S is the perimeter or circumference of the shape. The maximum stencil thickness for any given shape based on area ratio can be calculated from:

$$T_{max} = \frac{A_S}{\alpha_{min} \times P_S} = \frac{A_S}{0.66 P_S} = 1.5 \times \frac{A_S}{P_S}, \tag{82}$$

where, α_{min} is the minimum allowable area ratio (which is about 0.66).

Aspect Ratio. The *Aspect Ratio* is simply the ratio of the width of the aperture to its depth, or

$$\alpha = \frac{W}{T}. \tag{83}$$

The minimum allowable aspect ratio for successful solder paste printing is 1.5, so the maximum stencil thickness is:

$$T_{max} = \frac{W}{\alpha_{min}} = \frac{W}{1.5} = 0.67W. \tag{84}$$

Note that for a circular aperture, the area ratio and aspect ratio are identical:

$$\alpha = \frac{d}{T} = \alpha_A = \frac{d}{T} \tag{85}$$

Therefore, decisions based on aspect ratio are applicable to circular transfer patterns and decisions based on area ratio are applicable to transfer patterns that are more rectangular in shape.

Calculation of Stencil Thickness

Intrusive Soldering

Wave Soldering Fixtures

Selective Wave Soldering

Single Miniwave Soldering

5.14.4 Solder Paste Design

The design of solder paste stencils for SMT devices is straightforward and detailed elsewhere in this section. There are some special considerations for the design of stencil apertures for specialized devices. These are primarily for use with CBGA and QFN devices that have a thermal grounding pad.

5.14.5 Solder Paste Imaging

Solder paste stencil openings are normally created using a padmaster that contains all of the pads to which solder paste must be applied: normally, SMT pads. Not all surface mount device pads have solder paste applied. For example, BGA device lands are typically described as NSMD pads, that do not have solder paste applied. The solder balls that are attached to the bottom of the BGA provide the necessary material to form a proper joint. TH devices typically require wave soldering when combined with SMT on the same board, and do not have solder paste applied. There is a process for placing solder paste on the lead of TH components and thus assembling TH devices, SMDs and BGA in a single reflow cycle.[36]

Providing an image for creation of the stencil, therefore, consists of identifying the SMT pads that will have solder paste applied. The pcbnew module editor currently models pads that are SMT and require solder paste by making them "present" on the corresponding (front or back) solder paste technical layer with an opening that is based on the copper pad. The identified pads are typically flashed on an image, where the dark pads flashed on the image represent the solder paste stencil openings through which paste will be applied.

The positioning of solder paste on pads is not critical. Proper primary mask design and the balling effect of the solder during reflow will cause the solder paste to pull toward the pad and the terminal of the SMT. The floading and surface tension effects of the solder during reflow will also pull a slightly out of place component towards the reference position.

Of more importance is the volume of solder paste that is applied. The volume can be controlled in several ways: the stencil thickness, the relative size and shape of the opening in proportion to the land. The objective of the process is to form a fillet around the extremities of the SMT device terminal acting as an electrical and mechanical joint. The amount of paste applied must be sufficient to form the fillet. As a result, the amount of solder paste is not directly related to the area of the land, but is more related to the area of the land beyond the flush mounting of the terminal on the device (the amount of solder between the flush surface union of the terminal and land is typically far less than that included in the fillet portion of the joint). While a stencil can be made with shapes proportional to the land form most lands; creating shapes that do not apply too much paste for large thermal or ground pads under devices represents a challenge. Typically, cross-hatching or some other means of reducing the amount of paste delivered through the stencil needs to be applied to large lands that are not expected to form a fillet.

If a localized amount of paste is applied at the center of the larger land, and the component is inserted on top of the mound of paste, it could squeeze out in various directions unpredictably as the device is pressed to the surface of the board by the assembly robot. It is prefferable to have a distributed mesh of smaller deposits of paste that will better control the spreading process. One way to accomplish this is by placing criss-crossing segments over the land pattern to reduce both the overall volume of paste applied, and distributing the paste that is applied over a larger areas to avoid toothpaste-tube effects. Another way is to define smaller pseudo-pads on the solder paste layer that are distrbuted over the surface of the larger land. Historically, pcbnew could handle the later process (placement of smaller pseudo-pads), but could not handle the first process (criss-crossing of "clear" segments). However, historically pcbnew has not displayed the resulting stencil opening so that the designer can see the effective stencil opening.

Adjusting specific stencil openings to a given land pattern (footprint) is the business of libraries. Library components can be created with solder paste opeinings defined for given fabrication processes and materials. Once the opening shape has been defined for one process and material, it is typically a proportional calculation to adjust the opening for another process or material. IPC provides recommendations for land pattern and paste openings for a wide range of surface mount devices. Rather than check each and every component, land pattern and paste opening, it is typical for an assembler to simply require that the IPC standard land patterns and paste openings be used.

Imaging the solder paste stencil openings consists of generating basic pad-master-like artwork that describes the stencil openings using simple apertures.

There are three was of making a stencil:

1. chemical milling,
2. laser cutting; and,
3. mechanical milling.

Chemical Milling: For photo-imagable chemical milling of stencils, this artwork is sufficient and can be used to develop the chemical mask.[37]

Laser Cutting: For laser-cutting of sheet steel stencils, the arpetures need to be converted into a laser cutting path for

36. But I have no idea how to model that just yet...

37. See §F on page 153 for a description of chemical milling.

the laser-cutting tool that typically creates openings in the stencil. Ostensibly, laser-cutting machines are G-code programmed CNC machines that peform the cutting. Assembler CAM systems should be able to convert Gerber (RS-274X), DPF, or CAM format artwork into and optimized set of G code driven laser-cutting paths; however, if necessary we can make pcbnew generate the G code. Laser-cut stencils are often made of thin sheet steel.

Mechanical Milling: For machine milling, a G code control program would be used to mill out the openings using a CNC mill. Mechanically milled stencils are normally made of brass (much softer than steel), but any cheap material that can be found in 5-10 mil thicknesses will work (including beer cans). The walls of the stencil must be as smooth as possible. When mechanically milling stencils resulting in a slanted wall, the wall should be slanted outward toward the bottom of the stencil to avoid lifting the paste when the stencil is withdrawn from the board. As conical tools are often used for milling the brass stencil, the stencil should be milled from the back side (the side adjoining the board surface) to acheive a slant in the proper direction. This involves mirroring the image when creating G code. 5mil thick brass can be used for mechanical milling.

Stainless steel is the preferred metal for chemical etch and laser cut technology. Other metals, as well as plastics, may be specified. For electroform technology, a hard nickel alloy is preferred. Frames can be tubular or cast aluminium with the border permanently mounted using an adhesive. Some foils can be mounted into a tensioning master case and do not require a border or a permanent fixture from the foil to the frame. The stencil border is polyester standard material; but stainless steel is optional. The fabrication process for stencils may involve additive or subtractive methods. In additive processes such as electroforming, metal is added to form stencil foils. In subtractive processes, metal is removed from foils to create apertures. Laser cut and chemical etch are examples of subtractive processes.

Chemically etched stencils are produced using photo-imagable resist laminated on both sides of metal foils cut to specific frame sizes. A double-sided phototool, held in precise alignment usually with registration pins, is used to expose the stencil aperture image onto the resist. Aperture images exposed on the resists are reduced in size compared with the desired aperture dimensions, accounted by an etch factor. The etch factor describes the amount of lateral etching that takes place as the chemical etches through the thickness of metal foil. The exposed resist is then developed, leaving bare metal where apertures are desired. The metal foil is etched from both sides in a liquid chemical, creating apertures as specified. The remaining resist is then stripped away and a stencil foil is produced.

Laser cut stencils are produced from data run by software of the laser equipment. Unlike chemically etched stencils, no phototool is required. Since stencils are cut from one side only, tapered aperture all is an inehrent part of laser cut stencils. Unless otherwise specified, apertures are larger on the contact side that on the squeegee side.

Where a mixture of standard and fine-pitch assemblies are present on a board, the stencil fabrication process may be a combination of laser cut and checmical etch. The stencils produced are referred to as laser-chem combination or hybid stencils.

Electroforming is an additive stencil fabrication method utilizing photo-imagable resist and an electroplating process. Photo-imagable resist is placed on a metal mandrel. Thickness of the resist is resist than the final stencil thickness desired. The apertures are image onto the resist and the resist is developed, leaving resist pillars where apertures are desired. The mandrel with resist pillars is placed in a nickel plating tank where nickel is electroplated onto the mandrel. When the desired stencil thickness is reached, the mandrel is removed from the plating tank. Lastly, the resist pillars are stripped and the nickel stencil foil is separated from the mandrel.

Stencil Data See "Stencil Design Guidelines." [IPC-7525A]

Regardless of the stencil fabrication method used, Gerber data is the preferred data format. Possible alternative formats are GenCAM, DXF, HPGL, Barco, etc; however, they may need to be converted to Gerber format prior to the stencil manufacturing process. Gerber data describes the file that provides a language for communicating with the photo plotting system to produce a tool for chemically etched stencils. It is also used to produce the laser cut or electroformed stencils. While the actual data format may vary from file to file depending on the software package or designer, the data format commonly used by photo plotter and laser equipment is known as Gerber.

There are two standard Gerber formats available. RS-274D requires a data file listing the X-Y coordinates of the stencil where apertures are to be placed and formed, and a separate Gerber aperture list that describes the size and shape of the various Gerber apertures used to prepare the image. RS-274X in this format the Gerber aperture list is embedded in the data file. The aperture list is an ASCII test file containing D codes that define the size and shape of for ll apertures used within the Gerber file. Without the aperture list, the software an photo plotting system cannot read the Gerber data. Only the X-Y coordinate would be available with no size and shape data.

The solder paste layer data is necessary to produce a stencil. If fiducial marks are required on the stencil, they should also be included in the solder paste layer.

Data can be transmitted to the stencil supplier via modem, FTP (file transfer protocol), e-mail attachment or disk. To ensure data integrity after transmitting and due to the large size of data files, it is suggested that the files be compressed prior to sending data. It is recommended that the full data file (the solder paste, solder mask, silk screen and copper layers) sent to the printed circuit board manufacturer be supplied to the stencil manufacturer. This allows the stencil manufacturer to optimize or make recommendations on aperture sizes based on actual pad sizes for the SMT land.

In those cases where it is desired to have more than one image on the stencil, the stencil patterns will be panelized and included in the data file. In those instances where the data file does not already contain the panelized stencil design, a readme file, panel drawing or order information must specify the location of the two or more designs. This could be a reference from the edge of the frame, distances between patterns, etc.

In those cases where more that one image of the same design is to be printed, the data file for stencil fabrication should contain the stencil design in the step-and-repeat array. In those instances where the data file does not contain the step-and-repeat pattern, a readme file, panel drawing, or order information should specify:

- Total number of steps for the final array.
- Number of steps in the X-direction along with dimensions from a specific feature to corresponding feature (such as fiducial marks, component pad location, etc.)
- Number of steps in the Y-direction along with dimensions from a specific feature to corresponding feature (such as fiducial marks, component pad location, etc.)

In those cases where image orientation is not parallel to the frame or the step-and-repeat is not rectilinear (one ore more images is rotated), the data for stencil fabrication should contain the oriented image. In those coasts where it does not, a readme

file, panel drawing or order information should specify this information (X- and Y-offsets) referencing stencil features.

To accomodate specific printers, the stencil image may have to be located in different positions inside the frame:

1. center image.
2. center board/panel — requires board/panel outlines.
3. offset board/panel — requires board/panel outlines and reference locations.

In those cases where this data is not included in the Gerber data, a readme file, panel drawing or order information should specify this information referencing stencil features.

Stencil should contain identification information such as part number, revision number, thickness, supplier's name and control number, data and method of manufacture.

A general aperture design guidelines for SMT components is shown in *Tab. 5(94)*. Some of the factors influencing stencil aperture design are: component type, pad footprint, solder mask opening, board finish, aspect/area ratio, solder paste type, and user process requirement.

The volume of solder paste applied to the board is mainly determined by the aperture size and foil thickness. Solder paste fills the stencil aperture during the squeegee cycle of the print operation. The paste should completely release to the pads on the board during the board/stencil separation cycle of the print operation. From the stencil viewpoint, the ability of the paste to release from the inner aperture walls to the board pads depends primarily on three major factors:

1. the area and aspect ratios for the aperture design.
2. the aperture side wall geometry.
3. the aperture wall finish.

Both area ratio and aspect ratio are illustrated in *Fig. 65(90)*. A general design guide for acceptable paste release is ≥ 1.5 for aspect ratio and ≥ 0.66 for area ratio. The aspect ratio is a one-dimensional simplification of the area ratio. When the length is much greater than the width, the area ratio, $W/2T$, reduces to a factor of the aspect ratio, W/T.

When the stencil separates from the board, paste release encounters a competing process: solder paste will either transfer to the pad on the board or stick to the aperture side walls. When the pad area is greater that 0.66 of the inside aperture wall area, a complete paste transfer should occur.

$$\text{Aspect ratio} = \frac{W}{T}, \tag{86}$$

where W is the width of the aperture, T is the thickness of the stencil.

$$\text{Area ratio} = \frac{A_A}{A_W} = \frac{L \times W}{2 \times (L + W) \times T}, \tag{87}$$

where L is the length of the aperture, A_A is the area of the aperture, A_W is the area of the walls.

As a general design rule, the aperture size should be reduced compared to the board pad size. The stencil aperture is commonly modified with respect to the original pad design. Reductions in the area or changes in aperture shape are often desirable to enhance the process of printing, reflow, or stencil cleaning. For instance, reducing the aperture size will decrease the possibility of stencil aperture to board pad misalignment. This reduces the change for solder paste to be printed off the pad, which may lead to solder balls or solder bridging. Having a radiused corner for all apertures can promote stencil cleaning.

For leaded SMT's, e. g., J-leaded or gull-wing components with 1.3 to 0.4mm (51.2 to 15.7mil) pitch, the reduction is typically 0.03 to 0.08mm (1.2 to 3.1 mil) in width and 0.05 to 0.13mm (2.0 to 5.1mil) in length. For PBGA (Plastic Ball Grid Array), reduce circular aperture diameter by 0.05mm (2.0 mil). For CBGA (Ceramic Ball Grid Array), increase circular aperture dimension by 0.05 to 0.08mm (2.0 to 3.1mil) when this does not interfere with the solder mask and/or increase the stencil foil thickness to 0.2mm (7.9mil) and go one to one with the board pad. Refer to IPC-7095 for solder paste volume requirements. For FBGA (Fine-pitch Ball Grid Array) and CSP (Chip Scale Package) square aperture with the width of the square equal to, or 0.025mm (0.98mil) less than, the diameter of the pad circle on the board. The square should have rounded corners. A guideline is 0.06mm (2.4mil) radiused corners for a 0.25mm (9.8mil) square and 0.09 (3.5mil) corners for a 0.35mm (14mil) square. For chip components such as resistors and capacitors, several aperture geometries are effective in reducing the occurrence of solder balls. All these designs are aimed at reducing excess solder paste trapped under the chip component. The most popular designs are show in *Fig. 66(94)*. These designs are commonly used for no-clean processes.

For MELF and mini-MELF components, C-shaped apertures are suggested (see *Fig. 66(94)*). Dimensions of these apertures should be designed to match the geometry of component terminals rather than the land.

The glue aperture chip component has a glue stencil that is typically 0.15 to 0.2mm (5.9 to 7.9mil) thick. The glue aperture is placed int he center of the component solder pads. It is 1/3 the spacing between pads and 110% of the component width.

It is desirable to have a process where SMT and THT devices can both be:

1. provided with printed solder paste;
2. placed on or in the board; and,
3. reflowed together.

The objective of stencil printing of solder paste for the intrusive reflow process is to provide enough solder volume after reflow to fill the hole and create acceptable solder fillets around the pins. *Tab. 6(95)* shows process window for a typical intrusive soldering process.

Where more than one board or panel are placed on one stencil, a minimum of 50mm (2.0") is recommended between images.

Unless otherwise specified, additional design guidelines are:

- Minimum 20mm (0.79") border is recommended from the edge of the frame to the edge of metal.
- Minimum 50mm (2.0") from the inside edge of glue border to the edge of image is suggested for solder paste and squeegee travel.

5.14.6 Solder Paste Post-Processing

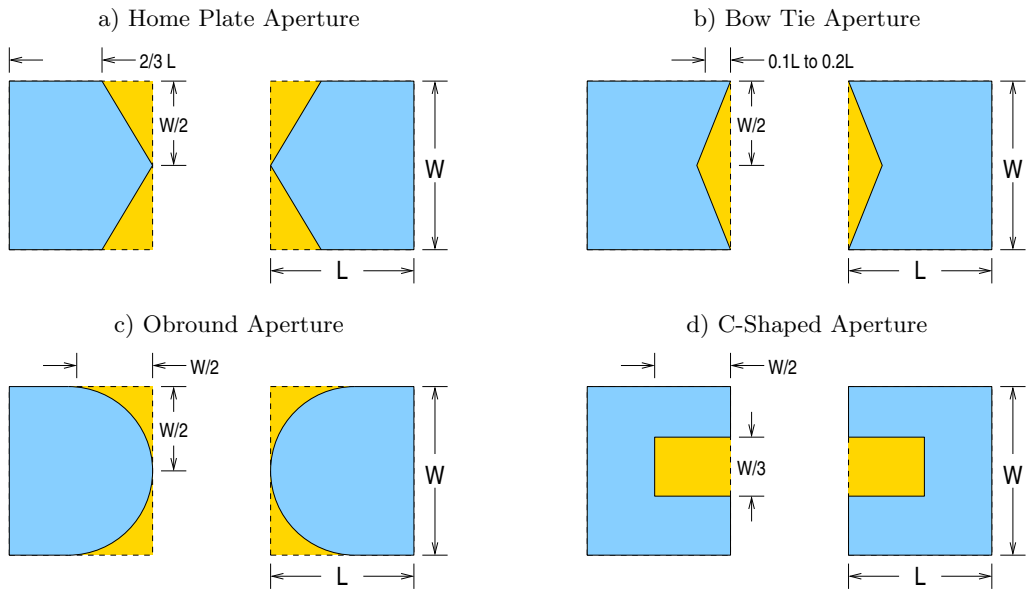

Figure 66: Various Chip Apertures

Table 5: Adjustments to Apertures

Part	Pitch	W	L	W_A	L_A	Thickness	Aspect Ratio	Area Ratio
PLCC	1.25mm 50mil	0.65mm 26mil	2.00mm 80mil	0.60mm 24mil	1.95mm 78mil	0.15–0.25mm 6–10mil	2.3–3.8	0.88–1.48
QFP	0.65mm 26mil	0.35mm 14mil	1.50mm 60mil	0.30mm 12mil	1.45mm 58mil	0.15–0.175mm 6–6.5mil	1.7–2.0	0.71–0.83
QFP	0.50mm 20mil	0.30mm 12mil	1.25mm 50mil	0.25mm 10mil	1.20mm 48mil	0.125–0.15mm 5–6mil	1.7–2.0	0.69–0.83
QFP	0.40mm 16mil	0.25mm 10mil	1.25mm 50mil	0.20mm 8mil	1.20mm 48mil	0.10–0.125mm 4–5mil	1.6–2.0	0.68–0.86
QFP	0.30mm 12mil	0.20mm 8mil	1.00mm 40mil	0.15mm 6mil	0.95mm 38mil	0.075–0.125mm 3–5mil	1.5–2.0	0.65–0.86
0402	n/a	0.50mm 20mil	0.65mm 26mil	0.45mm 18mil	0.60mm 24mil	0.125–0.15mm 5–6mil	n/a	0.84–1.00
0201	n/a	0.25mm 10mil	0.40mm 16mil	0.23mm 9mil	0.35mm 14mil	0.075–0.125mm 3–5mil	n/a	0.66–0.89
BGA	1.25mm 50mil	0.80mm 32mil	0.80mm 32mil	0.75mm 30mil	0.75mm 30mil	0.15–0.20mm 6–8mil	n/a	0.93-1.25
FPBGA	1.00mm 40mil	0.38mm 15.2mil	0.38mm 15.2mil	SQ 0.35mm 14mil	SQ 0.35mm 14mil	0.115–0.135mm 4.5–5.3mil	n/a	0.67–0.78
FPBGA	0.50mm 20mil	0.30mm 12mil	0.30mm 12mil	SQ 0.28mm 11mil	SQ 0.28mm 11mil	0.075–0.125mm 3–5mil	n/a	0.69–0.92

Note:

1. All units are millimeters (mils).
2. W and L are the width and length of the land.
3. W_A and L_A are the width and length of the (reduced) aperture.
4. It is assumed that the fine-pitch BGA pads are not solder mask defined.
5. N/A implies that only the area ratio should be considered.

Table 6: Process Window for Intrusive Soldering

Parameter	Maximum Limits	Desirable
Hole Diameter	0.65–1.60	0.75–1.25
Lead Diameter	FHS - 0.075	FHS - 0.125
Paste Overprinting	6.35	≤ 4.0
Stencil Thickness	0.125–0.635	0.15, 0.20 for fine-pitch

5.15 Adhesive Layers

The adhesive layers (front and back) represent adhesive (or glue) that is applied to the board before stuffing surface mount components to retain the components during reflow, particularly when wave soldering or the bottom side of a multilayer board is populated for a single reflow cycle.

The adhesive layers are assembly layers.

5.15.1 Adhesive Process

Processes for applying adhesive (or glue) consist of the application of glue dots to the surface of the board, positioned under devices that must be retained. Glue is typically applied using a dispenser that applies a glue dot to the board at specified locations. The amount of glue applied in a given glue dot is controllable.

There are three ways of applying SMA (Surface Mount Adhesive); these include, stencil printing, dispensing and needles.

Because stencil printing requires contact between the stencil and the board, and because stencil printing is also normally used for solder paste, stencil printing of SMA is only used for wave soldering all components. The advantages of reflow soldering and selective wave soldering for mixed technologies, means that wave soldering is only normally used for mass production where speed is its major advantage.

For double reflow (where each side of a multilayer board is reflowed independently and sometimes for miniwave reflow), SMA is most widely applied using dispensers. SMA is applied after stencil printing of the solder paste and before pick and place.

Needle printing is slow (because it supports typically only a small range of dot volumes, requiring a large number of dots per component), so is not popular.

Although stencil printing can achieve finished shapes other than a cylinder, design processes can easily treat all processes as delivering dots of a specified volume and variation.

5.15.2 Adhesive Materials

Adhesive materials typically consist of a non-conductive, epoxy-based glue.

5.15.3 Adhesive DFX Rules

Historically pcbnew has not provided any mechanisms for design rule checks associated with the application of adhesive. Several considerations for designing with adhesive are:

- Traces should not be run between SMT pads, particularly on thicker external copper layers, because of the irregularity that it creates in the contact surface under the SMT device.
- Bow-tie apertures should be used for solder paste with single-dot adhesive to avoid having adhesive contact the solder joint. Home-plat apertures should be used with double-dot.
- Vias (unless filled, made planar and covered over with mask) should never be placed under SMT devices, and particularly at a position that would interfere with the application of glue dots.
- For use in simultaneous double-sided reflow lines, a portion of the adhesive must be expose to UV light, so positioning of SMA dots along the periphery of a component might be necessary.

5.15.4 Adhesive Design

5.15.5 Adhesive Imaging

5.15.6 Adhesive Post-Processing

5.16 Courtyard Layers

The courtyard layers are assembly layers. The primary concern of the courtyard layer is the placement and spacing of components on the PCB in the PWA and therefore are primarily associated with the task of manual or automatic component placement. There are two types of courtyards:

Courtyard. The traditional IPC courtyard that consists of an contour drawn around the outline and pads of a component, adjusted by an excess to accommodate component tolerances and to a degree placement accuracy. This is the contour from which spacing calculations are performed. Courtyards are a component-attached concept.

Room. The popular "room", where an area is described and the candidacy for placement within the "room" is specified. Some properties of components that are considered are: height, power dissipation, power net, thermal model. "Room" are a specialized keep-in/keep-out, that is distinguished from the board level component keep-in/keep-outs that are described for the "Keep-out" layers (*Sec. 5.6(69)*). Rooms are only a board-, array- or panel-attached concept.

Historically, pcbnew did not support courtyard layers. Now there is a Courtyard layer for each of the Top (Front) and Bottom (Back) board sides for placement of traditional components, and Internal (Inner) layers for placement of embedded components. The courtyard layers provide information on the spacing of components and component keep-ins and keep-outs.[38]

SMT Placement Courtyard The placement courtyard was introducted by IPC in the IPC-SM-782 Surface Mount Design Land pattern Standard in 1987 and fine-tuned in the 2005 release of IPC-7351.

The primary use of the placement courtyard was to provide the PCB designer a guideline for placing land patterns next to each other with enough room to compensate for component tolerances.

Courtyard outlines are used to ensure that all parts will fit, but they do not compensate for assembly machine heads and manufacturing allowances. Each assembly manufacturer has their own unique processes that require various allowances. Placement courtyards are not meant to touch each other or overlap. They should have a space gap in-between them so when the PCB designer runs a design rule check for body-to-body clearance, there should be no errors found.

The standard courtyard line width is 0.05mm and it is placed on a layer dsignated by the CAD vendor. The placement courtyard is used as a CAD visual graphic aid for part placement and never post processed.

Land pattern courtyards are determined starting with the component pattern maximum boundary (which is a boundary box containing the component outline). Added to this is the courtyard excess to result in a courtyard minimum area which is another bounding box enclosing the first. An additional manufacturing allowance is added to this to result in the courtyard manufacturing zone.

It is the courtyard manufacturing zone that is critical for the assembly process. This is the body-to-body clearance that is set in design rules for design rule checking. The size of the manufacturing tolerance must come from the assembly shop that is going to be used to populate the parts on the PC board. Every assembly shop has different assembly tolerances, but the average is 0.1mm.

The assembly process makes it very difficult to determine placement courtyards for TH components. It is easy to determine SMT to SMT and even TH parts on the top side and SMT parts on the bottom side. Since TH parts require holes that go all the way through the PC board, the TH part top side courtyard would be different than the bottom side courtyard due to the wave solder process used to solder TH component leads. If wave solder is used for the TH component leads, the SMT parts mounted on the bottom side must have a 5mm (0.200") clearance between the edge of the TH pad and edge of SMT pad. If a selective wave is used there is a different tolerance between the pads depending on the assembly shop requirements.

Therefore, building in placement courtyards for TH parts is difficult due to many variables. A PCB designer must use common standard rules provided by the assembly shop when performing part placement. The assembly shop should always approve the part placement prior to routing any traces on the board. This is an official check point that must not be avoided.

Courtyards should have a height associated with them. Also, component placements should specify a height from the board surface. This is mostly to accommodate IDF interchange as well as to permit design rule checks on thermal objects such as heat sinks or other mounting hardware that might be offset from the board surface.

5.16.1 Courtyard Process

Orientation. There are a number of reasons to want to specify the orientation of a component within an area (or "room").

Selective Wave Soldering. Designing for selective wave soldering has some basic spacing rules that can be applied to permit the design of wave soldering fixtures and to design for selective wave or manual soldering, as discussed in *Sec. 3.6.6(45)*. The easiest of these rules is where the spacing of TH components from SMT components is dependent on the component height. One more complex rule is that two-terminal SMT components should always be oriented perpendicular to an adjacent zone of TH pins so that if the selective wave touches the component, only one terminal will melt and the other will not, keeping the component from being washed away. "Rooms" that specify allowable orientations can be used to direct an autoplacer or guide the designer to orient components with this orientation.

Assembly Verification Polarized two-terminal components have always presented challenges for assembly. Components reversed on their mounts can cause serious problems and component damage. Verification of polarity or anode and cathode is facilitated by markings on the silk screen and assembly drawings. Because these markings are typically viewed and verified by a human operator, it is advantageous to orient all similar polarized components, particularly SMT diodes, with the same orientation. The ability to specify a room with an orientation helps to facilitate this when directing an autoplacer or guiding the designer during component placement.

Power Dissipation When considering thermal models for estimating the board operating temperature as well as maximum junction temperatures, it is essential to first consider the minimum, nominal, and maximum power dissipation of each component. Of course, power dissipated in a component is a direct conversion of electrical power to heat. All heat sources on a typical PCB are from components. Regardless of the model for heat dissipation, however, the stead-state temperatures and temperature gradients experienced by the PWA are more a function of the power dissipated by the part. Therefore, a number of "rooms" support the ability to specify the maximum power dissipation of the room. This is what SPEECTRA DSN does. The usefulness of specifying a room's maximum power dissipation is questionable for several reasons:

1. It is power density that affects temperature gradients under steady-state conditions more than the power within a given

38. Sometimes referred to as "rooms."

area. For example, specifying the maximum power dissipation of the left half of a board, does not stop all high-power-dissipation components from being placed in one small corner.

2. Some components funnel their heat into the board (like QFN packages with thermal ground pads), and some into the ambient (like a flip-chip BGA), possibly through heatsink. It makes a difference, but SPEECTRA DSN rooms don't care.

It would have been far better to specify the maximum and minimum power density for the room, where the power density of a particular placement is determined as the power dissipation for the device divided by th area occupied by the component plus half of the space to other components: $p_D = P/(A_{comp}+A_{spac}/2)$ in units such as mW/□.

Mechanical Analysis

Power Dissipation and Heat

Components that dissipate their heat into the board.

Components that dissipate their heat into the ambient.

EMC

Voltage reference planes an EMC.

Placement Process The placement process is primarily concerned with board assembly, fixtures, and final product mechanical and thermal requirements, including board assembly, test, and environmental considerations. The placement process defines component placement zones that describe the electrical, mechanical and thermal requirements influencing the rules for placement of components. The layer is intended on capturing a simply set of rules that restrict the placement of components. This simple set of rules is then applied as part of a DRC (Design Rule Check) that is performed either externally, when exchanging information with a foreign system, off-line, when checking a design for its adherence to the rules, or on-line while manually editing the design. This set of design rules can be exchanged with other PCB design systems, whether electrically focused (ECAD), mechanically or thermally focused (MCAD), or manufacturing focused (CAM). Therefore,

(R) 38 (pcbnew) *The internal representation of placement courtyards in pcbnew must support the models provided by the intended exchange file formats: GenCAM [GenCAM], GenX [GenX], 258X [258X], IDF [Kehmeier and Makowski, 1998], and SPEECTRA DSN [SPEECTRA].*

IDFv3.0 [Kehmeier, 1996] defines placement areas using keep-ins as follows:

Placement outline. This section of the file defines a placement outline for the board or panel. Each placement outline specifies a region within which components must be placed, and consists of a simply closed curve made up of arcs and lines plus a height restriction. Portions of placement outlines on a panel that lie on a board in the panel are inherited by that board. Multiple placement outlines may be defined. The outline height is used to exclude components from the outline that, when mounted, exceed this height. If this field is missing, there is no height restriction on the outline.

Placement outlines of this type will be handled by component keep-ins in the "Keep-out" layer. See *Sec. 5.6(69)*.

Placement group area. This section of the IDF version 3.0 file specifies an area where a group of related components is to be placed. For example, it may be desirable to place all analog components in a particular area for thermal considerations. Each placement group area consists of a simple closed curve made up of arcs and lines along with a name designating the group of components to be placed in that area. Multiple areas are allowed.

Placement group areas are more closely related to "rooms" and will be modelled in the "Courtyard" layer.

IDFv4.0 [Kehmeier and Makowski, 1998] defines placement areas using keep-ins as follows: The keep-in entity is used to represent an area on the board or panel in which component instances or board features must be located. There are four predefined keep-in types, that are described as follows:

All components. An area in which all component instances must be placed. The keep-in may affect any valid component placement layers. The keep-in has an optional height to specify a maximum placement height value. All components with a molded height less than this value must be placed within the area. If a height is not specified, all component instances must be placed within the area.

Component keep-ins of this sort are handled by the "Keep-out" layer. See *Sec. 5.6(69)*.

Specific components. A component group area within which a specific set of component instances must be placed. The keep-in may affect and valid component placement layer or layers. The keep-in uses the group property to specify a list of component references that are to be placed within the area.

Specific component keep-ins are modelled as "rooms" in the "Courtyard" layer.

Component rotation. A component rotation area within which component instances must be placed according to as set of predefined orientations. The keep-in may affect any valid component placement layer or layers. It specifies a set of acceptable rotation values for the component placement direction (0 and 180 degrees, for example).

Rotation-restricted component keep-ins are modelled as "rooms" in the "Courtyard" layer.

IDFv3.0 [Kehmeier, 1996] defines placement keep-outs as follows:

Placement keep-out. This section of the IDF version 3.0 file defines a placement keep-out for the board or panel. Placement keep-outs specify the regions of the board where components cannot be placed. A keep-out can apply to all components, or to only those components above a specified height.[39] Placement keep-outs can exist on the top, bottom, or both top and bottom of the board or panel. Each keep-out consists of a simple closed curve made up of arcs and lines along with a height restriction. Portions of placement keep-outs on a panel that lie on a board in the panel are inherited by that board. Multiple keep-outs are allowed. A keep-out height is used to exclude components from the keep-out that, when mounted, exceed this height. A value of 0.0 indicates that all components are to be excluded.

Placement keep-outs of this form are implemented as keep-ins or keep-outs in the "Keep-out" layer. See *Sec. 5.6(69)*.

IDFv4.0 [Kehmeier and Makowski, 1998] defines placement areas using keep-outs as follows: The keep-out entity is used to represent an area on the board or panel in which component instances or board features may not be located. There are seven predefined keep-out types that are described as follows:

Components by height. An area in which no component instance having a mounted height greater than the specified keep-out is allowed. The component mounted height is equal to the maximum height of the shape of the component part

39. In this case, there is no difference between the keep-out and a placement outline.

plus the mounting offset of the component instance. The keep-out may affect any valid component placement layer or layers. It specifies a maximum placement height value, or when not specified (or zero), the restriction applies to all component instances, regardless of height.

Component height keep-outs are handled by the "Keep-out" layer. See *Sec. 5.6(69)*.

Neither GenCAM nor GenCAM-XML [GenCAM, GenX] define any component design keep-outs or keep-ins that are equivalent to placement keep-outs or "rooms".

SPEECTRA DSN [SPEECTRA] defines keep-outs a little different, as follows:

Component Keepins and Keep-outs. Keep-outs describe an area in which placement and certain autorouter rules are to be applied or not applied. It is not focused on placement but does support rudimentary place/no-place distinctions. These keep-outs are considered in the section on the "Keep-out" layers, *Sec. 5.6(69)*. However, SPEECTRA also defines a place boundary that defines an area of the PCB that permits component placement. This is equivalent to the component keep-ins described by other formats. The placement keep-in can also specify the component types (pin, SMD or area) where pin components are TH components, SMD is SMT devices, and areas are general keepous of placement keep-outs. A session structure descriptor contains a place boundary descriptor and a number of keep-out descriptors for the session. Keep-out areas can be modified with a window descriptor. Window descriptors subtract from an area and are roughly equivalent to cutouts. All keep-out areas can have cutouts. Keep-outs can specify the layers, or layer types, or layer aliases, to which they apply. It appears that all keep-outs are board-level areas: that is, they are fixed to the reference coordinates of the board and are not movable. All of these are handled by the "Keep-out" layer (*Sec. 5.6(69)*).

Rooms. SPEECTRA DSN also has a concept of rooms which is the more specific placement keep-in concept handled in the "Courtyard" layer. Rooms can have rules that restrict the minimum and maximum height, power disipation restrcitions, or power net. Specific components or component groupings can be included in a room or excluded from a room. Inclusion and exclusion from a room can be hard or soft: where hard means that the entire component or cluster must be included or excluded from the room; and where soft means that a portion of the component or cluster can remain included or excluded from the room. Rooms can describe placement restrictions on large, small, discrete or capacitor components, pin type (TH or SMD), permitted orientation, placement side, and components on the opposite side.

Spacing. SPEECTRA DSN does not model courtyards in the IPC way. It defines an outline that is the "true outline of the component" (whatever that is). The accuracy of the image outline specification in the design file is important to assure correct inter-component spacing. The outline must be defined from the top view of the design. If the image outline does not encompass all pins on the image, the outline expands to cover the pins. If an image outline is not defined in the design file, the tool generates a bounding box that includes all pins or pads on the component.

Courtyard Process The courtyard process is primarily concerned with board assembly, rework, test, and repair. First and foremost, the courtyard provides a component-specific keep-out zone that cannot be occupied by other components. While many PCB CAD systems treat the courtyard as two-dimensional (laid against the surface of the board), this courtyard is indeed 3-dimensional, particularly when considering mechanical components. For example, small SMT components can fit under a heat sink. The component is best modelled from an ECAD (Electrical CAD) and MCAD (Mechanical CAD) viewpoint as simple extruded closed contours. That is, as a z-axis projection closed contour that is extruded from a starting height from the board surface and ending at the maximum height. A component can be a composite of a number of such extrusions: again, for example, a heat sink that has two retaining pins can be modelled as one square extrusion for the heat sink, and two circular extrusions (cylinders) for the retaining pins. It is not useful to go overboard in attempting a full 3D model of the component, as simply describing an approximation of the envelope should suffice the task of placement. These extrusions that make up the component are modelled at a zero-reference mounting height and can be adjusted for different mounting heights.

To ease the library effort (e. g. combinations of heat sinks with various BGA packages, it should be possible to lock components together (from a physical or placement perspective). In that way they may be moved as a unit.

The assembly process depends on the placement and orientation of components. In its simplest form, a table containing the x and y coordinates of the components and fiducial marks are used as input to the assembly robot. This simple 2-dimensional view causes problems when considering heat sinks and other devices mounted beneath and over other components. A mounting height is required as well, and either the placement must be properly ordered so that components underneath other are populated first, or a way of describing *what* is underneath *what else*.[40]

5.16.2 Courtyard Materials

There are no courtyard materials, *per se*.[41] The attributes have more to do with purpose and illustration: the color and whether the courtyard is simply for placement, or also for rework or in-circuit testing.

5.16.3 Courtyard DFX Rules

The component courtyards need to be examined by the DRC subsystem to determine whether a component can be placed in proximity to another component. Using a closed contour extrusion model makes this a simple matter of checking for overlapping contours at various critical heights. An overlap constitutes a DFX rule violation. DFX errors should be able to be checked statically (i. e. a DRC check run on the board) or dynamically (while editing component placements). When editing component placements (i. e. moving, dragging, flipping components) it eases the PCB designer's task to resist placing components in such a way that they would overlap others. One of the more difficult aspects of manual component placement on high-speed and high-density digital logic boards is packing-in MLCC decoupling capacitors under a BGA via field, or lining up arrays of termination chip resistors. Almost always the PCB designer seeks to compress the placement into the smallest possible area. Component courtyards and dynamic DRC can ease this task by resisting further movement of the component when the courtyard butts up against another. This feature is similar in nature to the ability to lay tracks together as closely as possible.

Component courtyards do not only interact with each other. Component placement must also be restricted in certain areas due to fixtures, or due to EMI or shielding. The Keep-out Layers

40. In fact, the SMEMA Standard Recipe File Format [IPC-2531] took this approach to identify critical sequences of component placement.

41. Perhaps the material could be considered to be air, along with the mechanical and electrical properties of the atmosphere, but this can be assumed.

provide for component keep-ins and keep-outs at the board level. Courtyards can be viewed as component keep-outs that move with the component. Board-level keep-outs can be viewed as as keep-outs that are fixed to the board. Overlaps between component courtyards and board-level component keep-outs should be statically and dynamically checks for DFX rule violations as well. During editing, board-level component keep-outs should resist component placement in the same way as other component courtyards.

Components (and in particular, connectors) often extend beyond the edges of the board. Applying DFX rules for such overlapping components has two aspects: 1), checking that the land does not violate external copper and technical layer rules for placement; and, 2), checking that the courtyard is within the allowed component placement envelope.

Rework: Additional rules to permit manual soldering or rework is listed in *Tab. 7(101).*[42] The distance is normally measured from pad-edge to pad-edge; however, in cases where the component's body exceeds the pad dimensions (e.g. tantalum capacitors), the distance is measured from the body edge to the nearest feature (pad or component body). For pcbnew we use the courtyard to identify the extents of the component.

Additional secondary (solder, bottom) side restrictions include the following:

1. Component height restrictions on the secondary side of the board should be less than 120mil; the preferred height is less than 90mil.
2. Components taller than 100mil (e.g. tantalum capacitors or inductors) require 100mil clearance (land-to-land) in all directions to prevent solder defects (skips and opens).
3. Minimum spacing of 0.025" between SMT components (land-to-land) is required in both directions to prevent bridging.
4. Minimum spacing of 0.025" land-to-annular ring required between all SMT components and DIP pins requiring selective wave solder fixture. Press fit connectors are an exception and do not require this clearance on the secondary side.
5. On the solder (secondary) side of the PCB, the long edge of the PCB must have a clearance of 0.300" on both sides. No SMT components should be placed there. Through-hole components that are to be hand-soldered or waved are okay.
6. Never join SMT lands to form a common land. *Note: Pad to pad spacing applies to same-net pads.*

All polarized SMT and TH components should be placed in the same orientation and only one axis. This facilitates easier visual inspection. Where this is not practical, try to group components in the same orientation.

BGA and QFP/QFN: Some additional considerations for BGA and QFP/QFN are as follows:

1. BGA and larger QFP (Quad Flat Package) devices ($\geq$ 100 leads) should not be placed in the center of the PCB. The maximum board warpage tends to be in the center of the PCB. The results can be open solder connections. For a standard 0.062" PCB, this becomes a concern when the surface area exceeds 25 in^2 (e. g. 5"$\times$5"). *This requirement is difficult to meet. Perhaps a DFM calculation of the planarity of the board given board bow requirements should be made.*
2. If BGA are on both sides of the board, it is not recommended that the BGA are positioned on top of each other. This method makes rework of a BGA extremely difficult. Also, this method makes x-ray inspection of the solder balls difficult. *This requirement can most simply be met by not allowing, or warning about, placement of BGA on the secondary side.*
3. BGA should be placed on the top side of the PCB. This eliminates the possibility of open solder connections due to the weight of the part during second-pass reflow. *This requirement can simply be met by not placing BGA on the secondary side, or providing a DFM violation for BGA placement on the secondary side.*
4. CBGA (Ceramic Ball Grid Array) and QFN (Quad Flat Package No-Lead) should not be on the same PCB. The reason is that CBGA and QFN require different amounts of solder. (CBGA components require more solder paste than QFN components.) Because of this, a step-stencil or a separate solder stencil will need to be used for CGBA to place the required amount of solder paste. Should CBGA and QFN be placed on the same PCB, a minimum of 0.250" clearance around the CGBA is required to accommodate the solder stencil. *It is difficult to accommodate this requirement without simply adding 250mil to the CGBA courtyard.*
5. The traces that connect vias to BGA pads need to be masked off as a minimum to prevent solder from scavenging into the vias. The BGA pads are non-solder masked defined (i. e. solder mask opening is larger than the metal pad). *Several ways of meeting this requirement involve via capping, plugging, and encroachment.*

5.16.4 Courtyard Design

Module items that exist on a courtyard layer have differing function as follows:

Zones: Zones defined with the new module-zone define component courtyards in the traditional sense. Component courtyards are modelled as a simple extrusion that has an extrusion height (maximum component height) and a z-axis offset (mounting height). The extrusion models the tolerance boundaries of the physical component, or a portion of the physical component. The courtyard for a mechanical component is treated differently from an electrical component. Electrical components have a normal mounting height (also called a standoff) that is fixed with respect to the board layer surface. Mechanical components do not: they may have a normal mounting height, but for items such as heat-sinks, they might be super-imposed (mounted on top of) another component. Typically a single extrusion is sufficient; however, some exchange formats permit multiple extrusions.[43]

Draw segments: Draw segments serve only a documentation purpose and are treated as general purpose drawing elements.

Traces and vias (conductive board features) cannot exist on a component layer, just as they cannot exist on other technical layers. Other board items that exist on a courtyard layer also have differing function as follows:

Zones: Zones defined on courtyard layers represent board-level placement keep-ins and keep-outs (also commonly referred to as "rooms"). The zone defines a single z-axis extruded area. The extrusion extends from the mounting surface at the floor to a specified or unlimited ceiling. Components that do not fit within the room are excluded. Components

42. I think that these requirements are actually wave soldering requirements and not rework requirements.

43. Such as IDF [Kehmeier and Makowski, 1998].

Table 7: Clearance between Components

From/To	Chip	SOIC	Tantalum	SOT23	DIP	PLCC	QFN/QFP	CSP	BGA
Chip	40/1.0	—	—	—	—	—	—	—	—
SOIC	40/1.0	50/1.3	—	—	—	—	—	—	—
Tantalum	50/1.3	55/1.4	50/1.3	—	—	—	—	—	—
SOT23	50/1.3	50/1.3	75/1.9	35/0.9	—	—	—	—	—
DIP	60/1.5	60/1.5	60/1.5	60/1.5	100/2.5	—	—	—	—
PLCC	50/1.3	100/2.5	100/2.5	100/2.5	60/1.5	100/2.5	—	—	—
QFN/QFP	100/2.5	100/2.5	100/2.5	100/2.5	100/2.5	100/2.5	100/2.5	—	—
CSP	125/3.2	125/3.2	100/2.5	125/3.2	125/3.2	125/3.2	250/6.4	100/2.5	—
BGA	125/3.2	125/3.2	125/3.2	125/3.2	125/3.2	125/3.2	250/6.4	250/6.4	250/6.4

Notes:

1. Units are in mil/mm.
2. Sockets (for PLCC and DIP) and connectors should be away from BGA and CSP (Chip Scale Package) components to prevent solder joint cracking due to possible stress exerted during second loading/removal of add-on cards or IC (Integrated Circuit) components.
3. DIP separation is for component side only. For secondary side, 125mil/3.2mm clearance for all SMT components from DIP pins requiring selective wave solder fixture. Press fit connectors are an exception and do not require this clearance on the secondary side.
4. For chip to chip separations, when absolutely necessary, 0402 components can be 20mil/0.5mm apart; 0603 components can be 25mil/0.6mm apart. These number apply to special suppliers only and requires special setup. Supplier should be notified before the board is built.

---

can be excluded or included in the room by type, power dissipation, power net, reference designator, or other maximal or minimal property.

Draw segments: Draw segments on courtyard layers are for the purpose of documentation and can be used for generating general drawings.

Text: Text on courtyard layers is for the purpose of documentation or annotation and can be used for generating general purpose comments, or special assembly instructions,

Moirés: Moirés on courtyard layers serve no purpose except maybe to provide cut-marks for illustrations.

Dimensions: Dimensions on courtyard layers can be used for the purpose of documentation or annotation and can be used for generating general purpose drawings for assembly layers.

5.16.5 Courtyard Imaging

Display: Courtyards can be drawn (2D) and plotted as their z-axis board projection. In this case they are simply a collection of closed contours. The courtyards can be displayed as such on the PCB and module editor canvases and a control should be provided for their visibility. Here the simple Courtyard-Layer visibility controls should suffice.

Courtyards and rooms can be displayed during editing. This serves to identify components that are too close to another component or the boundaries of a room. Courtyards can be either displayed like silk (providing an thick-line outline) or like clearances (as thin clearance lines). Rooms can be displayed as are other zones, in outline or filled forms. The later is likely more amenable to manual placement when board editing; the former, to module editing. Courtyard outlines should be able to be made visible to the PCB designer. When moving or dragging a component with DFX rules enabled, the component should resist movement along the courtyard lines of adjacent components or the edges of rooms.

Plotting: When plotting specific courtyards on assembly drawings, the options for plotting the courtyard should be similar to (or the same as) those provided for plotting zones: that is, filled, outline only, full cross-hatch, perimeter cross-hatch, etc.

For assembly drawings, it should be optional whether the courtyards are draw on the assembly drawing. It is likely that general component courtyards should not be plotted on the assembly drawings because they will simply clutter the drawing. For example, the courtyards for 500 MLCC decoupling capacitors drawn on an assembly drawing: a), serves no purpose to the assembler; and b), clutters the drawing. On the other hand, a BGA courtyard whose intention is to provide sufficient clearance for rework might be more important to have drawn on the assembly drawing.

The manufacturing category of the layer is "Assembly". The information contained on courtyard layers is not normally plotted, but may be generated as thin lines or cross-hatched areas on assembly drawings. Courtyard layers could be printed separately for the purpose of documenting a design or providing special assembly instructions. Courtyard layers that correspond to external assembly layers (top and bottom) are associated with assembly. For more information on mapping layers to drawings, see *Sec. 5.24(116)*.

3D Display: One of pcbnew's distinguishing features is it 3D display. Component courtyards could be displayed in 2-dimensions at board level in a manner similar to silkscreen or component outlines, but this would not properly illustrate the mounting and extrusion height components of the courtyard. Displaying the extruded component courtyard on 3D display servers several purposes:

- The 3D display can help the PCB designer visualize both the contours and extrusions that have been exported to or imported from and MCAD system using IDF [Kehmeier and Makowski, 1998].
- The 3D display, particularly in conjunction with the full 3D model of the component) can provide for a visual check that the contours and extrusions indeed represent a proper envelope for the component and thus avoid errors in specification for library or board components.

3D fogging the contour extrusion on the 3D display (and providing an option to either display the fog or not) is the simplest way of accomplishing this.

5.16.6 Courtyard Post-Processing

As stated above, courtyards are not included on fabrication plots (they do not create a photoimagable plot involved in a process) but can be used for display purposes (assembly drawings). Nevertheless, the process information associated with component courtyards can be exported from the PCB CAD system for the purpose of mechanical design, auto-placement, or representation in CAM formats that serve both fabrication and assembly purposes. Component courtyards provide data that is included in the following CAD/CAM outputs:

There is a complication when it comes to pcbnew: historically, pcbnew has treated each component on the board as independent. This is inefficient from the point of board file format as well as internal data structures. Again, envision 500 to 1000 MLCC decoupling capacitors: each with their own entire copy of the library component. All of the reference data associated with each component is redundant, and 1000 unnecessary copies exist. Formats such as IDF provide description of a single library component for a given MLCC package and then keeps "placement" information separate. To create this structure would require that each component (module in pcbnew) be compared to each and every other until a match is found to collapse its internal structures into the form required for export. This deficiency of pcbnew is problematic. Therefore, the internal organization of components and the way that they are saved in the board file have been modified to separate the package information from the placement information.

Intermediate Data Format (IDF) Intermediate Data Format (IDF) [Kehmeier and Makowski, 1998] models components as a part prototype and instance.

Part: A part describes the geometry of a package as a collection of extruded closed contours with reference to the package local origin. Everything is modeled as a part in IDF: panels, boards, electrical components, mechanical components, fixtures.

Part Instance: A part instance represents the placement of a part. Components are placed with reference to a board position.

Therefore, IDF models component courtyards as non-overlapping collections of extruded closed contours. This is the format after which component courtyards are modelled. This makes generation of IDF component data particularly easy.

Unfortunately, components in the other CAM formats are modelled using a single extrusion. This cause complications for export of multiple extrusion information to formats supporting only a single extrusion such as GenCAM, GenX and 258X.

GenCAD GenCAD, the precursor to GenCAM, separates modules into component and device.

Component: A component is the data associated with a particular instance (placement) of a device. There is a free form attribute field that could be used to contain standoff information.

Device: A device, in GenCAD, describes primarily the pin shapes placements and function and associated artwork. There is a free-form attribute filed that could be used for height information.

Therefore, in general GenCAD can model courtyards as closed contours and single extrusions, however, it is not a formal part of the specification.

GenCAM (IPC-2511A) GenCAM [GenCAM] breaks the description of components into the following parts:

Packages: Under GenCAM, packages are similar to *kicad*'s modules, but also contain body artwork, package height, package standoff and package centroid. The package type is one of the defined and registered GenCAM package types. The package height is the maximum height of the package as measured from the finished surface of the board or panel on which it is mounted. The package standoff is the clearance height to the bottom of the package body as measured from the finished surface of the board or panel on which it is mounted. The package centroid is the centroid of the package as used by manufacturing equipment during assembly, inspection, and test. The package centroid is measured from the origin of the package. The body artwork consists of the physical shapes associated with the package. The body shape is that which would be drawn on a display screen by a CAD or CAM tool or on a GenCAM drawing. Another shape might provide a color coded image to be displayed by a repair station. The layer associated with the shape can be used to classify the function of a body artwork (for example, courtyard, component outline, silkscreen, etc.)

Therefore, GenCAM only provides for a simple closed contour artwork and extrusion defined by the height and standoff from the board's surface. This can be modelled using the model above, using a single extrusion.

Devices: Under GenCAM, devices are actual component selections that are associated with a package. The package describes just the packaging of the component, where as the device describes the part number, manufacturer and other BOM (Bill Of Materials) information associated with the populated component. The device associates part numbers from an enterprise (manufacturer or distributor) with a set of specific pin definitions and with a package definition. The geometry attributes of the device, including the location of the origin, are those of the package.

Mechanicals: Mechanicals define mechanical parts definition that associates part numbers from an enterprise (a manufacturer or a distributor) with non-electrical features of a fixture, board, panel, or assembly: that is, mechanical components such as heat sinks, bolts, or alignment pins. The geometry attributes are defined relative to a local origin for the part definition.

Components: Component are used for boards and panels. A component defines a component position relative to the origin of a board or panel. A component can also reference a device, pattern, or mechanical to define the makeup of the component at that location. Additional artwork and manufacturing instructions can be associated with the device or mechanical associated with the component. Under GenCAM, the component can also define keep-out information.

Therefore, when post-processing courtyards for GenCAM components, packages that describe the physical geometry must be generated and factored. GenCAM has the following deficiencies in its model:

- GenCAM chose to associate the standoff of the package with the package rather than the device (placement). This model cannot accommodate library packages that are placed at different heights, such as is the natural case for heatsinks.
- GenCAM can only represent a package with a single extrusion: this means that complex shapes (such as heat sinks with retainer pins) must use some other mechanism to ensure mechanical clearances.

GenCAM-XML (IPC-2511B) GenCAM-XML (GenX) [GenX] provides the same model as GenCAM [GenCAM], the

file representation is simply using GenX's XML format rather than GenCAM's EBNF (Extended Baccus-Naur Form) format. Therefore, the same post-processing considerations apply to GenX as were applied to GenCAM. That is, the format suffers from the same limitations described for GenCAM, above.

258X (IPC-2581) 258X (IPC-2581) [258X] breaks the description of components into *packages* and *components* in a manner similar to GenCAM, above. Packages are modelled using an artwork outline extruded in a single extrusion; however, 258X—intelligently—places the standoff in the component (placement) rather than in the package definition as does GenCAM. Nevertheless, it still properly separates placement from package. It also separates the various artworks into "Outline," "LandPattern," "SilkScreen," and "AssemblyDrawing." The "Outline" artwork is the closed contour (a close contour of line and arc segments) for extrusion, where "AssemblyDrawing" is the component outline for display on assembly drawings. Note that 258X can only define a single extrusion.

SPEECTRA DSN The Cadence SPEECTRA DSN [SPEECTRA] also also splits packages and their placement. Associated with placement is a minimum and maximum height that can be used to define a single extrusion. The format also defines an outline that can be modelled with a simple closed contour for extrusion. Therefore, the SPEECTRA DSN file suffers from the difficulties associated with single extrusions (inability to represent complex forms such as heat sinks with retainer pins). By placing both the standoff and height of the package in the placement specification causes some additional difficulties. Package height from standoff to package top is really a characteristic of the package not the placement. Therefore, initial placements must always be generated when writing SPEECTRA DSN files.

5.17 Component Layers

The component layers are assembly layers intended for assembly drawings, front and back. Historically pcbnew did not support component layers nor specialized assembly drawing layers.

> The assembly drawing outline should represent the maximum outline of the component body. Unlike the silkscreen outline which has to be created to avoid solder pads (a fake component outline); the assembly outline only gets placed on an assembly drawing that goes to the assembly shop. There is no need to fake this outline.
>
> Assembly outlines can be created with complex drawings to illustrate the actual physical component features or with simple rectangles. It makes no difference to the assembly shop that has to interpret the assembly drawing. Drawing complex shapes for the assembly outlines shows off the PCB designer's artistic creativity, but once the PCB board goes into production, it makes no difference because the assembly drawing is never used again. The CAD library should have 1:1 scale component outlines in the assembly drawing.
>
> The assembly polarity marking is sometimes totally different than the silkscreen polarity markings because the silkscreen must avoid touching the solder pad. Unlike the silkscreen, the assembly drawings can illustrate robust polarity markings to ensure that the component is inserted in the correct rotation.
>
> A good feature to add to the CAD library land pattern is a cross-hair placed on a documentation layer so that it is visible to the PCB designer. This is a part placement aid. When the PCB designer selects a component to move it, most CAD tools will take the part origin jump to wherever the cursor is located. This is very annoying when you are fine-tuning a part placement and your goal is just to tweak the placement. If you can visually see the cross-hair on the origin, you can select the part at the cross-hair location and the part will not jump. [Hausherr, 2006]

5.17.1 Component Process

The purpose of assembly drawings is as follows:

- Historically, assembly drawings were the only legend that was used in the assembly of components onto a PCB board. Application of silkscreen legends onto the board's surface came later. [Hausherr, 2006] The assembly drawings still serve the purpose of providing a graphical representation of the components on the boards surface.
- Assembly drawings provide a more detailed roadmap to the PWA (Printed Wiring Assembly). Assembly drawings act as a guide to assembling the board.
- Assembly drawings when exported to a usable graphics format, may be useful for documentation for repair stations or user guides for the product.

The component process is primarily concerned with display. First and foremost, it provides the component outline for display and assembly drawings. Whereas the silkscreen outline of a package must be adjusted for openings in the solder mask, component outlines do not.

Free-form drawings and text placed on a component layer will be printed on the corresponding assembly drawing, so any special assembly instructions or drawings can be made on a component layer.

Component outlines are typically contained within the silkscreen outline. Text associated with the outline (reference, value, other fields) must be legible and suitably located on the drawing. As a result, there is an opportunity to place the entire silkscreen legend also onto assembly drawings. This is detailed in the section on drawings, *Sec. 5.24(116)*.

5.17.2 Component Materials

There are no component outline or assembly drawing materials, *per se*.[44] The attributes have more to do with illustration: i. e. the colour, the minimum line thickness and default pen.

5.17.3 Component DFX Rules

There are no DFX rules directly associated with the component outlines. Minimum and default line thickness should be enforced.

5.17.4 Component Design

Editing: Component outlines should only consist of about 1 or 2 different thicknesses of lines. The default IPC library line thickness for an outline is 20 mil (0.05mm). Component outlines can be created by script when creating library components, or can be created by draw segments on one of the Component Layers using the pcbnew module editor. The component outline does not need to be a closed contour, as pcbnew does not use this contour for any purpose but display and mapping to assembly drawings. One can also draw and place text and zones on the component layer at the board level using the pcbnew board editor. These drawings and text will also be mapped to assembly layers along with the component outlines themselves.

Component outlines should be as representative of the z-axis projection of the component as possible. When a 3D wire-frame model has already been generated for a component, a script or program should be able to generate a 2D z-axis projected line drawing suitable for use as a component outline.

The component outline can be used for one thing: it can be used for automatically generating a placement courtyard from the component outline and lands. When used in this fashion, the component outline should always represent the maximum outline of the component considering the component's size tolerance. In this way the component outline can be closed bu joining any dangling lines to their closest vertexes, and then the component can be rotated by its maximum placement angular tolerance and then translated in each direction by its maximum xy-plane placement tolerance to obtain a maximal outline. The pads can be expanded by the xy-plane registration tolerance of the copper layer, and then a contour chosen that includes both the expanded pads as well as the rotated and translated component outline. This enclosing contour is the placement courtyard.

5.17.5 Component Imaging

Display: Component outlines can be drawn (2D) and plotted. Component outlines can be displayed as such on the PCB and module editor canvases and a control should be provided for their visibility. Here, the simple Component-Layer visibility controls should suffice. The only thing that needs to be handled is minimum line width for display.

Plotting: When plotting component outlines on assembly drawings, the options for plotting the component outline should be similar to (or the same as) those provided for displaying other draw segments and text.

Again, the only thing that might be dicey to handle is the line width of line segments. When generating component outlines as a part of a library, it is not clear at what magnification or scaling factor the component will be displayed. If the component dimensions are scaled without scaling the line width, the image could be corrupted. On the other hand, if line widths are scaled too narrow, the line can essentially disappear.

44. Perhaps the material could be considered to be ink, along with a minimum line thickness.

3D Display: There are no special requirements for 3D display. The layers can be displayed in the same fashion as comments and drawings layers: i. e, as zero-thickness ink just above the board exterior level. There is, in fact, an opportunity to plot back the other way: to take the 3D model and create a 2D projection along the z-axis. In this way, the 3D model could be used to automatically create the component outline. This is a way of converting library modules that have a 3D Model but no component outline yet.

5.17.6 Component Post-Processing

Component outlines are not normally included on fabrication plots. It is possible that a few component outlines of primary mechanical components could be necessary on a fabrication mechanical plot (i. e., a plot showing the overall dimensions of the board, any primary features, reference holes, brackets, etc.); but, normally these are displayed only on the assembly drawings. Nevertheless, the ability to copy component outlines to any fabrication or assembly drawing is provided per *Sec. 5.24(116)*.

5.18 Coating Layers

The coating layers are assembly layers.

5.18.1 Coating Process

5.18.2 Coating Materials

Conformal coating.

5.18.3 Coating DFX Rules

5.18.4 Coating Design

5.18.5 Coating Imaging

5.18.6 Coating Post-Processing

5.19 Edges Layers

The edges layers are board fabrication layers that describe edges of the board (other than drill holes), whether they be exterior or interior edges. Historically, pcbnew only supported one Edge layer. This is still true, however, various forms of route paths are now supported on this layer (route, score, bevel, perforate, etc.)

5.19.1 Edges Process

Board edges are typically fabricated using standard G code driven CNC mechanical milling and drilling equipment. Often the CNC machines are specially tasked to either perform milling or drilling but not both. Board edges (with the possible exception of break-away tabs) are performed completely on an CNC milling machine. Edges on the board can be created by one or more processes described in the paragraphs that follow.

Rout. Routing is a process of depanelization where the board edges are routed from the production panel. A preferred rout tool bit diameter is used throughout, unless the board edges include an inside radius that is not attainable with the preferred route tool When edges are not plated, routing is typically performed for depanelization in the final stages of board fabrication: the primary mask and silk screen typically have already been applied to the board, and routing is the last stage before inspection and testing. When edges must be plated, it is necessary to route the board before the final plating operation (the fabrication of PTH).

Perforations or Break-Away Tabs. Perforations are a form of break-away tab that is created by drilling holes through the production panel where rout paths almost meet. The purpose of the break-away tab is typically to provide assembly rails for assembly panel (arrays or matrices) that can be broken away from the board after assembly. Perforations or break-away tabs use a specialized pattern of drilled holes to perforate the tab so that when broken apart will not produce a projection beyond the otherwise straight edge of the rout path.

An example of a perforation or break-away tab, "as-built," is illustrated in *Fig. 67(109)*.

Scoring (V-Groove). Scoring is a process of creating a v-groove on either side of a production panel. It has the advantage over routing the a larger area of the production panel is available for use (boards can butt together at the score line); however, because only a single panel may be scored for each spindle, the process is more expensive than routing (where multiple panels can be routed by a single spindle in one pass).

An example of a score or v-groove is illustrated in *Fig. 68(109)*.

Bevel. Bevelling is the process of forming a bevel on the edge of a board for the purpose of forming an entry angle for edge connectors. Typically the angle and the breadth of the bevel is specified by the form factor for the add-in card. Bevelling is a separate step from depanelization.

An example of a bevel is illustrated in *Fig. 69(109)*.

Slots (Plated or Non-plated). Slots are formed in a board by routing an open contour where the slots are defined by the edges of the route path. Typically slots are routed without tool compensation of the route path. Slots can be plated or non-plated. Typically there are two choices for the fabrication of slots. Slots can be drilled or routed. Drilled slots use a specialized canned cycle in the CNC drilling machine to drill a sequence of holes that eventually form the slot. The tolerance of drilled slots is poorer than that of routed slots, due to the imperfect overlapping of the drilled holes that form the slot. Drilled slots that are plated are easier to fabricate. This is because drilled slots can be formed coincident with first-drilling procedures using the first-drilling CNC drill machine. Guaranteed non-plated slots can also be drilled. In this case they are either drilled as part of a second-drilling pass for NTPH fabrication, or, when of sufficiently narrow size, they can be first-drilled and then plugged before plating.

Cutouts (Plated or Non-plated). Cutouts are formed by routing an interior closed contour on the board, where the material interior to the contour is removed. Cutouts are normally performed with tool compensation and must satisfy minimum interior radius constraints. Cutouts can be plated or non-plated. When plated, routing is performed in conjunction with first-drilling for plate-though holes, before plating is applied to the completed laminate layers.

Wells (Plated or Non-plated). Wells are formed be milling or by forming a cutout through a sub-laminate sequence. The former is related to through-drilling of blind vias; the later, related to depth controlled drilling. Wells are more expensive than cutouts when formed using the depth-controlled milling approach because only one panel per spindle can be performed. When formed by cutting out a sub-laminate sequence, multiple panels per spindle can be routed. Wells can be plated or non-plated. Non-plated wells can only be formed using the milling approach after the final laminate is complete and after plate (that is, in conjunction with second-drilling).

Edge Painting or Plating. Edge painting or edge plating is a process of applying a coating (paint or plated copper) to the circumference of a board. It is applicable to routing and perforations, but not applicable to scoring. The purpose of edge plating is typically to shield the electromagnetic fields of the interior of the board for EMC. It is typically only used on aeronautical or military designs.

Chamfering. Chamfering of cutouts, wells, or slots, is related to bevelling of the board edge. The same process is applied to mill the edges of a cutout, well or slot. Chamfering is as expensive as scoring or bevelling in that only one side of the board per spindle can be chamfered at a time.

Edge Milling. Milling is a process whereby a depth of material is removed from the surface of the board. Edge milling is normally performed when a board is too thick to fit in its board guides: the edge of the board is milled down to the thickness of the board guides by removing material from the top, bottom, or both surfaces of the board near the edges. Edge milling can be performed either with a special purpose machine, or with a general purpose CNC mill.

Painting or Plating.

Filling cutouts, slots or wells.

Plating over filled cutouts, slots or wells.

5.19.2 Edges Materials

Can be painted or plated.

5.19.3 Edges DFX Rules

Routing Rules:

- Standard router bit is 0.0984", 0.0945", 0.0394", 0.100" and 0.0787".
- Radius shall be at least half of the router bit diameter plus 0.005".
- Minimum internal radius shall be 0.020".
- Unify slot width.
- Prevent different slot widths on same piece to increase productivity.

- Prevent different radius of corners, because of the need of different router bits to meet nominal radius.
- Other router bits available are 0.062".
- Router bits smaller than 0.094" will cost extra.
- Use the most generous tolerance that the product will allow to minimize board price.
- Use only one cutter size.
- The preferred cutter size for routing is 0.125", 0.0938" diameter. 0.062" is available.
- Avoid use of smaller cutters.
- Avoid routing through metal features. The result requires excessive hand de-burring and can cause quality defects.
- Conventional pin routing requires a minimum of two pins per board. Pin sizes to be greater than 0.062 inch and less than 0.251 inch.

Tab Rules:

- The preference is to set up parts for tab routing as a function of the tooling operation. To avoid unncecssary modifications to the mechanical drawing, it is preferred that customer provide only a note stating that the part needs to be shipped in panel formt, delta notes indicating where the tabs cannot be located. If the location of the parts in the panel is critical, the dimensions of the datums of the part to the component assembly locating holes must be provided.
- Locate tabs 0.350" minimum from any board corners.
- Place tabs 0.350" minimum from any board corners.
- Place tabs 0.350" minimum from datum holes, or directly on centre.
- A .125" cutter will be utilized, unless design requires otherwise.
- All cut paths that are not between boards will be 0.125" wide; preferred spacing between boards i 0.250", 0.150" minimum.
- Place tabs 3.00±0.50" inch appart from each other.
- Keep tabs in a straight line with X-Y axis if possible.
- Where there are component holes or traces close to the board edge, try to avoid tabbing in these areas to prevent the traces or hole walls from fracturing.
- Tab width is 0.125±0.010".
- Tab location dimension is ±0.025".
- Dimension tabs to the centre of the tab on a 0.025" grid.
- Place tabs 0.250±" minimum away from any radius on the outside board rdge.

Bevel Rules:

- Minimum bevelled edge-to-opposite-edge of the PCB dimesion: 2.5".
- Standard bevel depth: 0.015"/0.075".
- Bevel angles available: 20, 30, 45 and 60 ±5 degrees.
- Maximum recessed bevel from board edge: 1.25".
- Minimum space between gold edge connector and breakaway tab: 0.250".
- Do not partially edge chamfer the finger side of the board when the fingers are flush with the edge of the board.
- Edge bevelling may be performed on the outer edge of the board, a recessed segment of the board, or internal to the board.
- Inner layer plane layers must be recessed to avoid exposing the plane where the boards are bevelled.
- The following angles and depths may be achieved given sufficient board thickness:

  20±2° by 70 mil depth ±15 mil.
  30±2° by 50 mil depth ±10 mil.
  45±2° by 40 mil depth ±05 mil.

Edge Milling Rules:

- The milled edge is usually a "step" at the edge of the board.
- The depth of the step is variable from 10 mil remove to 32 mil remaining.
- The width of the step is variable from 20 mil to 375 mil. Milling requirements should be limited to simple cuts (i.e., two straight edges and simply corners).
- The path of the mill is limited to 90 degree truns and internal radii are controlled by cutter diameter (minimum 125 mil and common standard sizes).
- Geometries other than a stop are possible but need to be evaluated on an individual basis as processing time is prohibited.
- bf Double sided milling is strongly discouranged as edge thickness accuracy is reduced.
- The finished thickness of the milled edge can be held to ±8 mil for a single sided milled edge.
- For a double sided milled edge the finished thickness can be held to ±10 mil.
- The width of the step can be held to ±10 mil.
- Internal tooling pins are required.
- Tooling holes should be internal to the finished board and should be located as close as possible (but not actually in) the portion of the board to be milled.
- The finish produced by the mill process is similar to that produced by NC edge routing. No fractured glass fibres are produced.

5.19.4 Edges Design

When specifying edges, the type of edge must be specified. The default type is "board edges" which is the previous pcbnew behaviour.

Specification of route paths: Route paths are specified on the display and plotted on artwork as the edge of the board. No tool compensation is provided in the board edges. Tool compensation will be added when Excellon-2 router or G-Code milling files are generated. To assist pcbnew with determining board edges, route paths should be clockwise for board edges, and counter-clockwise for cutouts, and wells. To support wells, board edge drawings will require a layer-pair that specifies the layers through which milling is performed. The default is to route through the entire board.

Specification of perforations or break-away tabs: The fabrication results of a perforation with adjacent route paths is illustrated in *Fig. 67(109)*. Break-away tabs are specified as a perforation (that is, represented by a dotted line). The tool path specified is a straight line segment that is the length of the tab. Additional parameters specify the pattern of the mouse-bite that forms the break-away tab. Perforations or break-away tabs form a segment of the board edges for determining the extents of the board. The ends of the perforated line should coincide precisely with the end of the route path on either side of the tab.

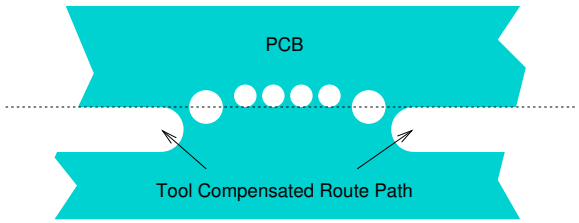

Figure 67: Perforation or Break-Away Tab

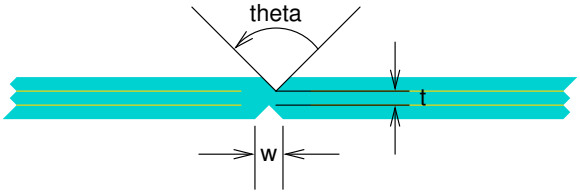

Figure 68: Scoring (V-Groove)

Specification of scoring (v-groove): A score (or v-groove) is illustrated in *Fig. 68(109)*. Specification of scoring is the same as route paths for board edges. The path specified for display and plotted on artwork is the centreline of the scrore mark (or v-groove). Additional specification includes the face-angle of the tool used to create the score mark; the width of the line is the diameter of the tool; and the minimum and maximum web between score marks on either side of the board must be specified.

Specification of bevels: A bevel is illustrated in *Fig. 69(109)*. Bevels that are required, for example, for add-in card connector bevelling, will be specified for display and plotting as the centreline of the bevel. The angle to the surface of the board must be specified and the depth of the tool is also specified. The diameter of the tool must be sufficiently wide that there is no perpendicular wall created at the point where the bevel meets the surface of the board. The thickness of the line will be twice the depth of the bevel, $2d$. Bevels must also be scored or routed. That is, the bevel itself does not form part of the board edges. The extent of the draw segment is the length of the bevel. The width, w, can be calculated from the depth, d, and the thickness of the board across the pads, T, as follows:

$$w = T - 2d \times tan(\alpha) \tag{88}$$

For example, the PCI-Express add-in card edge connector requires an angle, $\alpha = 20° \pm 5°$; the depth of the bevel, d, must be w =1.40mm±0.13mm (0.055”±0.005”).

Specification of chamferred edges: Chamferred edges are treated the same as bevels. The tool path specified is the edge of the slot or cutout that needs to be chamfered. The angle of the chamfer and the asssumed width of the tool are the same as for bevels. Again, chanferred edges do not form part of the board edges and they must be specified seperately.

Specification of edge painting or plating: Edge painting or plating is as simply as applying an attribute to draw segments on the board edges layer.

Specification of slots: Slots are handled as open contours or isolated draw segments.

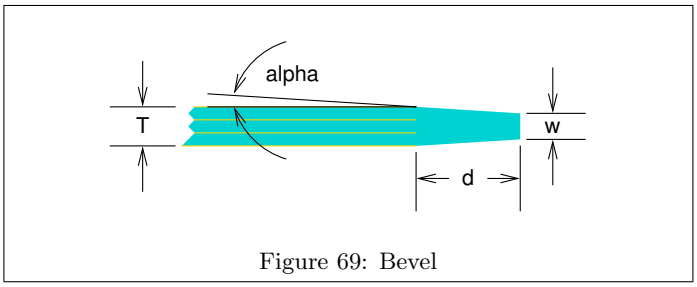

Figure 69: Bevel

Table 8: Typical Bevel Capabilities

Parameter	Class A	Class B	Class C
Angle (α)	$30 \pm 5°$	$20 \pm 3°$	$45 \pm 2°$ $60 \pm 2°$
Depth (d)	20 to 70 mil ±10 mil	15 to 75 mil ±7 mil	10 to 80 mil ±5 mil

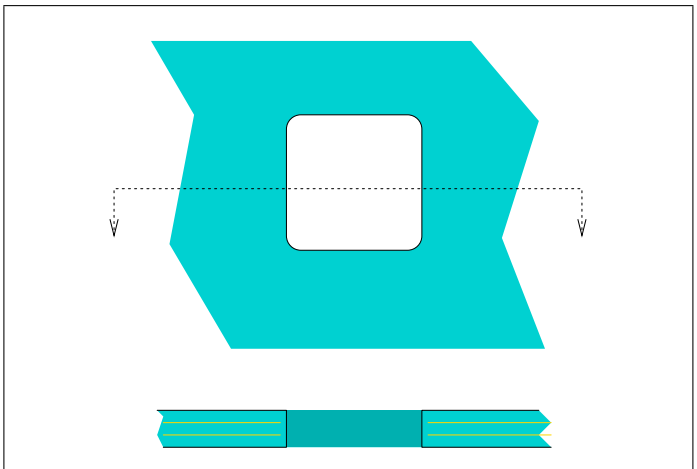

Figure 70: Cutout

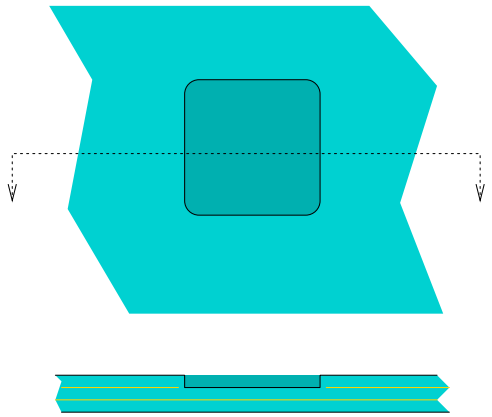

Figure 71: Well

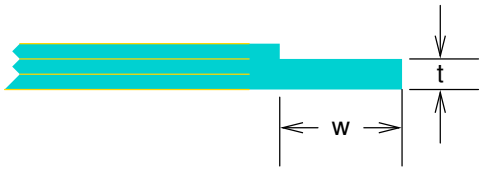

Figure 72: Edge Milling

Specification of cutouts: A cutout is illustrated in *Fig. 70(109)*.

Specification of wells: A well is illustrated in *Fig. 71(110)*.

Specification of edge milling: Edge milling is illustrated in *Fig. 72(110)*. Edge milling is a technique used to reduce the thickness the edges of a PCB so that it will fit in card guides that have a channel thinner than the thickness of the board. Material is typically removed with a rotary planer specifically designed for planing board edges. The finished thickness is that specified by t in the figure. Only a given finished thickness tolerance (typically ± 0.006") can be achieved. Material can be removed from either side of the PCB. Material is removed so that the milled edge has a width, w. All metal must be removed from the area to be milled.

Specification of chamfers: A chamfered cutout is illustrated in *Fig. 73(111)*.

5.19.5 Edges Imaging

Although it is quite easy to specify the rout path for board edges, cutouts, bevels, etc., using Excellon-2 programming, including tool compensation and or drilled vs. routed slots, board fabricators usually say to not specify the route path, but to only provide Gerber plots of the route path, using, say, a 0.020" line.

In addition, most board fabricators prefer that the board edges be plotted on all artwork layers. If nothing else, it provides verification that the artwork layer origin, offset, orientation and scaling are the same. Layers that have a different origin, or different offset, or are mirrored, or are rotated can easily be identified when there are board edges plotted on the artwork. When they are not, these differences are almost impossible to detect.

Strangely enough, the board edges serve a purpose on many artwork plots, but not on others. The purpose by layer is listed in *Tab. 9(110)*.

Table 9: Uses of Board Edge Plots on Artwork Layers

Layer	Use of Board Edges
Resistance	There is no use on these layers. Board edges must be removed from final artwork.
Copper	There is no use on these layers, unless edge plating is to be performed. Board edges must otherwise be removed from final artwork. The edges could be plotted as a clear overlay to remove copper from near board edges.
Plating	There is no use on these layers. Board edges must be removed from final artwork.
Hole Filling	Serves no purpose, must be removed from final artwork.
Keep-out	Layer not plotted.
Via Plug	Serves no purpose, must be removed from final artwork.
Mask	Can be expanded to retract mask from board edges. Should not be removed from final artwork when of proper line width.
Contacts	Serves no purpose, must be removed from final artwork. Could be plotted as a clear overlay to retract contacts from board edges.
Peelable Mask	Cannot be used as a dark overlay, must be removed from final artwork. Could be plotted as a clear overlay to retract mask from board edges.
Finish	Serves no purpose, must be removed from final artwork.
Legend	Serves no purpose, must be removed from final artwork. Should be plotted as a clear overlay to clip legend from board edges.
Heat Sink	Serves no purpose, must be removed from final artwork. Should be plotted as a clear overlay to clip heat sinks from board edges.
Solder Paste	Serves no purpose, must be removed from final artwork.
Adhesive	Serves no purpose, must be removed from final artwork.
Courtyard	Layer not plotted.
Component	Layer not plotted.
Coating	Serves no purpose, must be removed from final artwork.
Edges	This *is* the board edges layer.
Eco	Provides a board shape reference.
Probe	Serves no purpose, must be removed from final artwork. Should be plotted as a clear overlay to clip probe areas from board edges.
Fixture	Serves no purpose, must be removed from final artwork. Should be plotted as a clear overlay to clip fixture areas from board edges.
Comment	Provides a board shape reference.
Drawing	Provides a board shape reference.
Drill Map	Provides a board shape reference.

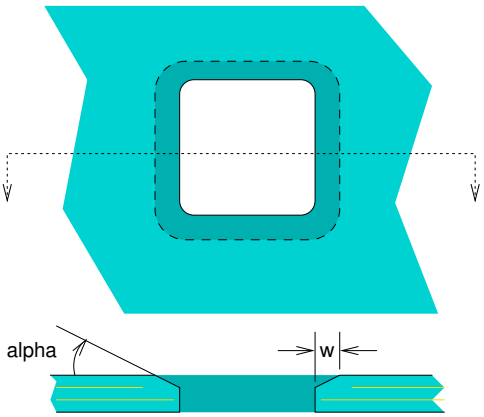

Figure 73: Chamfered Cutout

The origin of plots should always be the lower left corner of the board, matrix or panel. Plotting the board edges verifies this. pcbnew has historically provided the ability to choose between the CAD origin, or the auxiliary axis when plotting Gerber RS-274X files and NC drill files. While it is possible to set the auxiliary axis to the lower left of the board, it is certainly not easy to get it precise with the editor. On the other hand, the lower left of the board, matrix or panel is quite easy to calculate.

(R) 39 (pcbnew) *An enhancement will be made to allow the lower left corner of the board to be specified as an origin for plotting. The board centre will also be available as a choice.*

Not begin able to generate a rout path for board edges seems foolish for the following reasons:

- Gerber RS-274X files are not rich enough to express all the details related to boards that have anything on their edges, or are placed in an assembly panel (or matrix), or that have break-away assembly rails, or any number of other common features such as bevelled edges, chamfering, plating, painting, edge milling, etc. Thus, many callouts are required on the fabrication print and a number of detail sheets are required for each feature.
- Not creating a data file trains the designer to not specify the route path in detail. Therefore considerations such as perforations (break-away tabs or mouse-bites), and bevelling paths, are ignored with the attendant ramifications: everywhere these things can cause problems, problems will occur.
- Having someone layout a rout path at the fabricator is simply another NRE (Non-Recurring Expense), or fixed-cost associated with a prototype or short production run, that can be avoided.
- Fabricators say that the reason that they build their own route paths is because the user's tool paths are non-optimal. This is crap. Route paths are the easiest thing to optimize. Drill file optimization is $O(n!)$ (non-polynomial-hard) where n can be in the thousands. Any child could lay out an optimal rout path for a production panel. I have seen some 4 spindle routers cutting up production panels, and their paths were certainly not optimal.
- If you want to route your own production panels, an Excellon-2 NC route program is a good start. Converting that to G-Code, even by hand, is not that daunting a task. Not generating one means that one has to work from RS-274X, which is not as good a start.

(N) 40 (pcbnew) *An enhancement will be made to allow the user to request that NC route and drill files be generated for board edges, cutouts, bevels, etc.*

5.19.6 Edges Post-Processing

5.20 Eco Layers

5.20.1 Eco Process

5.20.2 Eco Materials

None.

5.20.3 Eco DFX Rules

5.20.4 Eco Design

5.20.5 Eco Imaging

5.20.6 Eco Post-Processing

5.21 Probe Layers

The probe layers are board fabrication test layers.

Design for Testability

- Keep components and test points at least 3.2mm (.125") from board edges (preferred 3.8mm or .150").
- Provide at least two unplated 3.2mm (0.125") diameter tooling holes, preferably in the opposite corners, and leave 3.2mm annular area around them free of components and test points. Consider using a "keying" pattern so boards can't be inserted backwards.
- A double-sided test fixture is more expensive, so try to put all test points on one side of the board, usually the bottom or the side with the least circuit complexity. If top-side of board must be used for probe sites, use top only for non-critical nets. Keep test points for clocks, control pins, programming pins, serial data and boundary scan on bottom.
- Test point sites can be through-hole leads, dedicated pads or small diameter vias, but avoid placing test points on surface mount lands or gold-plated edge fingers. Don't use larger via diameter as test probe sites. Via hole size should be 0.36mm (0.014") or less.
- For 100% testability, provide at least one test pad for each net.
- Provide two pads on nets tied to critical low impedance devices (four-wire Kelvin testing).
- Provide 2-10 probe sites for primary power, and two test points each for isolated power/grounds. For primary ground provide many probe sites, one test point for every twenty grounds, or consider a grid of at least one per square inch.
- Probe sites with 1.0mm (0.040) pad diameters are preferred. 0.9mm (0.035") is acceptable, 0.8mm (0.031") can be used if tooling holes are available for alignment, but smaller diameters will reduce contact repeatability.
- Try to probe sites at least 2.5mm (0.100) apart, center-to-centre. In reality, 0.9mm (0.035") pads spaced 1.8mm (0.070") is considered standard by many. Closer spacing is possible but will require using thinner, less reliable and more expensive probes.
- Test points should be evenly distributed over the surface of the board. High stress in congested areas can cause board to warp.
- keep tall components on the side that is not probed. The platen has to be cut out in places where components are over 6.4mm (0.255) tall on the probed side. For these, keep test points at least 5.0mm (0.200") away.
- For components taller than 2.6mm (0.100"), maintain minimum 2.0mm (0.080") clearance edge-to-edge. For all other components, keep test pads at least 1.0mm (0.040") from component body, edge-to-edge.
- If component through-holes are used for test probe locations, make sure leads are robust enough for compressive force (be careful using LEDs or some types of transformers). Also, make sure PTH leads will be present on all versions of assembly (not depopulated).

Much more information can be found in the SMTA (Surface Mount Technology Association) and TMAG (Testability Management Action Group) testability guidelines.[TP-101D]

In-Circuit Testing For test points, the following is preferred:

1. An area of 125mil wide must be free of components and test pads around the edge of the board.
2. The test pad should be at least 40mil in size; 75mil is preferred.
3. The test pads should be no closer that 50mil center-to-center; 100mil is preferred. At 50mil, there is an increase in false failures.
4. Test points should be on one side only. Two sides are possible but more expensive.
5. An 18mil annulus free of components around each test pad is required.
6. A test point should be no closer that 25mil (pad edge to center of via) from a via to prevent solder bridging.
7. Test points should be free of solder mask and ink.
8. Do not extend component pads to make test points. This may cause components to float.
9. 100% coverage for all nets on the design, including power, ground, and unused gates.
10. Fill through-holes with solder or solder resist; this is to provide a vacuum seal.

Flying-Probe Testing When ICT requirements cannot be accommodated, flying-probe test points should be attempted. Caution: flying-probe testing takes very long and is not for high volume modules or to be used on a production basis.

The following is less than ideal:

1. No component or test pad within 100mil of the perimeter of the PCB.
2. Minimum size of the test pad is 28mil. 20 is the lower limit before you start getting an increase in false failures. 10mil is very unreliable and the assembly house should should be contacted for a review of the locations.
3. Minimum pad spacing is 50mil center-to-centre (i.e. 25mil spacing, pad edge-to-edge). This is to prevent bridging during wave soldering. If wave soldering will not be used, the minimum spacing is 50mil (20mil spacing between pads).
4. 18mil annulus free of components around each test pad.
5. Test points should be free of solder mask and ink.
6. Do not extend component pads to make test points. This may cause the component to float.
7. 100% coverage for all nets is preferred, including power, ground, and unused gates. Engineering judgement should be used regarding which traces to verify (critical nets) if 100% coverage is not possible.

5.21.1 Probe Process

5.21.2 Probe Materials

None.

5.21.3 Probe DFX Rules

5.21.4 Probe Design

5.21.5 Probe Imaging

5.21.6 Probe Post-Processing

5.22 Fixture Layers

5.22.1 Fixture Process

5.22.2 Fixture Materials

None.

5.22.3 Fixture DFX Rules

5.22.4 Fixture Design

5.22.5 Fixture Imaging

5.22.6 Fixture Post-Processing

5.23 Comment Layers

The comment layer in pcbnew has historically been used for annotations. These annotations are typically not post-processes and do not appear in any fabrication or assembly print. This layer was typically used just for private designer annotations.

5.23.1 Comment Process

5.23.2 Comment Materials

None.

5.23.3 Comment DFX Rules

5.23.4 Comment Design

5.23.5 Comment Imaging

5.23.6 Comment Post-Processing

5.24 Drawing Layers

Historically, pcbnew has used the "drawing" layer for the fabrication print master drawing, and used the "silk screen" layers for assembly drawings. The "comment" layer was used typically for design annotations. Drill maps were plotted separately and not in synchronization with the other sheet references. This approach was not terribly useful or successful for a number of reasons:

1. Fabricators prefer that primary PTH, NPTH and tooling holes and the associated drill maps and stack-up be provided on the master fabrication print wherever possible. The fabrication master print should contain all of the information necessary to successfully quote and steer fabrication of the entire job. Detail sheets should only provide the minute details of a particular item. The major difficulty with the historical approach is that the primary drill map is plotted separate, and yet the hole sizes, counts and map are expected to be on the master print.
2. Assembly drawings should be printed as viewed from the appropriate process side of the board. That is, the bottom side of the board should be printed as viewed *from the bottom of the board*. This is different from fabrication prints that are *always printed as viewed from the top of the board, through the board*. A feature was provided in plots to mirror the layer and add sheet references, but unfortunately, the file names of the silkscreen layers used for assembly drawings were not able to be renamed in the plot dialogue. This required manipulation of files (error prone) or requiring the assembler to use silk as assembly drawings.
3. Detailed fabrication drawings could not be provided on separate sheets for detail that cannot be provided on the master print (because it would make it too busy to be legible). Items that need to be printed on fabrication or assembly detail sheets include detailed drill maps for blind, buried, back-drilled and microvias; select NPTH; the location for UL logos and such; details on cutouts or wells; description of gold masks; selective electroplated gold contacts, carbon contacts, and secondary finishes; mechanical drawings of brackets and bezels; stack-up drawings that cannot fit on the master print; and on and on...
4. The sheet number on sheet references could not be controlled by the designer so that other tools can be used to provide sheets for insertion into the sequence. For example, the designer should be able to number sheet 2 of 10. Otherwise, it should be possible to import and export XFIG, SVG, or EPS drawings, and have them able to be plotted properly on RS-274X Gerber output.
5. To avoid error, it should be possible to embed process and material parameters into textual comments so that the designer cannot neglect to update fabrication print comments. This, of course, would require some sort of macro or metalanguage for textual comments.
6. Providing lists of comments on fabrication prints is tedious. A paragraph or enumerated list format, such as HTML, should be possible for text. Perhaps a new text-box board item should be created.
7. Some items, such as drill map legends, should be treated as a block or group so that they can have their position controlled by the designer.

To provide the flexibility required for fabrication and assembly master prints and detail sheets, a facility needs to be added to technical layers that identifies to which fabrication or assembly sheets the information is to be copied. This should be possible for all board and technical layers. In conjunction with copying the technical layer to a fabrication or assembly drawing, there should also be attributes that can be defined such as default or minimum pen width, colour, etc. A reasonable default should be provided for each layer. For example, it reasonable that the silk screen (or legend) layers be printed on the corresponding primary assembly drawings (`Assem_Front` and `Assem_Back`) for the back and front of the board. It is reasonable that the component outline layers be printed on the corresponding primary assembly drawings (`Assem_Front` and `Assem_Back`). *Tab. 10(117)* lists the possible default mapping of technical layers onto fabrication and assembly drawings.

There should also be a facility for allocating new fabrication or assembly print layers, to which the designer can add free-form text and drawings.

5.24.1 Drawing Process

5.24.2 Drawing Materials

None.

5.24.3 Drawing DFX Rules

5.24.4 Drawing Design

5.24.5 Drawing Imaging

5.24.6 Drawing Post-Processing

5.25 Drilling Data

5.25.1 Drilling Process

Plate-Through Holes

Non-Plate-Through Holes

Tooling Holes

Breakaway Tab Holes

Through Via Holes

Blind Via Holes

Buried Via Holes

Microvia Holes

Back-Drilled Via Holes

Depth-Control-Drilled Via Holes

Depth-Control Drilling

Laser Drilling

Back-Drilling

5.25.2 Drilling DFX Rules

5.25.3 Drilling Design

5.25.4 Drilling Imaging

5.25.5 Drilling Post-Processing

Table 10: Default Mapping of Technical Layers

Layer	**Print/Description**
Resistance	*None.* (Represented by artwork plots.)
Copper	*None.* (Represented by artwork plots.)
Plating	*None.* (Represented by artwork plots.)
Hole Filing	*None.* (Represented by artwork plots.)
Keep-Out	*None.* However, keep-outs that are defined for thieving should be copied to the fabrication print.
Via Capping	*None.* (Represented by artwork plots.)
Mask	*None.* (Represented by artwork plots.)
Contacts	*Fab#1, Assem#1,#2.* Contacts are printed, by contact finish and board side, on fabrication prints. Contacts are also printed on assembly drawings (as are pads). The default is to print all contacts on the fabrication master drawing and on the assembly drawings front and back.
Peel Mask	*None.* (Represented by artwork plots.)
Finish	*Fab#2.* Secondary finishes printed on a fabrication detail sheet by default.
Legend	*Assem#1,#2.* Legend (silk screen) is printed by default on the assembly drawings for front and back. This is for backward compatibility. Along with printing the silk, printing of module fields (reference, value, etc.) should also be controllable.
Paste	*None.* (Represented by artwork plots.)
Adhesive	*None.* (Represented by artwork plots.)
Courtyard	*Assem#1,#2.* Component courtyards are printed by default to the assembly drawings for front and back to demonstrate proper spacing for rework and other assembly requirements.
Component	*Assem#1,#2.* Component outlines, and anything drawn or written on the component layers, are printed by default on the assembly drawings for front and back. In fact, the component front and back layers *are* the assembly front and back layers.
Coating	*None.* (Represented by artwork plots.)
Edges	*All.* Edges are printed on all fabrication and assembly sheets by default.
Eco	*Assem#1,#2.* Eco1 and Eco2 layers are printed by default on assembly front and back sheets. This is for possible backward compatible usage.
Probe	*Assem#1,#2.* The normal plot for the probe layer is the pad master front and back, so the default is to print the pads on the assembly front and back drawings.
Fixture	*Assem#3,#4.* Printed on its own assembly detail sheet for front and back by default.
Comment	*None.*
Drawing	*Fab#1.* The single historical drawing layer is printed by default on the fabrication master print for backward compatibility. The drawing layer from modules are printed by default on the fabrication master print. Module drawings should be separately controllable.
Drill Data	*Fab#1, #n.* Primary drill maps (through holes) and summary drill reports are printed by default on the fabrication master print, Fab#1, and secondary drill maps are printed each on a separate fabrication detail print.
Sheet Ref.	*All.* Sheet reference is printed by default on all fabrication and assembly prints and detail sheets.
Other	*Select.* All drawings should be selectable as to the fabrication or assembly print or detail sheet to which they belong.

6 Manufacturing Outputs

Fabrication outputs as referred to as design files, or the design database. There are largely two ways of providing the fabrication outputs to a board fabricator or manufacturer: as the traditional fabrication outputs (Gerbers, NC drill files, netlist), or as a CAM formatted database (GenCAM, GenX, ODB++, 258X). Historically pcbnew has provided only traditional fabrication outputs (and without a netlist).

6.1 Fabrication Outputs

Traditional data requirements for board fabricators are described in the subsections that follow.

README. The classic README file. This file contains customer and contact information, a table of contents to the other files in the package, their names, general contents and basic format; general instructions. The README file requirements are detailed in *Sec. 6.1.1(118)*.

Fabrication Drawings. Fabrication drawings consist of the master drawing and any number of detail drawings. The fabrication print must provide all information concerning acceptance, fabrication specifications, materials, special instructions and reference to supporting artwork, etc. The fabrication print and detail requirements are detailed in *Sec. 6.1.2(118)*.

Assembly Drawings. Assembly drawings consist of component outlines and placement. Assembly drawings are detailed in *Sec. 6.1.3(120)*.

Artwork. Artwork file provide all of the supporting artwork for each process contained in Gerber RS-274X, Barco DPF files, or other supported formats. Artwork files and content are detailed in *Sec. 6.1.4(120)*.

NC Drill and Rout. Excellon drill and route files. NC drill and rout files and their contents are detailed in *Sec. 6.1.5(123)*.

Netlist. Netlist information for artwork verification, flying-probe and ICT. Netlists and their requirements are detailed in *Sec. 6.1.6(123)*.

It is normal to zip the entire set of files into a single zip directory. This ensures that the data set is delivered complete, that corruption has not occurred, as well as compressing the data set for transmission. Note that fabricators might prefer the DOS-compatible zip format, or the UNIX tar gzip format. Many fabricators should accept either. Zip is the most common.

Bad or unknown compression. Electronic files are generally compressed prior to being sent. This minimizes memory size and transmission times. The file cannot be decompressed if it has an unknown compression technique. The file may fail decompression, which will also make it unusable. Bad or unknown compression is a critical *package completeness* issue. A new file must be sent to correct this issue.

Corrupted files. The corrupted file has been decompressed successfully, but it does not conform to a valid protocol or format. Corrupted files are considered a critical *package completeness* issue. A new file must be sent to correct this issue.

6.1.1 README

The classic README file. This file must contain the following information:

- Customer identification information.
- Contact information of responsible people.
- A list of the design files in the package, their file names, general contents and basic format information.
- General instructions (but no duplication of instructions contained on the fabrication print).

In particular, when RS-274X and NC drill and rout files are provided, this README file acts as a table of contents to the files in the data package. When the design files are in a CAM format (GenCAD, GenCAM, GenX or 258X) all of the data is contained within the one file and contains its own general information, depending on the level of the file.

Missing README file. The README file is created in ASCII and provides information on what each of the electronic files is used for and how many should be included. It may also include special instructions for how to combine files. A missing README file is considered a critical *package completeness* issue.

Sender not identified. The README file shall include text that indicates who sent the data to the fabricator. It should include the name of the company, division (if applicable), contact name, telephone number, fax number and e-mail address. When this information is not specified it is considered a *package completeness* issue.

6.1.2 Fabrication Prints

Fabrication drawings consist of the master drawing and any number of detail drawings. The master drawing should contain the following:

- notes providing the acceptance requirements for board fabrication;
- special instructions for advanced features with reference to the corresponding artwork files;
- stack-up (but the stack-up may also be placed on detail sheets);
- drilled hole sizes and types (primarily drill data on through-holes, NPTH, tooling, that pass through the entire board: drill maps for buried, blind, micro, depth-control-drilled and back-drilled vias should be contained on detail sheets);
- critical board dimensions and distance to datum or reference holes;
- panels and arrays (but these can also be contained on detail sheets);
- identification of all materials not fully specified by the artwork or CAM files, such as solder mask materials, legend, laminate, copper foil or cladding (especially things such as roughness for high-speed boards);

In addition to the master drawing, detail sheets are required for any information that is not contained or fully specified in the artwork or CAM files such as:

- drill maps for blind, buried, micro-, depth-control-drilled and back-drilled vias;
- locations for placement of logos, UL certification, and other legend markings to be added by the fabricator;[45]
- detail of assembly rails, bevels, perforations, and other rout path (even when NC rout path files are provided);
- identification of dimensions and placement of any thieving or venting keep-outs (especially when thieving and venting is provided by the fabricator);

All fabrication and detail sheets, when plotted, should use the same axis, origin and scaling as the artwork. All fabrication detail sheets should be as viewed from the top of the board.

The fabrication drawing fulfills several important functions to support the circuit board manufacturing process:

45. Typical mark is 0.011" tall × 0.026" wide. The requirement for UL marking must be identified by the customer on the fabrication print. The marking is placed on the outer layer (in copper).

- It provides enough information about the design for the bare-board fabricator to prepare a quote.
- It adds all the extra details about the build that aren't easily incorporated into the data files, like material and final finish.
- It lists the criteria by which the finished product will be evaluated for acceptability.
- It serves as a tool to be used during final inspection.
- It is a record or document to store the history of a product by title, part number and revision; physical dimensions, and lists the name of the designer and possibly several other supporting entities as well as the company name and address, etc.

Several different types of information can be found on a typical fabrication drawing, including design parameters, board details including construction, material and process specification, conductor and clearance minimums. Fabrication allowances, marking requirements, soldermask, silkscreen and final finish. Test coupon requirements, test requirements, performance requirements, and a graphical representation of the board with critical dimensions.

Mose of these topics are covered by a few "building blocks" that you will find on most drawings. The drill chart, a layer stack-up or construction detail, and a set of typical drawing notes.

The notes are like an instruction manual for manufacturing you product. Many of them will reference acceptability requirements. For example:

```
NOTES -- UNLESS OTHERWISE SPECIFIED:
```

1. FABRICATE TO MEET OR EXCEED THE REQUIREMENTS OF IPC-6012 FOR CLASS 3 AS DEFINED IN IPC-6011.
2. MINIMUM TRACE WIDTH .15, MINIMUM CLEARANCE .15
3. MATERIAL: LAMINATE FLAME RETARDANT EPOXY-GLASS PER IPC-4101/126, 170 DEGREE C MINIMUM Tg, DECOMPOSITION TEMPERATURE 340 DEGREES C MINIMUM, T288 DELAMINATION TIME OF 35 MINUTES MINIMUM, MAXIMUM THICKNESS EXPANSION OF 3% FROM 50--260 DEGREES C.
4. PREPREG MATERIALS PER /126 SHALL MEET THE SAME REQUIREMENTS.
5. INNER LAYER FOIL PER IPC-4562, TYPE E, GRADE 3, CLASS 2.
6. PLATING: ELECTRO-DEPOSITED COPPER, HOLE WALL PLATING AVERAGE MINIMUM .030, NO LESS THAN 0.025
7. FINISH: IMMERSION SILVER PER IPC-4553.
8. REGISTRATION: MINIMUM ANNULAR RING .0254, NO BREAKOUT ALLOWED, TEARDROPPING ALLOWED IF MIN CLEARANCE MAINTAINED.
9. FABRICATION: NON-FUNCTIONAL INNER-LAYER PADS SHALL NOT BE REMOVED FROM LAYERS 1, 2, 3, N-2, N-1, and N.
10. THIEVING IS ALLOWED IF 2.54 MINIMUM CLEARANCE TO CONDUCTIVE FEATURES IS MAINTAINED
11. SOLDER MASK: LPI BOTH SIDES OVER BARE COPPER PER IPC-SM-840, CLASS T, .02-.05 THICK, COLOR: MATTE GREEN. ALL FIDUCIALS, LANDS AND HOLES, EXCEPT VIAS, SHALL BE FREE OF MASK MATERIAL.
12. SILKSCREEN: WHITE NON-CONDUCTIVE EPOXY INK MuST WITHSTAND PEAK TEMPERATURE OF 260 DEGREES C 60 SECONDS, 3 CYCLES WITHOUT DISCOLORATION.
13. MARKING: VENDOR LOGO FOLLOWED BY 94V-0 AND FOUR-DIGIT DATE CODE.
14. FINISHED THICKNESS 1.6 +/-10%

```
ADDITIONAL SAMPLE NOTES INCLUDE:
```

1. BOARDS SHALL MEET OR EXCEED THE REQUIREMENTS OF IPC-6012 QUALIFICATION AND PERFORMANCE OF RIGID PRINTED BOARDS FOR CLASS 3.
2. ALL BOARDS WILL BE PERMANENTLY MARKED WITH THE PART NUMBER AND REVISION LEVEL THAT CORRESPONDS TO THE PART NUMBER AND REVISION LEVEL OF THE FABRICATION DRAWING. BOARDS WILL BE PERMANENTLY MARKED WITH MANUFACTURER'S REGISTERED UL LOGO AND 4-DIGIT DATA CODE.
3. A NETLIST WILL BE GENERATED FROM THE GERBER DATA (SUPPLIER CAM DEPARTMENT), AND COMPARED AGAINST A NETLIST PROVIDED WITH THE DESIGN DATA. ANY DISCREPANCY IN THE TWO NETLISTS WILL BE COMMUNICATED AND RESOLVED BEFORE STARTING THE BARE BOARD MANUFACTURING PROCESS.
4. ALL MULTI-LAYER BOARDS MUST PASS 100% ELECTRICAL TEST.
5. IF A LOT EXCEEDS 1000 BOARDS, THE LOT WILL BE DIVIDED INTO SUB-LOTS OF NO MORE THAN 1000.
6. A MICRO-SECTION AND FIRST ARTICLE INSPECTION (FAI) REPORT SHALL ACCOMPANY EVERY LOT OR SUB-LOT.
7. QUALITY CONFORMANCE TEST COUPONS WILL BE MANUFACTURED WITH EVERY LOT, AND DELIVERED WITH THE CIRCUIT BOARD PACKAGE AND FAI REPORT.

Missing fabrication drawing/file(s). The fabrication file presents mechanical board information and critical dimensioning. Fabrication notes are also included. It is required in conjunction with the artwork files. Missing fabrication drawings or details is considered a critical *package completeness* issue.

Non-electronic fabrication drawing. A non-electronic fabrication print will require that paper copies be distributed for production. Production or tooling will stop if the print is damaged or lost. Electronic prints allow multiple copies to be readily created and the master can be electronically archived. Non-electronic fabrication prints are considered a non-critical *package completeness* issue.

Fabrication drawing is not legible. Electronic fabrication drawing illegibility is determined when it is plotted or printed. The plotted quality must maintain good legibility after being copied. Hard copy fabrication drawings, such as fax's and plots, must maintain good legibility after two generations of copies are created. It is quite easy to cram too much information into too large a sheet size resulting an illegible fabrication print when printed. Illegible fabrication prints is considered a critical *data quality* issue.

Missing fabrication drawing information. All mechanical and non-artwork information that is not included in the README file or specifications provided, must be on the fabrication print. This includes items such as all printed board dimensions and tolerances and reference specifications. Missing fabrication drawing information is considered a critical *data quality* issue.

Missing PWB material information. The material requirements must be provided in the fabrication drawing or in the specifications provided. Missing material information is considered a critical *data quality* issue.

Missing layer sequence and stack-up. The layer sequence shows the order that the layers are fabricated in the printed board (e. g., Layer 1 is top copper, Layer 2 is next, Layer 3, then Layer 4 is the bottom side copper layer). It is preferred that the CAD layer file names are references to the actual board layer numbers.

This will ensure that layers are not inadvertently transposed. The stack-up also provides the layer spacing and tolerance information. Impedance values and tolerances are also included in the stack-up information. Missing layer stack-up information is considered a critical *data quality* issue.

PWB data does not match fabrication drawing. Conflicting information between the printed board graphical files and the fabrication drawing includes items such as wrong dimensions, missing layers and additional layers. These are considered critical *DRC* issues.

PWB array data does not match fabrication drawing. The fabrication drawing specifies a different sub-panel array than the PWB data provides. For example, a 2 x 4 array is specified on the drawing and a 4 x 8 is provided in the data. This also includes sub-panel rail dimensions and attributes. These are considered critical *DRC* issues.

OEM and Assembler specification conflict. The OEM provides a fabrication acceptance specification. The assembly is performed by one or more sub-contract assemblers, each with their own fabrication acceptance specification. There should not be any conflicting specifications when these documents are compared. This is considered a critical *DRC* issue.

PWB data conflicts with OEM specifications. The PWB database violates requirements that are established in the OEM fabrication specification. The database must be updated or a waiver must be provided. This is considered a non-critical *DRC* issue.

Inefficient fabrication panel utilization. The board size cannot be placed into the fabrication panel efficiently. This could include items such as actual board dimensions and coupon requirements. These problems are considered non-critical *DFM* issues.

Non-conformance to PWB thickness tolerance. The board cannot be fabricated per the specified thickness tolerance. This can be a result of too many layers begin included for a given thickness. A tolerance may be specified too tight. There may also be insufficient thickness provided to allow compliance with the impedance requirements. These problems are considered non-critical *DFM* issues.

Non-conformance to impedance tolerances. Compliance with the specified impedance and/or tolerance cannot be maintained due to line width or dielectric spacing restrictions. These problems are considered non-critical *DFM* issues.

6.1.3 Assembly Drawings

Assembly drawings consist of component outlines and placement. Any materials or details concerning assembly and acceptance should be placed on the assembly drawings. Assembly drawings depicting the top of the board should be printed as viewed from the top of the board. Assembly drawings depicting the bottom of the board should be printed as view *from the bottom of the board*. This can be most easily accomplished by reflecting the board on the x-axis about the center of the board.

PWB data conflicts with PWA specifications. The PWB database violates requirements that are established in the PWA fabrication specifications. The database must be updated or a waiver must be provided. This is considered a non-critical *DRC* issue.

6.1.4 Artwork

Classic formats are Gerber RS-274D and RS-274X. All fabricators can accept these formats. Also, may fabricators also support Barco DPF. Note that all three formats (RS-247D, RS-274X and DPF) require one file per image layer. Image layers to be included are all board fabrication image layers, as well as master drawing and fabrication drawings for board fabrication. Assembly drawings are necessary for manufacturing.

Most fabricators appear to prefer Gerber plots for board edges (however, note that IPC-2524 [IPC-2524] requires a route path for board edges (mechanical) layers). Also, many request that board edges be plotted on all layers. Some care about out-of-bounds data (such as sheet reference frames); others do not.

In addition, there should be board fabrication master drawings, as well as drill maps for each drill file¡ also plotted in Gerber or DPF format.

Copper: One file each for each of the conductive layers in a multilayer board. Files should be named according to layer. The fabrication print or stack-up detail should provide file names for copper layers on the stack-up diagram. Special apertures should be used for any traces requiring impedance control.

Dielectric:

Plate:

Fill: Separate data is required for each sub-laminate stack-up as well as for the entire board. Data consists of Gerber or DPF where the aperture size is the FHS of the hole to fill. Hole filling materials for each hole filling layer pair must be identified on the fabrication print. Data is not always required when holes are not filled selectively. When NPTH and PTH are mixed and first-drilled, hole filling is always selective for the full board and corresponding data files must be provided.

Conductive filled holes should have a finished hole size $\geq$ 0.010". Conductive filled holes should be indicated on the fabrication print with a note stating:

```
13. PROVIDE VIA FILL WITH DUPONT CB100 OR EQUIVALENT
    CONDUCTIVE MATERIAL FOR THE F.H.S. HOLES PROVIDED
    IN THE BOARD-HOLEFILL-XXX.GBR FILES.
```

Non-conductive filled holes should have a finished hole diameter $\geq$0.006". Non-conductive filled holes should be indicated on the fabrication print with the following note:

```
14. PROVIDE VIA FILL WITH PETERS PP2795-SD OR EQUIVALENT
    NON-CONDUCTIVE MATERIAL FOR THE F.H.S. HOLES PROVIDED
    IN THE BOARD-HOLEFILL-XXX.GBR FILES.
```

Conductive filled holes should have a F.H.S. $\geq$0.010". Capped or plated over blind or through-hole vias with conductive filled holes should be indicated on the fabrication print with a note stating:

```
15. PROVIDE VIA FILL WITH DUPONT CB100 OR EQUIVALENT
    CONDUCTIVE MATERIAL FOR THE HOLES IDENTIFIED WITH
    F.H.S. APERTURE IN FILE BOARD-HOLEFILL-XXX.GBR, AND
    PLATE THE SURFACE WITH 0.0005" MINIMUM OF COPPER.
```

Non-conductive filled holes should have a F.H.S. $\geq$0.006". Capped or plated over blind or through-hole vias with non-conductive filled holes should be indicated on the fabrication print with a note stating:

```
16. PROVIDE VIA FILL WITH PETERS PP2795-SD OR EQUIVALENT
    NON-CONDUCTIVE MATERIAL FOR THE HOLES IDENTIFIED WITH
    F.H.S. APERTURE IN FILE BOARD-HOLEFILL-XXX.GBR, AND
    PLATE THE SURFACE WITH 0.0005" MINIMUM OF COPPER.
```

The following note should be incorporated into the fabrication print for buried via holes filled with conductive filling and plated over:

Table 11: Required Artwork Layers

Type	Layer	Name	Ext.	Comments
Cu	Top			Top, front, component, primary copper layer information.
Cu	Inner			Inner copper layers, one file for each layer.
Cu	Bottom			Bottom, back, solder, secondary copper layer information.
Fill	Top			Hole filling from the top of the board.
Fill	Pairs			Hole filling for internal buildup layers.
Fill	Bottom			Hole filling from the bottom of the board.
Pads	Top			
Pads	Bottom			
Contacts	Top			
Contacts	Bottom			
Finish	Top			
Finish	Bottom			
Plug	Top			
Plug	Bottom			
Mask	Top			Top primary (solder) mask layer.
Mask	Bottom			Bottom primary (solder) mask layer.
Silk	Top			Silk screen top layer required when top layer legend is to be printed. Fabrication print should identify colour, minimum line width, minimum character height, whether clipping is required or has been performed.
Silk	Bottom			Silk screen bottom layer, required when bottom layer legend is to be printed. Fabrication print should identify colour, minimum line width, minimum character height, whether clipping is required or has been performed.
Paste	Top			Required when SMT components are present on the top layer.
Paste	Bottom			Required when SMT components are present on the bottom layer.
Glue	Top			
Glue	Bottom			

```
17. PROVIDE VIA FILL WITH DUPONT CB100 OR EQUIVALENT
    CONDUCTIVE MATERIAL FOR THE 0.0XX" HOLES AND PLATE
    THE SURFACE WITH 0.0005" MINIMUM OF COPPER.
```

The following note should be incorporated into the fabrication print for buried via holes filled with non-conductive filling and plated over:

```
18. PROVIDE VIA FILL WITH PETERS PP2795-SD OR EQUIVALENT
    NON-CONDUCTIVE MATERIAL FOR THE 0.0XX" HOLES AND PLATE
    THE SURFACE WITH 0.0005" MINIMUM OF COPPER.
```

The following note should be incorporated into the fabrication print for buried via holes filled with resin:

```
19. FILL ALL VIAS USING THE PCB RESIN.
```

— Holes to be plugged must be identified on the drill chart to prevent errors.

— Component, solder, or both sides can be plugged. Plugging is not recommended for covering both sides of the via on the board; use filled holes instead. Air or moisture could become trapped inside causing blow-out during assembly.

— With holes $\leq$ 15mil, simply screening LPI over the via, the via will "tent." However, after assembly (thermal stress, washing, vibration, etc.), there is no guarantee that the tent will remain intact. PCB manufacturers recommend either plugging or filling vias. Delamination of the solder mask may result. Additionally, PCB vendors indicated that chemicals could become trapped in the tented via causing reliability problems and corrosion later in time.

— There are some differences in terminology. Some vendors say "plugged" to mean one or both sides of the hole are plugged with epoxy. Some make the distinction that "plugged" is where only one side of the via is covered with epoxy; whereas "filled" is where both sides of the via are covered with epoxy. On fabrication prints it is best to use the word "capped" to mean covered on one side with secondary mask; and, "filled" to mean the hole completely filled.

— Plugged holes: The minimum hole diameter is 8mil. The smaller the hole, the more difficult it is for the epoxy to flow into the hole. Maximum hole diameter is 20mil. The larger the hole becomes, the more difficult it is to maintain 100% coverage. PCB vendors cannot guarantee that solder resist will not seep onto the other side of the board through the via. Therefore, this should not be used for covering the component side vias where access to the secondary side is required (as in ICT). Use filled vias with conductive epoxy for such applications.

— Filled holes (an even more expensive option): Minimum hole diameter is 10mil. Maximum hole diameter is 18mil. Conductive or non-conductive epoxy can be used. Conductive epoxy can be used for areas requiring the component side to be covered but the solder side vias accessible for test points.

— Possible alternatives to plugging is using dry film to tent the vias on one side. Very few PCB manufacturers use dry film so it may be difficult to find such a manufacturer.

Pads: This is a pad master that includes all primary mask exposed pads. It can be used by the assembler to develop solder paste stencils. It can also be used to develop solder mask openings. It can also be for board testing probe point or ICT fixture development. *Note that this layer should not be provided in the data package, but should be copied to the necessary layer artwork.*

Contacts: This layer identifies any contacts such as goldfinger (edge tip). The contact material and finish must be specified on the fabrication print with a note such as:

```
20. EDGE CONTACTS ARE TO BE ELECTROPLATED WITH A MINIMUM
    OF 100 MICRO-INCHES OF NICKEL FOLLOWED BY A MINIMUM OF
    20 MICRO-INCHES OF GOLD.
```

Finish: Finish material and process must be specified on the fabrication print with a note such as:

```
21. FINISH: IMMERSION SILVER PER IPC-4553.
```

Where multiple finishes are required, plating masks should be plotted on this layer for top, bottom, or both, wherever selective finishes are required. Where selective finishes are required (typically nickel-gold), a fabrication detail sheet should be provided.

Selective non-bussed areas can be plated with nickel and gold before etching, by the use of a gold mask. These gold masks should be supplied in the data package. Simple gold masks can be created by the fabricator when detailed information is supplied as to which features will be plated.

Plug: Secondary mask (via capping) material must be identified on the fabrication print, as well as colour, whether clearances have been applied or plotted 1:1. When clearances have been applied, what clearance value was used must also be specified on the fabrication print.

Where via plugging (capping) is performed on one side of the board, non-selectively, a note can be placed on the fabrication print specifying that all vias are to be capped, the material/process, and the side of the board from which vias are to be capped. Via plugging artwork might not be necessary in this case. The fabrication print should identify whether via capping is selective or non-selective and which sides are involved.

When selective capping is performed, or some vias are capped from one side of the board; others, from the other side; it is always necessary to provide data files for the capping. Where capping is provided before primary mask, data will include apertures that are the via FHS for the capped vias (possibly including clearances). The primary mask data will cover over these vias completely on the capped side and provide full clearance or encroachment on the uncapped side. Where capping is provided after primary mask, the primary mask layer will have apertures of FHS (possibly increased for clearance) on both sides of the board, and the plug data will have via pad size (possibly increased for clearance) for the caps.

When the primary mask is plotted 1:1 and fabrication notes are provided instructing the fabricator to adjust primary mask clearances to their processes, it is necessary to do the same for the via plugging (secondary mask) layers. That is, caps are plotted as via F.H.S. When capping is before mask; and via pad size when capping is after mask. When primary mask is plotted with clearances, solder dams and gang relief, it is necessary to do the same for the via plugging. That is, caps are plotted as via F.H.S. expanded by primary mask clearance when capping is performed before primary mask; and via pad size expanded by primary mask clearance when capping is after mask.

As with hole filling, vias that are capped should also be identified on the primary drill map.

Some example fabrication notes follow:

Mask: Primary (solder) mask material must be identified in the fabrication print, as well a colour, whether clearances have been applied or plotted 1:1, solder dams and gang relief instructions, hole tenting, encroachment. When clearances have been applied, what clearance value was used.

For encroachment, the fabrication print should have a note specifying the via coverage rules. For example:

```
22. THE VIAS WITH PRIMARY MASK APERTURES THE SAME SIZE AS
    F.H.S. AND EXPANDED BY 0.003" PRIMARY MASK CLEARANCE
    WERE DESIGNED WITH SOLDER MASK ENCROACHED ON THE
    VIA-PADS. MANUFACTURER CAN ADJUST THE SOLDER MASK
    CLEARANCE SIZE PER THEIR PROCESS TO ASSURE VIA HOLES
    ARE OPEN FOR IMMERSION-SILVER FINISH, WHILE KEEPING
    THE PADS ENCROACHED WITH SOLDER MASK AS MUCH AS
    POSSIBLE.
```

Silk: Fabrication print should identify colour, minimum line width, minimum character height, whether clipping is required or has been performed. Example fabrication note:

```
23. LEGEND: WHITE NON-CONDUCTIVE EPOXY INK MUST WITHSTAND
    PEAK TEMPERATURE OF 260 DEGREES C FOR 60 SECONDS, 3
    CYCLES WITHOUT DISCOLOURATION.

    DESIGNED MINIMUM LINE WIDTH 0.007", MINIMUM CHARACTER
    SEPARATION 0.007", MINIMUM CHARACTER HEIGHT 0.035".
    LEGEND ARTWORK HAS BEEN CLIPPED WITHIN 0.003" OF
    SOLDER MASK OPENINGS.
```

Peel:

Paste:

Glue:

General data requirements: The following are general data requirements that speak to the quality of the artwork:

- For plot-and-go, fabricators require a data package that contains a full set of artwork files. Extraneous, redundant or additional files not required for fabrication should *not* be provided.
- For plot-and-go, all artwork must be provided in Gerber RS-274X or DPF format, with no reliance on fabrication print details.
- For full-service fabrication, fabrication print details must not conflict with provided artwork.
- Never use RS-274D; always use RS-274X or DPF instead. ASCII aperture reports are useful for quotation or prototype services, but not for production, as they must be compared with artwork for conflicts. Aperture reports must not be included for plot-and-go services.
- All artwork layer file names must be short and expressive of the content of the file. Proper Gerber extensions can be optionally used. The image name should be specified within the file and should be consistent with the file name.
- All thermals should be included in the database, and must be included for plot-and-go services. Thermals should not have to be manually typed in by the fabricator.
- All pads must be flashed. Complex pad shapes must have an aperture macro (RS-274X) or complex aperture (DPF) defined. Complex apertures should be reduced to a simple aperture wholly contained within the complex aperture for the pad master (probe) layer.
- No merge layers, or layers that are to be used for multiple purposes should be provided. All layers that are used multiply should be copied to a separate file and named accordingly. No merge files should be provided. All image information must be self-contained within a single file with no included files.
- All copper zones must be represented with contours and must not be “painted” with segment draws.
- Gold mask artwork must be provided for selective gold finish. Peelable mask layers should also be provided, or the gold mask artwork must also be usable as a peelable mask layer.

Missing PWB artwork file(s). The PWB artwork files are required to provide the electronic data required to fabricate each layer, such as copper, drill and solder mask layers. Missing PWB artwork files is considered a critical *package completeness* issue.

Missing aperture information/file for 274-D artwork files. The aperture file is used to provide D-Code information for the 274-D Gerber files. This is not required for 274-X or other data formats. When required, missing aperture information is considered a critical *package completeness* issue.

Non-electronic aperture list. A hard-copy aperture list will be manually entered into the CAM system. An incorrectly entered aperture may not be discovered during the rest of the tooling process. Non-electronic aperture lists are considered a non-critical *package completeness* issue.

Rename layers, align artwork layers. Layers should be numbered in the order and system that the fabricator uses in their CAM software. All of the layers should automatically be aligned to each other when the files are extracted. Individual layers should not require shifting or scaling. Unaligned or scaled, or misnamed artwork layers are considered a non-critical *data quality* issue.

Copy solder mask layer or create a layer. In many designs a single layer is used multiple times in the same printed board. This is common for solder mask layers. The layer should be copied and renamed in the CAD database prior to sending the database to the fabricator. No layers should be created by the fabricator. Layer creation or merge requirements are considered a non-critical *data quality* issue.

Substitute flashes for draws of test points and pads. CAM systems utilize a significant amount of intelligent, automated editing. Most of the automation requires that printed board features, such as component pads and test points, be created with flash apertures and not drawn. Test fixture software assigns test points to flashed pads that do not have solder mask over them. Drawn pads must be manually converted into flashed pads during data entry or test fixture program generation. Drawn pads is considered a non-critical *data quality* issue.

Type in thermals. All thermal pads should be included in the database. If they are excluded, it will require manual entry during the data input operation at the fabricator. Missing thermals are considered a *data quality* issue.

Delete fabrication lines (only when tips are connected). Fabrication lines (tie bars) are required to be manually deleted for boards that have gold contacts (tips) which will be electroplated. Tie bars are considered a non-critical *(*data quality) issue.

Fabrication drawing conflicts with OEM specification. The printed board fabrication drawing violates requirements that are established in the OEM fabrication specifications. The drawing must be updated or a waiver must be provided. This is considered a non-critical *DRC* issue.

Fabrication drawing conflicts with PWA specification. The printed board fabrication drawing violates requirements that are established in the PWA fabrication specifications. The drawing must be updated or a waiver must be provided. This is considered a non-critical *DRC* issue.

Does not conform to fabricator critical reliability threshold. A feature, or combination of features, may create an unreliable condition in the fabricated printed board. This could include plated hole to power plane clearances, trace to edge of board clearances, etc. A decision must be made to fabricate the board with the attribute or modify the design. These problems are considered critical *DFM* issues.

Non-conformance to non-critical feature tolerances. The printed board cannot be fabricated to the specified tolerances. Common issues are finished hole location tolerances, slot tolerances, and rout location tolerances. These problems are considered non-critical *DFM* issues.

Non-conformance to solder mask requirements. The solder mask requirements cannot be complied to with the provided design database. Common issues are solder mask clearance and web widths. These problems are considered non-critical *DFM* issues.

PWB data file size is very large. The database is too large for the software. Common issues are large plane areas utilizing drawn vectors (0.025mm) and too many layers. Plane areas on mixed plan and signal layers should be contoured versus filled with drawn vectors. Drawn planes will substantially increase automated design rules check run times. These problems are considered non-critical *DFM* issues.

Incomplete surface finish requirements. No clear indication of type or tolerances for metallic or organic surface finish requirements for lands for component mounting, board edge connector contacts, and other component mounting technologies. These problems are considered non-critical *DFM* issues.

Critical parameter is outside of fabricator technical capabilities. A printed board feature or technology is outside of the fabricator's capability. This may also include a combination of features or characteristics of the design database. These problems are considered critical *DFM* issues.

6.1.5 NC Drill and Rout Files

Classic formats are Excellon-1 and Excellon-2 (IPC-NC-349) drill files. Some fabricators also support Gerbers (flashing is drilling, drawing is routing). Drill files to be included are those for all PTH (including DCD blind, back-drilled, PTH buried, NPTH).

Most fabricator appear to prefer Gerber plots for routing such as board edges. Some would like to see NPTH, route path, and other mechanicals in the same file. Files should be matched to the artwork files in terms of grid and precision.

Drill file(s) and layer connectivity (blind and buried via layer combinations) not specified. Drill files are required to provide all drilled hole locations. Printed boards with blind and buried vias should have a note on the fabrication print that specifies which files are utilized between specific layers. Not specifying these files and data where requires is considered a critical *package completeness* issue.

Missing rout/profile file. A missing rout or profile file will require complete manual creation. Missing rout/profile files are considered a non-critical *package completeness* issue.

Drill layers do not match PWB artwork files. The drill files and summary hole count table on the fabrication print must match. This includes the total hole quantity. All individual hole quantities must also match. This is considered a critical *DRC* issue.

Non-conformance to finished hole requirements. Drilled hole diameter and tolerances cannot be met for either supported or unsupported holes. This is usually related to hole to board aspect ratio, pad to hole tolerances, or location tolerance specification. These problems are considered non-critical *DFM* issues.

Non-conformance to rout, bevel, or score requirements. The board cannot be fabricated with the specified rout, bevel, or score requirements. This generally involves tolerances or non-standard combinations of features. These problems are considered non-critical *DFM* issues.

6.1.6 Netlist

Netlist information is only necessary when the artwork is not provided in a format that supports embedded netlists. Note that DPF provides netlist information, whereas, RS274D and RS274X do not. Nevertheless, DPF does not really provide all of the

information needed by bare-board testing (such as whether things are covered by primary mask or not, and from which side of the board they are accessible), and so an IPC netlist file is still necessary.

Netlists can be provided in IPC-D-356 (with or without an accomapnying NTD file), IPC-D-356A, IPC-D-356B.

Missing netlist, when a netlist test is required. A netlist, extracted from the original CAD database, is required when netlist test or data verification is required. The fabricator may extract a netlist from the interpreted data file and compare it back to the provided netlist. This verifies that no connectivity data corruption has occurred during the initial data extraction. A missing netlist file where one is required is considered a critical *package completeness* issue.

Missing netlist, if netlist test is not required. A netlist that is extracted from the CAD system should be sent with the CAD database. A netlist test can also be used for incoming data netlist comparison, electrical test, or both. This issue is considered non-critical because this requirement is not required for tooling and can be waived. A missing nelist file where one is not required is considered a non-critical *package completeness* issue.

Extracted netlist does not match supplied netlist. The fabricator planner extracts a connectivity netlist and compares it to the netlist that is provided with the PWB database. When the two netlists conflict, it must be resolved. Failure to resolve this confict could result in creating a PWB database that is electrically non-functional. This is considered a critical *DRC* issue.

6.1.7 Data Quality

A standard for the quality of the dataset specified by IPC [IPC-2524], provides a fairly good indication of the expectations of fabricators. Some additional expectations on the data set are as follows:

— All files should reference the same origin.

— All files should use the same units (Inches or Millimeters).

— All files should use the same resolution (typical is 2:4 for Inch and 3:5 for Metric).

— All Gerber or Barco plots should use the same master aperture list.

— Only those apertures that are used on a Gerber or Barco files should be defined in the file.

— No zero-size apertures allowed. (I can fool CAM systems into thinking that a pad or line is located where there is none.)

— As it is easy to give the traces or pads of certain apertures special treatment, a separate aperture number should be used for flashes or draws that cooresponds to special items, such as impedance controlled traces, special contact surfacing or special finishes.

— All NC drill files should use the same tool list.

— Only those tools used in an Excellon-2 drill file should have their sizes included in file.

— PTH, NPTH and VIAs of the same size should have different tool numbers because different plating compensation applies to each.

Package completeness problems include the following:

Missing PWB artwork file(s). The PWB artwork files are required to provide the electronic data required to fabricate each layer, such as copper, drill and soldermask layers.

Missing fabrication drawing/file(s). The fabrication file presents mechanical board information and ciritical dimensioning. Fabrication notes are also included. It is required in conjunction with the artwork files.

Missing README file. The README file is created in ASCII and provides information on what each of the electronic files is used for and how many should be included. It may also include special instructions for how to combine files.

Missing aperture information/file for 274-D artwork files. The aperture file is used to provide D-Code information for the 274-D Gerber files. This is not required for 274-X or other data formats.

Missing netlist, when a netlist test is required. A netlist, extracted from the original CAD database, is required when netlist test or data verfication is required. The fabricator may extract a netlist from the interpreted data file and compare it back to the provided netlist. This verifies that no connectivity data corruption has occurred during the initial data extraction.

Drill file(s) and layer connectivity (blind and buried via layer combinations) not specified. Drill files are required to provide all drilled hole locations. Printed boards with blind and buried vias should have a note on the fabrication print that specifies which files are utilized between specific layers.

Sender not identified. The README file shall include text that indicates who sent the data to the fabricator. It should include the name of the company, division (if applicable), contact name, telephone number, fax number and e-mail address.

Bad or unknown compression. Electronic files are generally compressed prior to being sent. This minimizes memory size and transmission times. The file cannot be de-compressed if it has an unknown compression technique. The file may fail decompression, which will also make it unusable. A new file must be sent to correct this issue.

Corrupted files. The corrupted file has been de-compressed successfully, but it does not conform to a valid protocol or format. A new file must be sent to correct this issue.

Non-electronic fabrication drawing. A non-electronic fabrication print will require that paper copies be distributed for production. Production or tooling will stop if the print is damaged or lost. Electronic prints allow multiple copies to be readily created and the master can be electronically archived.

Missing netlist, if netlist test is not required. A netlist that is extracted from the CAD system should be sent with the CAD database. A netlist test can also be used for incoming data netlist comparison, electrical test, or both. This issue is considered non-critical because this requirement is not required for tooling and can be waived.

Non-electronic aperture list. A hard-copy aperture list will be manually entered into the CAM system. An incorrectly entered aperture may not be discovered during the rest of the tooling process.

Missing rout/profile file. A missing rout or profile file will require complete manual creation.

Data Quality problems include the following:

Fabrication drawing is not legible. Electronic fabrication drawing illegibility is determined when it is plotted or printed. The plotted quality must maintain good legibility after being copied. Hard copy fabrication drawings, such as fax's and plots, must maintain good legibility after two generations of copies are created.

Missing fabrication drawing information. All mechanical and non-artwork information that is not included in the README file or specifications provided, must be on the fabrication print. This includes items such as all printed board dimensions and tolerances and reference specifications.

Missing PWB material information. The material requirements must be provided in the fabrication drawing or in the specifications provided.

Missing layer sequence and stack-up. The layer sequence shows the order that the layers are fabricated in the printed board (e. g., Layer 1 is top copper, Layer 2 is next, Layer 3, then Layer 4 is the bottom side copper layer). It is preferred that the CAD layer file names are references to the actual board layer numbers. This will ensure that layers are not inadvertently transposed. The stack-up also provides the layer spacing and tolerance information. Impedance values and tolerances are also included in the stack-up information.

Rename layers, align artwork layers. Layers should be numbered in the order and system that the fabricator uses in their CAM software. All of the layers should automatically be aligned to each other when the files are extracted. Individual layers should not require shifting or scaling.

Copy solder mask layer or create a layer. In many designs a single layer is used multiple times in the same printed board. This is common for solder mask layers. The layer should be copied and re-named in the CAD database prior to sending the database to the fabricator. No layers should be created by the fabricator.

Substitute flashes for draws of test points and pads. CAM systems utilize a significant amount of intelligent, automated editing. Most of the automation requires that printed board features, such as component pads and test points, be created with flash apertures and not drawn. Test fixture software assigns test points to flashed pads that do not have solder mask over them. Drawn pads must be manually converted into flashed pads during data entry or test fixture program generation.

Type in thermals. All thermal pads should be included in the database. If they are excluded, it will require manual entry during the data input operation at the fabricator.

Delete fabrication lines (only when tips are connected). Fabrication lines (tie bars) are required to be manually deleted for boards that have gold contacts (tips) which will be electroplated.

DRC problems include the following:

PWB data does not match fabrication drawing. Conflicting information between the printed board graphical files and the fabrication drawing includes items such as wrong dimensions, missing layers and additional layers.

PWB array data does not match fabrication drawing. The fabrication drawing specifies a different sub-panel array than the PWB data provides. For example, a 2 x 4 array is specified on the drawing and a 4 x 8 is provided in the data. This also includes sub-panel rail dimensions and attributes.

OEM and Assembler specification conflict. The OEM provides a fabrication acceptance specification. The assembly is performed by one or more sub-contract assemblers, each with their own fabrication acceptance specification. There should not be any conflicting specifications when these documents are compared.

Extracted netlist does not match supplied netlist. The fabricator planner extracts a connectivity netlist and compares it to the netlist that is provided with the PWB database. When the two netlists conflict, it must be resolved. Failure to resolve this confict could result in creating a PWB database that is electrically non-functional.

Drill layers do not match PWB artwork files. The drill files and summary hole count table on the fabrication print must match. This includes the total hole quantity. All individual hole quantities must also match.

PWB data conflicts with OEM specifications. The PWB database violates requirements that are established in the OEM fabrication specification. The database must be updated or a waiver must be provided.

PWB data conflicts with PWA specifications. The PWB database violates requirements that are established in the PWA fabrication specifications. The database must be updated or a waiver must be provided.

Fabrication drawing conflicts with OEM specification. The printed board fabrication drawing violates requirements that are established in the OEM fabrication specifications. The drawing must be updated or a waiver must be provided.

Fabrication drawing conflicts with PWA specification. The printed board fabrication drawing violates requirements that are established in the PWA fabrication specifications. The drawing must be updated or a waiver must be provided.

DFM problems include the following:

Critical parameter is outside of fabricator technical capabilities. A printed board feature or technology is outside of the fabricator's capability. This may also include a combination of features or characteristics of the design database.

Does not conform to fabricator critical reliability threshold. A feature, or combination of features, may create an unreliable condition in the fabricated printed board. This could include plated hole to power plane clearances, trace to edge of board clearances, etc. A decision must be made to fabricate the board with the attribute or modify the design.

Non-conformance to non-critical feature tolerances. The printed board cannot be fabricated to the specified tolerances. Common issues are finished hole location tolerances, slot tolerances, and rout location tolerances.

Non-conformance to soldermask requirements. The soldermask requirements cannot be complied to with the provided design database. Common issues are soldermask clearance and web widths.

PWB data file size is very large. The database is too large for the software. Common issues are large plane areas utilizing drawn vectors (0.025mm) and too many layers. Plane areas on mixed plan and signal layers should be contoured versus filled with drawn vectors. Drawn planes will substantially increase automated design rules check run times.

Incomplete surface finish requirements. No clear indication of type or tolerances for metallic or organic surface finish requirements for lands for component mounting, board edge connector contacts, and other component mounting technologies.

Non-conformance to finished hole requirements. Drilled hole diameter and tolerances cannot be met for either supported or unsupported holes. This is usually related to hole to board aspect ratio, pad to hole tolerances, or location tolerance specification.

Inefficient fabrication panel utilization. The board size cannot be placed into the fabrication panel efficiently. This could include items such as actual board dimensions and coupon requirements.

Non-conformance to rout, bevel, or score requirements. The board cannot be fabricated with the specified rout, bevel, or score requirements. This generally involves tolerances or non-standard combinations of features.

Non-conformance to PWB thickness tolerance. The board cannot be fabricated per the specified thickness tolerance. This can be a result of too many layers begin included for a given thickness. A tolerance may be specified too tight. There may also be insufficient thickness provided to allow compliance with the impedance requirements.

Non-conformance to impedance tolerances. Compliance with the specified impedance and/or tolerance cannot be maintained due to line width or dielectric spacing restrictions.

6.2 Assembly Outputs

CAD systems outputs are also generated to provide information for the assembly and assembly test phases of the product manufacture.

There is a stark contrast between treatment of the customer by board fabricators and assemblers. Board fabricators demand specific and detailed information, specific and particular data file formats, adherence to particular standards, and will place your order "on hold" as punishment for not meeting their demands. Assemblers, on the other hand, demand little and accommodate many particulars without complaint. They accept just about any kind of data format and will adjust to about every need. Perhaps the major difference between the two is than an error inside a multilayer board makes it unusable, where an error in attaching a component can almost always be repaired, manually if necessary. I think that maybe the most onerous requirement that I have seen stated by an assembler is: *"Thou shalt use IPC footprints."*

Most CAD systems generate a series of outputs that are only intended for assembly and have no bearing on board fabrication. In fact, many board fabricators offering a plot-and-go service, do not even want to see these assembly data outputs.

The fundamental assembly outputs are as follows:

BOM: BOM (Bill Of Materials) (mandatory). The BOM provides the part data necessary for ordering and loading pick-and-place machines. This is not really a full BOM in the CAM sense, but a *Parts List*.

See *Sec. 6.2.1(126)* for details.

Netlist: IPC-D-356 (preferred). IPC-D-356 can be used reliably for ARE (Automatic Reverse Engineering) of footprints and component placement. When not available, the silk or assembly drawings are used for ARE of footprints and component placement in an ever more increasingly manual process.

See *Sec. 6.2.2(127)* for details.

CAD Data: CAD File (IPC-D-356 or GenCAD) (preferred). Any of a number of CAD file formats are used for extracting footprint and component placement data.

See *Sec. 6.2.3(127)* for details.

Array Data: When the PCB outlines are beneath a certain size, they need to be placed into an array and panelized to create an assembly panel. Most turn-key and short production run assemblers have similar requirements. Full production assemblers might have unique requirements.

See *Sec. 6.2.4(128)* for details.

Centroid File: A pick-and-place or XYRS data file is typically provided in a CSV (Comma Separated Values) format generated by the CAD system. This file is often used simply to verify data.

See *Sec. 6.2.5(130)* for details.

Exterior Copper: Gerber (RS-274X) artwork for exterior copper layers containing pads (mandatory). Provides the information concerning pad geometry and scaling for ARE of footprints and placement. Does not provide component orientation information unless reference pins are marked in copper. Also, does not provide polarity information unless polarity is marked in copper.

See *Sec. 6.2.6(131)* for details.

Silk Screen: Gerber (RS-274X) artwork for silk screen (preferred with no clipping). When the other data is complete, silk screen data is simply used for verification of part placement, orientation and polarity. When other data is missing, the legend can be used for ARE (Automatic Reverse Engineering).

See *Sec. 6.2.7(131)* for details.

Solder Paste: Gerber (RS-274X) artwork for solder paste. Turn-key and short production run assembly shops will use this artwork to create solder paste stencils.

See *Sec. 6.2.8(131)* for details.

Assembly Drawings: Assembly drawings (additional). When no CAD data or legend is provided, the assembly drawings must be used to Reverse Engineer part placement, orientation and polarity. When full data is provided, they can still be used to verify or clarify placement. In addition, assembly drawings may contain special assembly instructions.

See *Sec. 6.2.9(131)* for details.

Solder Mask: Gerber (RS-274X) artwork for solder mask layers (additional). Some assembler want to see the solder resist layer; some do not. When Gerber copper data contains too much information (traces, via pads, contacts, etc.), and CAD data or IPC-D-356 are not provided, the solder mask layer can be used to determine the location of the exposed pads.

See *Sec. 6.2.10(131)* for details.

NC Drill Files: NC Drill Files for holes that transfix the board (optional).

See *Sec. 6.2.11(131)* for details.

Some additional comments about assembly outputs are as follows:

1. Almost no turn-key assembly shops list DPF as a possible format for artwork. RS-274X is the only choice.
2. The common denominator for CAD data appears to be IPC-D-356 or GenCAD, with some assemblers CAM systems preferring GenCAD. All systems appear to accept IPC-D-356, because ARE using Gerber and IPC-D-356 is well known and reliable. GenCAD might require review to see whether it is formatted properly. It appears that GenCAD was the favorite, but IPC-D-356 has gained dominance. IPC-D-356 data is the popular format for bare-board testing, so it is likely the best choice of the two for KiCad.
3. Surprisingly enough, turn-key and short production run assembly shops do not particularly want to see assembly drawings. They prefer the silk. Also, they prefer boards with silk (contrary to what many board fabricators would tell you). The reverse appears to be true for full production assembly shops.
4. It appears that turn-key and short production run assembly shops simply use the solder paste information as-is, as some provide recommendations on QFN thermal pad stencil aperture design. This means that it is even more important the proper stencil design is performed by the CAD system.

6.2.1 BOM

The BOM provides the part data necessary for ordering and loading pick-and-place machines in turn-key assembly shops, as well as identifying supply requirements for full production assemblers. When a turn-key assembly shop says "BOM", they do not mean a Bill Of Materials in the traditional sense of manufacture, which is a list of *all* the necessary materials for the product, including copper-clad laminates, copper foil, LPI, etc. *The assembly shop's BOM is actually just a Parts List itemizing the components that are to be attached to the board.*

The BOM tells the assembler what parts are required to be placed (and which are not in a partial-assembly) as well as where they should go. Assemblers can accept a BOM in a number of formats: *.xls*, *.xlsx*, *.csv*, *tab-delimited*. All of these format can be read and created in Microsoft Excel. A number of turn-key

assemblers have slightly different requirements for fields in the BOM, but the following is a common denominator:

1. *Line/Item Number:* this makes it easier to talk on the phone about issues with a particular line. Some require it, some do not.
2. *Quatity Per* The number of instances of a single part. This is really just for checking orders. The quantities are part of the data set.
3. *Reference Designator* This is required for CAM software: it is the only thing that links a part placement on a board to an actual ordered part sitting in a reel or tube. This is the placement identifier.
4. *Part Number* Manufacturer or Supplier part number. This is required for CAM software: it is the only thing that links a part in a reel or tube to a part placement on a board. This is the part identifier.
5. *Part Description:* This is just for clarity, both for ordering and ARE.
6. *Type:* Such as: SMT, Thru-Hole, Fine-Pitch, or BGA/Leadless. This is just to assist with the ordering and ARE process when necessary.

Some order these (4), (5), (2), (3) instead, and skip (1) and (6). Many CAD systems will put (3), (2), value, and then (4). Additional data includes:

7. *Manufacturer:* This is to assist with ordering, when parts are to be ordered by a turn-key assembly shop.
8. *Manufacturer P/N:* This is to assist with ordering, when parts are to be ordered by a turn-key assembly shop.
9. *Vendor:*
10. *Vendor P/N:* Also called Distributor P/N. This is to assist with ordering, when parts are to be ordered by a turn-key assembly shop.

Conspicuous by its absence is a boolean field that says Install/DontInstall. Assemblers want you to highlight these in red instead for partial assembly orders. One of the reasons is that BOM reading software is highly customizable and has not traditionally added predefined fields for partial assembly, meaning that there needs to be strong visual clue to the operator when importing the BOM. If the KiCad user is going to do partial assembly orders, it would be best to import the CSV into Excel and highlight is as appropriate.

Some optional fields include:

11. *Part Marking:*
12. *Package:* The footprint or decal. It would be good if this field followed IPC footprint naming conventions [PCB-2010].
13. *Population Type:* Whether the device is populated. The values of this field should indicate DNS (Do Not Supply), DNI (Do Not Install).
14. *Customer P/N:* Sometimes the customer wishes to assign their own internal part numbers. This part number does not necessarily identify the reel or tube.
15. *Bin Location:* Identifies a supplier bin location.
16. *Stock Code:* Identifies a supplier stock code.
17. *Software Version:* Identifies the software version of software applied to programmable parts that can be ordered with different programmings.
18. *Part Revision:* Identifies the die revision when it is important to function or cost of the component.
19. *URL:*
20. *Cost:* The approximate cost of a component. This can be used as a verification check for ordering. It can also be used to identify critical (expensive) components. Sometimes the overage is also based on the approximate cost of the part.
21. *Height (mils):*
22. *Polarity (Y/N):*
23. *Socketed (Y/N):*

Additional custom fields can be supported by most assembler CAM software, but are not normally used for turn-key orders or short production runs.

6.2.2 Netlist

Many assembler CAM tools can use an IPC-D-356 netlist for ARE (Automatic Reverse Engineering) or verification of the footprints and component placement on a PCB. CAD files can be used for this purpose as well, however, some assemblers are concerned about the IP (Intellectual Property) issues of importing CAD data directly. IPC-D-356 netlists are normally required for bare-board testing and most turn-key and short production run assembly shops also require full bare-board-testing be performing on all boards that they assemble.

The principle data contained in the IPC-D-356 file that applies to ARE (Automatic Reverse Engineering) is the identification of packages and pins and the precise location of each land with reference designator and pin number attached. Using this data, it is quite easy for the CAM program to identify the placement and orientation of each component by reference designator and pin number. The format of the IPC-D-356 file is detailed in *Sec. 7.4(136)* and *Apx. E(151)*.

All current assembler CAM systems accept IPC-D-356. At one time, GenCAD [GenCAD] was preferred by assembly shops and some CAM tools (such as CircuitCAM) recommend exporting GenCAD files in preference to IPC-D-356. Some tools (such as UniCAM and FabMaster) also support GenCAM [GenCAM], the successor format to GenCAD. Interestingly enough, this was also the case for bare-board testing at one point.

To support bare-board testing, assembly shop ARE (Automatic Reverse Engineering) and ICT (In-Circuit Test), pcbnew should export IPC-D-356 files either at one of the revision levels (Lavenir Format 2 (IPC-D-356), Lavenir Format 4 (IPC-D-356/NTD), IPC-D-356A or IPC-D-356B. Some older CAM tools might not support IPC-D-356B. Any of the formats has sufficient information for ARE.

6.2.3 CAD Data

From the perspective of performing ARE and verification, the data included in an IPC-D-356 netlist is sufficient. However, vendors of CAM software prefer to have CAD data for the following reasons:

- It is the richest format and the most likely to contain the most information.
- It always contains sufficient information for ARE: in fact CAM software does not really perform ARE on CAD data, but simply scans the CAD data for the necessary information.
- Native CAD data is the least post-processed data and is therefore less prone to errors in translation.

Because of the foregoing, CAM software vendors prefer CAD data in the following order, where the first listed is the highest preference:

1. Native CAD Formats.
2. Neutral CAD Formats.

3. IPC-D-356 Output.
4. Native CAM Formats.
5. Neutral CAM Formats.

However, assemblers that are worried about the legal stigma associated with providing CAD data, will order their preferences: (3), (5), (4), (2), (1). This means that the least common denominator (at the time of writing) is IPC-D-356. Nevertheless, these various formats and the reasons for their preference are detailed in the paragraphs that follow:

Native CAD Formats: Native CAD formats are the first in preference due to the reasons explained previously. Usually on the ASCII forms, and even the neutral native format (such as Mentor Neutral File) are accepted. The suitability of the data for use with assembler CAM software might still require review. Also, some assemblers do not prefer CAD formats based simply on appearance (legal concerns over IP (Intellectual Property) and NDA (Non-Disclosure Agreements)).

It appears that nobody has written in import wizard for KiCad, and maybe that is not a bad thing: KiCad's file formats are subject to change without notice.

It should be stressed that pcbnew should not simply use the specification of someone else's CAD format here (such as Protel 98E/99, which is a publicly available specification) because the conversion would ruin the desire to get the data directly *from the horse's mouth.*

Neutral CAD Formats: The one neutral CAD format that achieved broad adoption by assemblers and thus their CAM software has been GenCAD [Jones, 2001]. Being a *dead language*, GenCAD v1.4 is well supported neutral CAD format that is a well supported by various CAD and CAM tools. Again, some assemblers do not want neutral CAD data due to appearances (legal concerns mostly). The GenCAD format is detailed in *Sec. 7.9.1(140).*

Most assembler and their CAM software suppliers prefer GenCAD to IPC-D-356 simply because the former is older than the later, and thus—at one point—there was more industry experience with working with GenCAD than IPC-D-356. However, that tide has turned, and from a KiCad development perspective it is far easier to generate a correct IPC-D-356 file than it is to generate a correct GenCAD file. (For example, GenCAD file generation is currently broken.) Nevertheless, should the GenCAD support in KiCad be fixed (and an attempt was made to do so in this development), pcbnew should also provide GenCAD output to be included with the assembly output data.[46]

IPC-D-356 Output: IPC-D-356 was addressed previously under *Sec. 6.2.2(127).* IPC-D-356 is suitable for ARE, does not have the legal stigma associated with GenCAD or native CAD formats, and is quite popular and widely adopted for bare-board testing. The Lavenir Format 2 (IPC-D-356 w/o NTD) version of the netlist likely has the least legal issues, because it is most difficult to recreate a full board design from the data; whereas, when the NTD (Network Trace Data) is included, such as in Lavenir Format 4 (IPC-D-356 w/ NTD), IPC-D-356A or IPC-D-356B, it becomes easier to fully recreate a board design from the data.

As a result, IPC-D-356 is the favoured format of assemblers wishing to avoid the legal stigma of CAD formats.

Native CAM Formats: Native CAM formats are the next in preference, but only when they can be successfully imported and come directly from a native system. Typically support by assembler CAM software for other CAM formats is simply to try to attract customers. The matter is moot because very few CAM formats are publicly available. The only native CAM format of any importance is ODB++, but due to its proprietary nature, it unlikely that KiCad will ever support it.

Neutral CAM Formats: Neutral CAM formats such as IPC CAM formats are—maybe not suprisingly—last on the list. Most of these formats are rather complex and not well supported by CAD tools. They are accepted by production manufacturing facilities and production CAM software. Because they can be interpreted differently by different software, most of these formats are unusable for turn-key or short production run assembly.

Recommendations: When generating assembly outputs, the IPC-D-356 netlist that was provided to the board fabricator for bare-board testing should also be provided to the assembler. Also, when a usable GenCAD export is achieved, this file should also be provided to the assembler where there are no concerns about any of the legal issues associated with CAD data.

6.2.4 Array Data

Where assembly is performed on complete production panels, sub-panels or arrays, the fabrication print must provide information on panelization and arrays (if any). Where panelization was performed by the board fabricator, the details need to be forwarded from the board fabricator to the assembler. As a matter of fact, any other parameters decided by the board fabricator, such as the Tg of the dielectrics used, final design thickness, etc., must be forwarded to the assembler.

- There is a minimum panel size that can be handled by most assemblers. Typical minimums are: 16 sq. in.; 2"×2"; 3.15"×2". PCB sizes smaller than the minimum must be panelized into assembly panels.[47]
- There is a maximum panel size that can be handled by most assemblers. It has to do with the maximum panel size that can be accepted by equipment within the manufacturing line. Panel sizes should be no larger than about 18"×14" (including the 0.200" (3mm) border on the long edges of the panel). This tends to rule out using an entire board production panel for assembly.
- Turn-key assembly shops look for an optimal assembly panel size of about the size of a sheet of letter paper (about 9"×11"). This is sufficient for somewhere between 2 and 6 PCIe add-in cards. For cards with gold fingers, it might be advantageous to combine boards rotated for edge tip plating into the same assembly panel (array). Optimal panelization requires a trade-off between assembly panel (arrays) and production panel board counts and sizes.
- Assembly shops will accept data that is one-up, or for a full array. Typically it must be specified when there is a different between one-up data and an array (that is, the array consists of more than one board, or even of different boards).

46. Note that most turn-key and short production run assembly shops take the attitude that more information is better: that is, more information, even when redundant, simply provides more points of verification and clarification. Full production facilities tend to take the opposite approach: a single source of non-conflicting information is best. This is why full production facilities prefer integrated CAM formats, whether ODB++ or 258X.

47. Note that the board fabricator calls these "arrays", whereas the assembler calls them "panels." The board fabricator uses the term "panel" to refer to a fabrication production panel. To avoid confusion, this document refers to assembly panels as "assembly panels" or "arrays" and refers to fabrication panels as "production panels" or "panels."

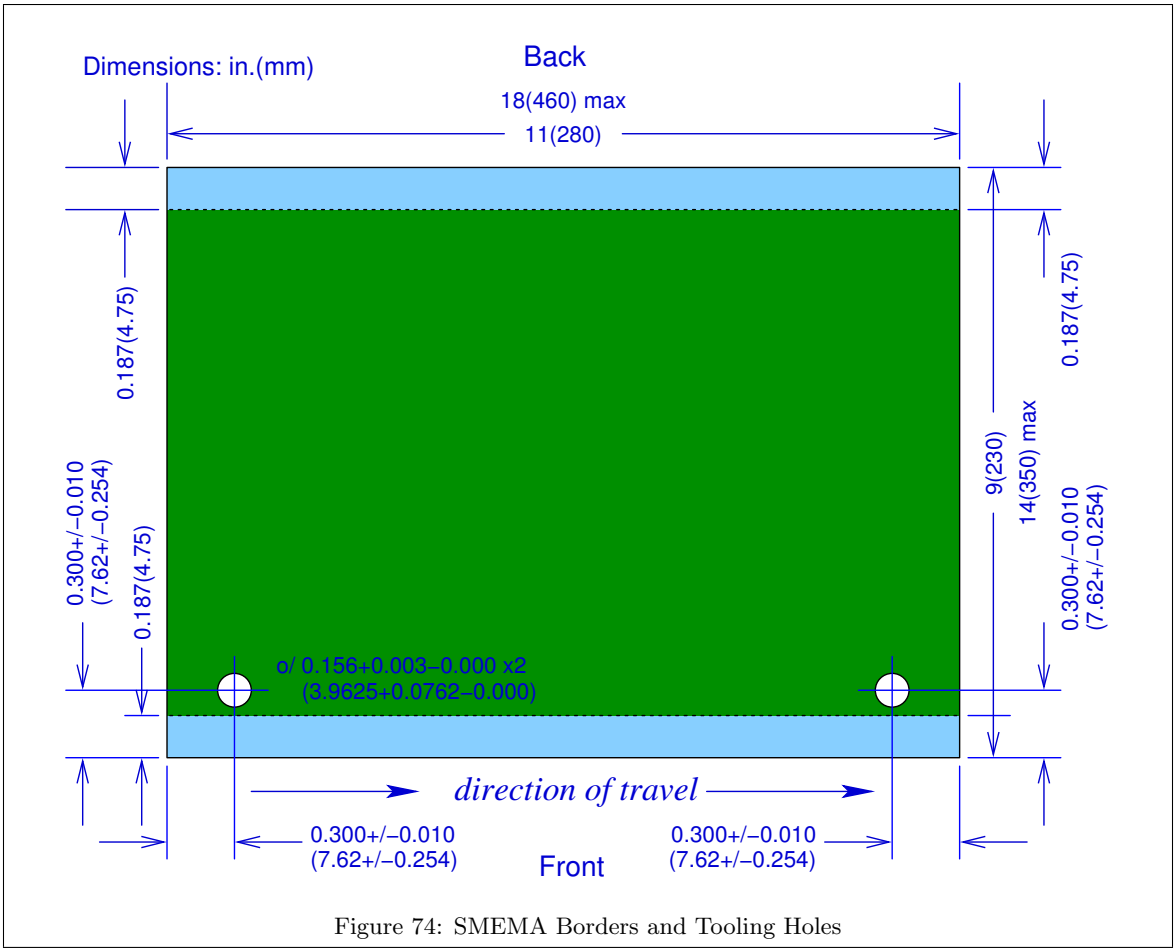

Figure 74: SMEMA Borders and Tooling Holes

- Most turn-key assembly shops require that boards be rectangular, or that they have assembly rails attached to make them rectangular. At least two parallel edges are required to move the assembly panel down the conveyor line. These edges can be added as break-off rails or can be part of the board itself. Some assembly shops require only the two long sides to be parallel. Others require a full rectangular shape. Varying requirements in this regard have to do with the conveyor used to transport the assembly panel down the line, and whether a fixture is attached to each array. For short and full production runs, panel fixtures are avoided where possible. When bare assembly panels are run down a conveyor, the front and back edges of the assembly panel must be longer than the side edges: this is to avoid having the assembly panel jam in the conveyor.
- When an array is required, board must be placed within the array so that no part (that is assembled while the array is intact) overhangs the edge of the array. If a component overlaps the edge of an assembly pane, an assembly rail needs to be provided to remove the overlap.
- Individual boards should also be arranged so that no part (that is assembled while the array is intact) overlaps an adjacent board.
- Arrays and boards within the array must be separately marked with fiducial marks. Each array and board must be marked with at least 3 fiducial marks on opposite corners.
- Standard fiducial marks are used [SMEMA 3.1]. Some assemblers require the fiducial mark to be 0.025”(0.635mm) from any trace.
- The distance from a fiducial mark to the edge of a printed circuit board or fabrication panel shall not be less than the sum of 4.75mm/0.187” (the SMEMA standard transport clearance[SMEMA 1.2]) and the fiducial mark clearance ($2R$ which is typically 0.5mm/0.020”).
- A fiducial mark must be placed within 0.025”/0.635mm of any high-pitch BGA, QFN/QFP or SOP. There must be two fiducial marks on opposite corners for these high-pitch devices. If there is room, the fiducial mark can be placed underneath the device. These requirements are not quite in fitting with the SMEMA (Surface Mount Equipment Manufacturers Association) standard [SMEMA 3.1]. The standard says to place both inside the perimeter of the footprint, or to place one at the centre of the footprint and one elsewhere on the board.
- Usually assemblers require special instructions to separate arrays into individual boards. Often the default procedure is to leave the array together.
- Assembly panels (consisting of one or more boards) must have a border provided for conveyor fingers along the long sides (front and back edges) of the assembly panel. No components are allowed in this border. This standard border is 0.187”/4.75mm [SMEMA 1.2], and the standard border is illustrated in *Fig. 74(129)*.

  Some assemblers require all 4 sides of the assembly panel to have these borders. (Likely so that they can ultimately choose the travel direction of the assembly panel later.) Some assemblers require a 0.180”/3mm border, and others a 0.200”/5mm border. To ensure that a range of assemblers can be used, use the 0.200”/5mm border.
- The SMEMA transport standard[SMEMA 1.2] states that two tooling pins must be on the front edge of the board

(next to the fixed transport rail). A recommended hole size for the tooling pins is 0.156(+0.003,−0.000)"/3.9624-(+0.0762,−0.000)mm. This corresponds to a 5/32", #21, or 4mm drill bit. The distance from the edge is 0.300±0.010"-(7.62±0.254mm). This standard tooling is illustrated in *Fig. 74(129)*.

Some assembly panels require that four 0.158"/4mm tooling holes be provided and that they be positioned 0.197"/5mm in both directions from all corners of the board (meaning that they are 0.118"/3mm from any edge). These assembly panels only use 0.118"/3mm borders. This is rather different from the SMEMA (Surface Mount Equipment Manufacturers Association) standard [SMEMA 1.2].

(Well, ... no.) If boards are given assembly rails and are of a sufficient size, no arrays are required. Where arrays are provided, they are actually a subset of the panel and should be included in the fabrication print.

For short production runs, any PCB smaller than 16 sq. in. or non-standard shape (any shape other than square or rectangle) must be panelized into an assembly panel (board array). The optimum assembly panel (array) size is about 9"×11", or about the size of a sheet of American letter paper. The following are additional requirements:

- Assembly panels should be roughly 9"×11".
- Fiducial markers must be placed on all Panels and boards.
- Boards must have a minimum "keep-out" areas of 0.200" from the edge of the board to the edge of components (when not possible, an assembly rail must be added).
- Individual boards should be spaced within an array so that overhanging components do not overlap adjacent boards.
- Each array and boards must have at least 3 fiducial markers.
- Each fiducial marker should be at least 0.025" away from any solder mask or traces.
- Fiducial markers comprised of round unmasked etched pads 1mm/0.040" in diameter are preferred.

Some other assembler requirements:

- If the PCB is 2×2 or less, they must be in an array/panelized format, or it could increase the assembly cost.
- All PCB must be ordered with electrical testing for all assembly orders.
- For assembly of non-rectangle PCB at least two parallel edges are required. These can be added as break-off rails or can be part of the PCB itself. The same applies to a circular PCB.
- All vias under BGA and leadless type devices are tented with solder mask to ensure there are no solder shorts created during the reflow process.

Some other assembler requirements:

- PCB should be no larger than 18"×14" including a 0.200" border on the long edges of the PCB.
- If PCB is small or cannot provide a 0.200" border, it is recommended to place multiple PCB on a panel. This will enable easier handling of the PCB.
- The minimum panel size is 3.15"×2", including border.
- All panels should have at least 3 fiducial marks near the edges of the panel but at least 0.250" from any edge. You cannot have too many fiducial marks. Also there should be at least one fiducial mark (two recommended) within 0.250" of any fine pitch part (this fiducial can be underneath the part if space is limited).
- Our recommended fiducial is a 0.040" diameter round pad surrounded by a 0.020" radius empty space.
- Because a fiducial is used to provide a frame of reference for the locations of the pads, a drill hole cannot be used as a fiducial as the drilling is a separate process from the etching.
- Remember to leave the fiducial mark off of the paste layer; a fiducial with paste on it is *not* usable.
- A 0.158" (4 mm) tooling hole should be provided 0.197" (5 mm) in both directions from all corners of the board.
- A 3 mm keep-out zone must be provided for conveyor fingers along the long sides of the board.
- Components taller than 0.75" (19 mm) must be kept 5 mm from the long edges of the panel.
- Components cannot be taller than 1.0" (25.4 mm). All components outside these ranges will have to be placed by hand after both sides of the board is complete.
- Use eyelets (rivets) instead of screws wherever possible (heatsinks, TO-220, etc.).

6.2.5 Centroid File

Also called Pick-and-Place File, CPL (Component Placement List), XYRS (XYRS Data), or simply XY data, this file contains information about the placement, orientation and polarity of components.

There is no formal standard for the format of CPL files. There is, however, an industry wide *de facto* format. They are normally formatted as ASCII files with information that is suitable for import into an Ecxel spreadsheet: that is, a general CSV (Comma Separated Values) formatted file. Care should be taken that this information is correct and that all of the necessary information is contained in the single file so that no manual additions to the data set will have to be made.

Assemblers use this file for just about everything. It is used to develop solder paste or adhesive stencils or dispenser programs. An assembler expects to see the correct vendor part number, manufacturer's part number, reference number and footprint. The assembler expects IPC footprints, so use the IPC naming and footprints provided by IPC. The locations of all board and local fiducial marks need to appear in this file.

A portion of the cost and time needed when using a production assembly house may include a hidden 'one time' costs for the hand placing of *all SMT 'parts'*. This is done so the pick-and-place machine can learn the XYRS information. This placement effort is usually based solely on Gerber land patterns. Some fabricators avoid this costs and human intensive operation by getting the information directly from the CAD system output. Every CAD system can generate this information. The XYRS is simply the information needed to place (or find) a component on the assembly. The (XY) information is in Inches, MM, or Mils; and is measured from a given board origin. The (R)otation is in degrees clock-wise (0,45,90,etc.). The (S)ide is usually defaulted to the top side, but can be expresses as Top/Bot or even 0/1. Note that the XY information should represent the centroid of the part as opposed to the Pin 1 location as that is what the pick-and-place machine needs to use.

CPL (Component Placement List) is electronic format must contain the following: X/Y Centroid location, rotation, reference

designator, and board side (top/bottom). Most turn-key assembly shops prefer the CPL in Microsoft Excel or CSV (Comma Separated Values) format.

Centroid files (aka pick-and-place or XY data): are the machine file in ASCII Text format which should include X, Y, Theta, Side of Board (Top or Bottom), and Reference Designator. If you cannot generate this file, some times the assembler can generate the file from the Gerbers at an extra charge, either manually or using a program. Note that quotes around data are not necessary. It is acceptable to have reports for top and bottom sides separately. The preferred unit is inches.

For more information on pick-and-place files, see *Sec. 7.7(138)*.

6.2.6 Pad Master

The pad master is two Gerber or DPF plots, one front one back, that plot all of the pads on the board. This can also be the primary mask plotted 1:1, on, in our case, the probe layer. These plots are identified as pad-master front and back. These should likely be printed on the assembly prints as well. Also, the assembly print should identify the design parameters of the footprints: i.e. whether footprints on the solder-side have been designed for wave soldering or reflow.

Assembler needs pad data, but they do not need traces. Either a pad master needs to be provided or the external copper layers containing components.

6.2.7 Silk Screen

Silk screen artwork should be supplied so that documentation for manual stations, rework or repair can be developed. Silk screen artwork must clearly identify the orientation of polarized components. Some assembly shops require silkscreen: well, they say that they do not require it but they prefer it (the wish): when submitting boards without silkscreen, the assembly shop will need to have some other form of documentation to verify component placement and rotation. Order with boards that have no silkscreen need to be reviewed (the punishment).

This is really required. Note that an assembly print that is aligned with the artwork could also be used (or useful).

6.2.8 Solder Paste

The solder paste artwork should provide solder paste stencil apertures for assembly. Again, these plots are unlikely to be used by the assembler. The assembler might demand IPC footprints for a good reason. Chances are that the assembler will want to develop their own solder paste stencils from the pad master, so that pad master should always be provided. The assembly prints should identify the assumptions for the artwork. Some assembly shops might use a pick-and-place machine that has integrated solder paste and adhesive dispensers. In these cases, the solder paste and adhesive patterns might be programmed against the part number or IPC footprint. Custom footprints might require extensive programming. Note that solder paste dispensers do not work like stencils: that is, they deposit and array of solder paste dots (Hershey's Kisses) of several sizes instead of using apertures. They may be able to accept stencil Gerber data, they may not.

Necessary for stencil generation.

6.2.9 Assembly Drawings

The main assembly drawings, front and back, showing component outlines and possibly also the silk screen legend, board edges, special instructions or handling. Additional detailed drawings can be provided for some mechanicals, such as bezels, and brackets. It is important that the polarity of all polarized components (capacitors, diodes, transformers) are clearly marked by their outlines on the assembly drawing, front and back.

For mixed technology boards with TH pins, press-fit pins, and SMT devices, top and bottom, that all TH pins and press-fit pins are identified.

Special instructions should be supplied with the data package. Often, special instructions will be supplied in an Assembly Instructions document, usually in PDF or notepad depending on the need for visual aids.

Drawings of the finished assembly are requested and will help ensure correct component placement. If not available, documentation must be provided that indicates orientation and polarity of all components, detailing any safety critical components. Sample assembled units, if available, are appreciated as well.

6.2.10 Solder Mask

6.2.11 Drill Data

It is important for assembly shops to be provided the NC drill data for all holes that transfix an external layer and are not filled or tented completely with primary or secondary solder mask. This includes both first-drilled PTH and NTPH as well as second-drilled NPTH and tooling holes.

(Not necessary.)

6.2.12 Additional Data

Perhaps strange enough, most turn-key and short production run assembly shops are *not* interested in the following data (although some full production run assemblers might still be interested in them):

Fabrication Print: As a number of the details of the board required by the assembler might be contained on the fabrication print, the full production run assembler likely needs a copy of the fabrication print to identify these details. Some things include the board surface finish, the general dimensions and outline of the board, the position and geometry of assembly rails, the location of gold fingers, the location of tooling holes.

For turn-key builds, this is not required. Any special assembly instructions that might appear on the fabrication master should be included instead in the special instructions on the assembly print, or in a separate assembly document.

Contacts Artwork: The contacts artwork (for contact finishes) should be provided, perhaps as an assembly drawing detail, particularly when contacts need to be protected from soldering processes for full production runs and the assembler is to apply Kapton tape or some other protection for contacts.

For turn-key and short production runs, any protected contacts need to be protected by the board manufacture. Corfin or Peel-off Mask should be used to protect contacts and left in place for removal after assembly. Otherwise, it needs to be addressed with special instructions on the assembly drawings and detail.

Finish Artwork: The finish artwork (for secondary finishes) should be provided, particularly when these finishes need to be protected from soldering processes.

(Well, no, not really, unless detail for special assembly instructions. Any protected secondary finish needs to be protected by the board manufacture (unless, maybe, it is just simply edge finger taping).)

Secondary Mask Artwork: The secondary mask (via capping) artwork should be provided so that the assembler can evaluate the suitability of specific solder approaches and assembly line configurations, as well as ICT fixture development.

(Well, no, some fabricators say don't send this data. It could be used for review as some assemblers have requirements for via capping under BGA fields, QFN and the like.)

Primary Mask Artwork: The primary mask artwork should be provided so that the assembler can evaluate the suitability of specific soldering approaches and assembly line configurations. Problem areas might be identified on an assembly detail sheet.

(Well, no, some fabricators say don't send this data.)

Silk Screen Artwork:

Peel-off Mask Artwork: Any removable mask or peel-off mask that is left on the board after fabrication for assembly purposes should be identified with the peel-off mask artwork plots. Chances are an assembly detail sheet should be provided as well indicating the purpose of these masks (e.g. to protect gold finger contacts from wave soldering).

(Not necessary.)

Solder Paste Artwork:

Adhesive Artwork: Although we have an adhesive layer and we can generate plots of glue dots or stencil openings on the adhesive layer, most assemblers do not use this information. They use the pick-and-place files and footprints listed there to generate glue dot programs where necessary. For boards designed for wave soldering, all solder-side components will need to be glued. For simultaneous double-sided reflow, only components over a certain mass will need to be glued on the solder-side. An integrated pick-and-place machine will identify the components that need adhesive by their part number and the glue dot pattern will be programmed against the part. Even when the adhesive and pick-and-place machine is not integrated, the pick-and-place files are used to program adhesive dispenser by using the part numbers and/or footprints identified in the pick-and-place file. For more information on adhesive layers, see *Sec. 5.15(96)*.

(Not necessary.)

NC Drill Data:

Pick and Place Files:

Net List: For In Circuit Tests, the netlist (IPC-D-356) must be provided, including any defined test points. Also, the probe layer defining test points that can be accessed should also be provided. This file can be used (ala WISE software) to automagically reverse engineer every footprint on the board. The reason is that the IPC-D-356 file contains reference designator and exact pin number and location on every package. This can easily be compared with flashed apertures on the Gerber files to automatically identify every footprint.

Note that CircuitCAM considers IPC-D-356 to be a sufficient CAM file format. Without these layers, an unclipped legend layer is required.

CAD Data: For short to full-scale production runs, assembly shops required CAD data. Many of the ASCII formats provided by CAD tools are supported. The common denominator for most of these systems appears to be GenCAD???!. [GenCAD].

ODB++ or CAD data is also acceptable pending data review in lieu of Gerber data and CPL (Component Placement List).

Program Files: Any program files necessary for programming EEPROM, FPGA, NVRAM, SPI-based EEPROM chips, etc., should be provided. Also any test programs for boundary scan or JTAG should also be provided.

BDSI Files: BSDI definition files for devices that are to be boundary scanned need to be provided.

JTAG Files: JTAG files for programming or test need to be provided.

SMBus, I2C, SPI Files: Any programming files for SMBus, I2C, SPI files or images to be transferred need to be provided.

Gerbers: Send all Gerber files generated by the CAD program in RS-274X format (as if you were sending them for PCB fabrication). At a minimum, assemblers need silkscreen, copper (trace) and solder paste (or stencil) layers for assembly.

6.3 CAM Outputs

Newer CAM formats include:

GenCAD. GenCAD [GenCAD] is an EBNF formatted CAM format that was a precursor to the GenCAM format. The GenCAD format is described in *Sec.* **??***(??)*.

GenCAM. GenCAM [GenCAM] is an EBNF formatted CAM format with strong similarities to GenCAD. The GenCAM format is described in *Sec. 7.10.1(140)*.

GenCAM-XML. GenCAM-XML [GenX] (or sometimes simply referred to as GenX) is an XML formatted CAM format largely equivalent to GenCAM, but formatted with XML instead of EBNF. The GenX format is described in *Sec. 7.10.2(141)*.

ODB++. ODB++ is a proprietary CAM database format developed by Valor[48] that includes a directory of primarily XML coded files. The ODB++ format is described in *Sec. 7.10.3(141)*.

258X. 258X (IPC-2581) [258X] is a vendor-neutral XML-formatted CAM database format that draw from both GenX and ODB++. The 258X format is described in *Sec. 7.10.4(141)*.

Note that all CAM formats provide the entire database. *ODB++* keeps the database in a zipped directory structure; however, all of the others provide the entire database in a single file.

Some of the output formats for pcbnew have been redone. Of specific concern was the Gerber and NC drill file formats. Other output formats have been added, such as the GenCAM [GenCAM] and other formats in the IPC family. This section describes the changes made to existing formats and the additional formats now supported by pcbnew.

48. See *Apx. D(150)*.

7 File Formatting

7.1 Gerber

Fabrication data requirements for Gerber artwork. Following are some points and requirements concerning Gerber data:

- Make sure the artwork layers and drilling files are created using the same format (If you don't know what to use, you probably can't go wrong specifying absolute, no zero suppression and either 2:6 coordinates in inches or 3:3 if the units are in metric.
- Create a separate layer for the board outline, or plot it only for the soldermask layers. You don't wan etched copper board outlines on every layer, or on the finished silkscreen.
- Unless you really know what you are doing, generate solermask openings 1:1 with the pad sizes (except for fiducial marks). Unless the designer is known to ave paid careful attention to expanding the soldermask apertures appropriately, the manufacturer would rather do it using a known starting place. Starting with mask that matches the pads makes the modification process much easier for the manufacturer.
- If you are okay with removing non-functional pads, leave them in the design so they can be used for drill/annular ring checking, and the CAM operator can remove them afterwards.
- A unique aperture size used to draw planes, which is a different size than any of the thermal spoke widths, can be very helpful to CAM.
- Make sure to differentiate which holes are plated and which are unplated.
- If you are leaving it up to CAM to create an assembly array or pallet with breakaway tabs, make sure you leave spaces free of traces and components for the tabs.
- README files are encouraged! Make sure to list any known netlist discrepancies.
- Never try to perform your own etch compensation!
- Review your final output before you send it out! You can find several free tools by searching the internet for "Gerber Viewer".

Extraneous files. Do not provide any additional files such as original CAD data, Graphicode GWK files, PDF files, Word files (doc), Microsoft Excel files (xls), part lists, placement and assembly information, etc.

Verified plots. Check Gerber and Excellon data with a Gerber viewer before sending to a fabricator.

File naming. Use clear and easy to understand file naming and try to avoid long filenames. Make sure that it is easy to determine the layer function from the filename.

Scaling. Do not scale Gerber or Excellon data. All data provided must use 1:1 (100%) scale.

Zero-sized apertures. Ensure that Gerber files do not contain apertures with a zero-size (size = 0.00mm) and that Excellon data does not have zero-sized tools (size = 0.00mm). This can result from round-off errors to the resolution size. When the size of an aperture is beneath the minimum size representable using the file format, the resulting size can be zero. Also, some CAD system may have seen fit to use a zero aperture for describing contours. See also *"Flash pads, draw lines"* below.

Zero-length draw segments. Due to round-off errors on extremely small lines (such as those making up fancy-drawn characters), it is possible that the start and end position are identical. Some fabricators hate these because it is difficult to determine whether the segment is a mistakenly draw (instead of flashed) pad, or whether it is a zero-length line segment. See also *"Flash pads, draw lines"* below.

Zero-length arcs. Due to round-off errors on extremely small arcs, it is possible that the start position and end position of the arc will wind up being identical. This is not good in G75 360-degree circular interpolation mode because an arc that has the same beginning and end point is a full circle instead of a point. Small arcs of this kind resulting in zero-length arcs should either be rendered using G74 90-degree circular interpolation, or not plotted at all.

Plot offset. Use the same offset for all Gerber layers and the Excellon drill data. Preferrably use no offset at all.

Units. Use the same units (mm or inch) in all Gerber and Excellon output files as in the CAD PCB design software. This will eliminate conversion and rounding errors.

Resolution (grid). Use the same resolution (grid) for Gerber and Excellon data to allow a perfect match. Also make sure that the resolution (grid) used for Gerber and Excellon is at least a factor of 10 better than the resolution (grid) used in the CAD PCB design software. Because of round-off errors, fabricators recommend that the Gerber and Excellon data be presented at a resolution that is 10 times greater than the resolution of the CAD system. See also *"Zero-sized apertures," "Zero-length draw segments," and "Zero-length arcs,"* above.

Mirroring. All data should be supplied as seen from the top to bottom through the PCB. Do not mirror (or reflect) any data layer – image or drill.

Flash pads, draw lines. Gerber files should only flash (and not draw) pads. Gerber files should only draw (and not flash) lines. See also *"Zero-length draw segments"*, above.

Extraneous apertures. The aperture list should only contain apertures used in the design.

Special use apertures. A separate aperture should be used for drawing impedance-controlled traces so that they can easily be identified (by D-code).

Board outline in plots. Include the board outline in all layers. This will help the fabricator properly align all layers in case of an offset problem. Also include the board outline in a separate Gerber file.

Octagons. There was a change in Rev D that altered the definition of how octagons were sized and positioned. Prior to Revision D octagons where measured across the flats and two vertices were on the Y-axis. In Revision D and beyond, octagons were measured across the vertices and two vertices were on the X-axis. I don't know whether this was an error in only the specification, or whether the behaviour of true Gerber photoplotters changed.

Arcs. pcbnew historically output a sequence of small line segments to represent an arc. This is not plot-and-go as arced segments should be represented as arcs so as to not fool the import functions of CAM systems.

Zones. pcbnew historically painted zones with line segments. This is not plot-and-go. Many CAM systems attempt to convert painted zones back into contours. Zones should be represented as contours so as to not fool the import functions of CAM systems. (And many fabricators express that

contours are preferred to painting; also, specifying the minimum overlap and variation in line width permitted when painting zones against recommendations.)

7.1.1 RS274D (Gerber)

pcbnew historically did not support RS-274D (Gerber) outputs, only RS-274X (X-Gerber). There are several reasons for not supporting RS-274D outputs:

- The form is archaic.
- Although some fabrications say that they will *accept* RS-274D Gerbers, they all say that they *prefer* RS-274X.
- The actual Gerber RS-247D specifications are lost to the web. Although hardcopies might exist somewhere, the old Gerber format manuals are either lost to the web or were maybe never generated in electronic format. Yes it is that old. The wayback machine doesn't go earlier than RS-274X Rev. B.
- Most CAM systems that say that they support RS-274D Gerbers probably don't. There are archaic commands in early RS-274D, such as parabolic and cubic interpolations that are probably not supported on current tools correctly.
- The format was not so much a standard. Many photoplotter implemented the format, but, of course, each modified and embellished the language for thier particular plotter (even by model).
- Many of the commands were purposeful only for controlling aniquated photoplotters, and placed restrictions on coding that are no longer pertinent to photoploters made in the last two decades or so.
- The format cannot be read and verified properly by various Gerber viewers.
- It is really not worth the effort, because the use of RS-274X is so much more widespread.

Of these reasons, the most significant one is that the format is not electronically available on the web.

Therefore, for pcbnew, an approximation to the RS-274D format has been created using the following principles:

- RS-274D is plotted identical to RS-247X with the following exceptions:
  - — No *mass parameters* are placed in the file.
  - — Apertures are described in a separate text file.
  - — Aperture reports are the same.
  - — 360° circular interpolation is never used.
  - — Zones are painted with line segments and not polygons.
  - — Complex apertures are "drawn" using standard apertures.
- Aperture numbers, ranges, limits, wheels, etc. are for raster or laser plotters and not antiques. That is:
  - — The aperture list is simply removed from the plot file and placed in the separate *.rep file.
  - — Aperture number are not restricted to 28, nor 255.

These differences are quite simply not enough to make the result usable on an antique photoplotter, but, *so what*?

7.1.2 RS274X (X-Gerber)

pcbnew previously supported plots used RS274X; however, the precision, format, coding, arc representation, etc. were extermely limited.

> *Gerber data* is a simple, generic means of transferring printed circuit board information to a wide variety of devices that convert the electronic PCB data to artwork produced by a photoplotter. [Almost] every PCB CAD system generates Gerber data because all photoplotters [used to] read it. It is a software structure consisting of X, Y coordinates supplemented by commands that define where the PCB image starts, what shape it will take, and where it ends. In addition to the coordinates, Gerber data contains aperture information, which defines the shapes and sizes of lines, holes, and other features.
>
> *Gerber Format*, which is the format in which Gerber data is expressed, actually [*sic*] is a family of data formats that are subsets of EIA Standard RS-274D. *Extended Gerber Format*, which is also called RS-274X, provides enhancements that handle polygon fill codes, positive/negative image compositing, and custom apertures, and other features. RS-274X also encapsulates the aperture list in the header of the Gerber data file and therefore allows files to pass from one system to another without the need to re-input the aperture table. RS-274X produces a variety of Gerber data called *X data*.
>
> RS-274X is a *superset* of the EIA Standard RS-274D format. RS-274X supports some of the parameter data codes (G codes) and aperture codes (D codes) contained in RS-274D, as well as codes referred to as *mass parameters*. Mass parameters are plot parameters that define characteristics that can affect an entire plot, or only specific parts of the plot, called *layers*. Mass parameters extend the capabilities of Gerber Format. Their presence makes the Gerber data *X data*.
>
> RS-274X is maintained by *Gerber Systems Corporation (GS)*, a leading supplier of CAD/CAM systems, large-area plotting systems, and precision cutting systems since 1965.[49]

There are currently many versions of RS-274X. Fortunately, the differences between the revisions are extremely minor and has to do more with companies changing hands and being renamed than any substantive change to the specification. Several revisions of which I have copies are as follows:

- [RS-274X/C] RS-274X Revision C, September 21, 1998.

  This file was downloaded from *http://www.photoplotstore.-com/pages/rs274xc.pdf* on September 6, 2010. Part Number 414 100 014 C, but strangely enough, the document properties say Revision B. Is it possible that this document was simply rereleased by Barco?
- [RS-274X/D] RS-274X Revision D, March 1, 2001.

  Strangely this version's copyright notice is dated 1998. This version removes the line "RS-247X is maintained by Gerber Systems Corporation (GS), a leading supplier of CAD/CAM systems, large-area plotting systems, and precision cutting systems since 1965." from Revision C. In several places *Barco Gerber Systems Corporation* is changed to *Barco Graphics ETS North America*. There is a difference on page 22, figure 5, polygons. The $5 parameter is drawn in Rev C as the distance between the centerpoints of two opposite sides of the polygon; Rev D draws the same $5 parameter as the distance between two opposite vertices. Also, the center of a face is shown on the x-axis in Rec C, whereas the first vertex is shown on the x-axis in Rev D. Other than that, there are no substantive changes. I have seen posts on EDA

49. Introduction from RS-274X Revision B.[RS-274X/C].

mailing lists where old hats complain about one or the other form of the polygon.

7.2 NC Drill and Route Files

Fabrication data requirements for NC drill files. Fabricators impose some requirements on NC drill files that were not previously met by pcbnew as follows:

1. One of the most honerous requirements is that a single tool list be used for the entire set of NC drill files. Using a single (unified) tool list permits the use of a single, unified `nc_tools.txt` file and drill report. It also (and this is the important point) permits the merging and splitting of NC drill files without needing to reassign tools.
2. Fabricators also require a separate NC drill file for NPTH, whether or not they are created using a primary or secondary tooling, that was not historically generated.
3. Some fabricators do not care whether NPTH and PTH are combined into the same file, provided that a separate tool number is used for PTH and NPTH of the same size.
4. For depth-controlled-drilling, fabricators require a separate NC drill file for each depth and each sub-laminate drilled.
5. For back-drilling, fabricators require a separate NC drill file for each depth back-drilled.
6. For through-drilling of sublaminated, fabricators require a separate NC drill file for the layer-pair associated with each sub-laminate.
7. For drill bit selection, fabricators require the identification of holes as PTH, NPTH or VIA, so that appropriate plating compensation and tolerances can be applied.
8. Fabricators often also require the identification of any press-fit PTH, because, depending on the finish, press-fit PTH might require special treatment.
9. Fabricators require that the tolerance of all holes be specified.
10. Some low-cost fabricators require that drill sizes correspond to a standard drill rack.
11. Many fabricators charge according to the number of drill bits used. It behoves the designer to limit the number of drill bits used.
12. Many fabricators charge premiums according to the drilling density (i.e. number of hits per square inch).
13. Some (a few) fabricators expect or accept NC drill data in Gerber or DPF format! That is, the drill data and neither drill maps nor reports.
14. NC drill data must be on the same grid, with the same units, and with the same precision as the Gerber or DPF data. NC drill data must be plotted with the same offset as Gerber or DPF data.
15. Some fabricators *prefer* EXCELLON-1 drill format over EXCELLON-2.
16. Data should be in absolute steps. (This makes rearranging the file easier and less error prone).
17. Data formatted in ASCII. (What else is there? EBCDIC? EIA-244?)
18. *Some fabricators require a <LR><CR>, that is x0D x0A at the end of each line. Is this really necessary? Perhaps we should add an option to the Gerbers and NC drill files to select the line termination.*
19. Metric in the format 0000.00? Why? It seems that some fabricators want to do their pannelization in the same units as the one-up data. Imperial in the format 00.0000? Well, ok.
20. *Hole sizes are to be arranged from **smallest** to **largest** and labelled T01, T02, T03, etc.*
21. *Non-plated holes, including snap holes, are to be assigned to different tools than plated holes of the same diameter.*
22. *Zero datum on the inside bottom left hand corner mark, to make all coordinates positive.*
23. A detailed drawing, showing all hole sizes, must eb supplied for each board. Drawing should also show position of datum point toegher with XY directions.
24. A drill file should never have mutliple hits on the same position.
25. The drill file should not place holes too close together. A minimum web of dielectric must be maintained between drilled holes to minimize tool breakage. In fact DRC should check that this minimum is met. One of the easiest places to violate this rule is in the drilling of perforations for breakaway tabs. Bear in mind that the required design web increases when holes are second-drilled because of the poorer positional tolerance of second-drilled holes to first drilled holes and the possibility of the worst-case addition of positional tolerances.
26. No canned cycles or step and repeats.
27. Holes that are filled require a Gerber or DPF plot where the aperture is the same size as the FHS (Finished Hole Size). This image will be used to generate hole-filling stencils.
28. Holes that are capped require a separate Gerber for each capping side where the apertures are either the same size as the FHS or the pad.

Historically, pcbnew did not address any of these requirements.

7.2.1 EIA 244

7.2.2 Sieb & Meyer

The Excellon and Sieb & Meyer drill formats are designed to drive CNC drilling and routing machines. They are broadly similar, differing only in minor details.

Sieb and Meyer is another drilling and milling machine programming language developed by Sieb & Meyer for their CNC machines. It is seldom used for PCB fab submissions because it is a binary file, and has to be translated to the specific Sieb and Meyer machine in use. Even companies who use S&M machines prefer to get the Excellon ASCII file and have their software convert it.

7.2.3 EXCELLON-1

pcbnew did not previously support EXCELLON-1, only EXCELLON-2, and even thn only the drill press portion of the specification.

7.2.4 EXCELLON-2

EXCELLON-2 was standardized in the IPC specification IPC-NC-349.

7.2.5 IPC-NC-349

EXCELLON-2 was standardized in the IPC specification IPC-NC-349. IPC-NC-349 provides for a subset of Excellon defaults. The following codes are supported:

Table 12: EXCELLON-1 Control Codes

EXCELLON-2 Code	EXCELLON-1 Code
G05	G81
M00	M02
M01	M24
M02	M26
M06	M01
M08	M27
M09	M00
M02X#Y#M70	M26X#Y#M23
M72	M70
M02X#Y#M80	M26X#Y#M21
M02X#Y#M90	M26X#Y#M22
R#M02	R#M26

7.3 Barco DPF (Dynamic Process Format)

DPF standards for "*Dynamic Process Format*" and is a file format that was used by *Barco* to represent layer information in printed circuit boards in its CAM database format.[50]. The format describes the image of the layer, including pads, tracks, holes, power and ground planes, and also electrical netlist information as well as additional free-form product information represented with user-assignable attributes. [DPFv5]. DPF was designed specially for the Electronics Manufacturing industry, in contrast to Gerber's RS274D and RS274X formats, which were based on the RS274D G-Code used by CNC mills and lathes. Nevertheless, Barco's DPF supports a similar set of Gerber RS274X extensions, such as: embedded aperture definitions, reverse objects, contour for outline descriptions, and block apertures to represent step-and-repeat items. [DPFv7]

There are currently many versions of DPF. Fortunately, the difference between the latter revisions are extremely minor and has to do more with companies changing hands and being renamed than any substantive change to the specification. Several revisions of which I have copies are as follows:

- [DPFv5] DPF Version 5, November 1998.
- [DPFv7] DPF Version 7, March 2009.[51]

Many CAM tools (and therefore fabricators) can import and export Barco/Ucamco DPF format in addition to Gerber. Barco DPF format has the advantage that it is directed at raster scanners rather than the traditional Gerber machine. It can accept negative polarity overlays (i.e. REVERSE apertures). [DPFv7] This allows features such as clearance holes to be added to the end of a plot (like postcript or clear layers in Gerber). HPGL and DXF file formats are incapable of this [DXF].[52] The DPF format can include netlist information directly in the file format. This makes Barco DPF equivalent to Gerber plus IPC-D-356 and NTD or IPC-D-356A, however, the information on a layer-by-layer basis is contained in a single file, resulting is less data mismatching.

Some of the advantages of the Barco DPF format are:

1. File contains netlist information.

Some of the disadvantages of the Barco DPF format are:

1. It is still necessary to have a file per layer.

7.4 IPC Netlist

Many CAM tools (and therefore fabricators) can import IPC-D-356/NTD, IPC-D-356A or IPC-D-356B netlist data for feeding information to flying probe bare-board testers as well as for creating test harnesses for bed-of-nails testing, accomodating board repair, and verifying artwork.

IPC Netlist files are only necessary when RS-274D, RS-274X, IPC-D-350, HPGL or DXF formats are used for artwork. This is because these formats do not support embedded netlist information. As *RS274X* (Geber) does not support netlists, many fabicators list Gerber + IPC-D-356 + NTD as the minimum database requirements.

All recent versions of DPF and all CAM formats (GenCAD, GenCAM, GenX and 258X) support embedded netlist information.[53] Although *DPF* [DPFv7] and *GenCAM* [GenCAM] serves this purpose as well; however, it never hurts to have many widely adopted formats for output.

There was sufficient information on the web about the IPC-D-356 and NTD file format and the precursors to IPC-D-356A to be able to write an export function without purchasing documents. IPC-D-356B has very little freely available information (IPC requires that documents be purchased). An option was added to the fabrication outputs menu to generate extended IPC-D-356, NTD files, IPC-D-356A and IPC-D-356B. I purchased a copy of the IPC-D-356, IPC-D-356A and IPC-D-356B specifications. Fortunately they only cost about $50.00 each.

There are currently three (or four) revisions of IPC-D-356:

- IPC-D-356, March 1992
- IPC-D-356 + NTD (Network Trace Data)
- IPC-D-356A (Revision A), January 1998
- IPC-D-356B (Revision B), October 2002

I do not have copies of these documents, but from what I can tell, IPC-D-356 is the original file format; [54] IPC-D-356 + NTD is an *ad hoc* extension to the IPC-D-356 standard. IPC-D-356A is the backward compatible version of the IPC-D-356 format with the additional NTD information embedded. IPC-D-356B is not backward compatible with IPC-D-356A nor IPC-D-356 and has its own format. The file formatting is similar in that the VER parameter can be used to distinguish between an IPC-D-356, IPC-D-356A and IPC-D-356B file.

These standards only cost about $50.00 each from authorized distributors of IPC documents.

7.4.1 IPC-D-356

There are 4 versions of IPC-D-356:

- Lavenir Format 2 (IPC-D-356).
- Lavenir Format 4 (IPC-D-356 and NTD)
- IPC-D-356A.
- IPC-D-356B.

IPC-D-356 (Netlist). This file is used to export the IPC-D-356 netlist and NTD (network trace data) file for Gerber verification, bare-board testing and repair.

7.4.2 IPC-D-356A

As with 356, 356A carries information required for electrical test:

- Location, size, and accessibility of pads and holes on the board.
- Netlist information, showing what pads and holes should be on the same net.
- Header information giving the name and revision of the board, and so forth.

50. Note that Barco NV, is the company that purchased Gerber NV

51. Yes, 2009. *Barco* and the *DPF* specification are now part of "The Evil Empire" (see *Apx. D(150)*).

52. FIXME: this is not true: DXF can create masks or clipping.

53. In fact, GenCAD, which mogrified into GenCAM and GenX was developed by Mitron which was a developer of board tester software. Mitron was bought and made a subsidiary of GenRad, itself later bought by Teradyne in 2001, all whose focus over the years have included board testers.

54. Note that the freely available IPC-2515A [BDTST] has an example of a 1992 revision IPC-D-356 file.

Table 13: EXCELLON-2 Control Codes

Code	Description
M48	Beginning of a Part Program Header
M95	End of a Part Program Header
%	Rewind Stop
A#	Arc Radius
B#	Retract Rate
C#	Tool Diameter
F#	Table Feed Rate; Z-axis Infeed Rate
G00X#Y#	Route Mode
G01	Linear (Straight Line) Mode
G02	Circular CW Mode
G03	Circular CCW Mode
G04	X# Variable Dwell
G05	Drill Mode
G07	Override current tool feed or speed
G32X#Y#A#	Routed Circle Canned Cycle (CW)
G33X#Y#A#	Routed Circle Canned Cycle (CCW)
G34,#(#)	Select Vision Tool
G35(X#Y#)	Single Point Vision Offset (Abs)
G36(X#Y#)	Multipoint Vision Translation (Abs)
G37	Cancel Vision Translation Offset (from G35 or G36)
G38(X#Y#)	Vision Corrected Single Hole Drilling (Abs)
G39(X#Y#)	Vision System Autocalibration
G40	Cutter Compensation Off
G41	Cutter Compensation Left
G42	Cutter Compensation Right
G45(X#Y#)	Single Point Vision Offset (to G35 or G36)
G46(X#Y#)	Multipoint Vision Translation (to G35 or G36)
G47	Cancel Vision Translation (from G45 or G46)
G48(X#Y#)	Vision Corrected Single Hole Drilling (to G35 or G36)
G82(G81)	Dual In Line Package
G83	Eight Pin L Pack
G84	Circle
G85	Slot
G87	Routed Step Slot Canned Cycle
G90	Absolute Mode
G91	Incremental Input Mode
G93X#Y#	Zero Set
H#	Maximum hit count
I#J#	Arc Center Offset
M00(X#Y#)	End of Program - No Rewind
M01	End of Pattern
M02X#Y#	Repeat Pattern Offset
M06(X#Y#)	Optional Stop
M08	End of Step and Repeat
M09(X#Y#)	Stop for Inspection
M14	Z Axis Route Position With Depth Controlled Contouring
M15	Z Axis Route Position
M16	Retract With Clamping
M17	Retract Without Clamping
M18	Command tool tip check
M25	Beginning of Pattern
M30(X#Y#)	End of Program Rewind
M45,long message\	Long Operator message on multiple part program lines
M47,text	Operator Message
M50,#	Vision Step and Repeat Pattern Start
M51,#	Vision Step and Repeat Rewind
M52(#)	Vision Step and Repeat Offset Counter Control
M02XYM70	Swap Axes
M60	Reference Scaling enable
M61	Reference Scaling disable
M62	Turn on peck drilling
M63	Turn off peck drilling
M71	Metric Measuring Mode
M72	Inch Measuring Mode
M02XYM80	Mirror Image X Axis
M02XYM90	Mirror Image Y Axis
M97,text	Canned Text
M98,text	Canned Text
M99,subprogram	User Defined Stored Pattern
P#X#(Y#)	Repeat Stored Pattern
R#M02X#Y#	Repeat Pattern (S&R)
R#(X#Y#)	Repeat Hole
S#	Spindle RPM
T#	Tool Selection; Cutter Index
Z+# or Z-#	Depth Offset
%	Beginning of Pattern (see M25 command)
/	Block Delete

Table 14: IPC-NC-349 Control Codes

Code	Description
%	Rewind and stop
X#Y#	Moved and drill
T#	Tool selection
M30	End of program
M00	End of program
M25	Beginning of pattern
M31	Beginning of pattern
M01	End of pattern
M02 X#Y#	Repeat pattern
R#M02X#Y#	Multiple repeat pattern
M02 X#Y# M70	Swap axis
M02 X#Y# M80	Mirror image X axis
M02 X#Y# M90	Mirror image Y axis
M08	End of step and repeat
N#	Block sequence number
/	Block delete
R#X#Y#	Repeat hole
G05, G81	Select drill mode
G04 X#	Vairable dwell (ignored)
G90	Absolute mode
G91	Incremental mode
G92 X#Y#	Set zero
G93 X#Y#	Set zero
M48	program header to first %
M47	Operator message CRT display
M71	Metric mode
M72	English-imperial mode
Snn	Spindle speed (rpm)
Fnn	Z axis feed speed (rpm)

IPC-D-356A adds considerable information:

- Location and width of conductors.
- List of nets which are adjacent to each other. (This shows where shorts are likely.)
- Actual test probe locations, as opposed to just the location of the pad.
- The size, shape, and location of the soldermask openings, independent of the pad.
- Support for blind and buried vias.
- Longer net names (up to 57 characters) are allowed.
- Support for buried resistance.
- A variety of non-test features, such as fail marks.
- The board outline.
- Support for high voltage tests.
- Support for impedance tests.
- Support for stepped images.

IPC-D-356A also changes the handling of some information from 356:

- In 356A, there are stringent requirements for the order of data in the file. IPC-D-356 does not have these requirements.
- IPC-D-356 supports an optional "*variable record*" format. Support for this format has been dropped in 356A.
- IPC-D-356 uses fixed-column records. Trailing spaces must be included in the file. IPC-D-356A allows trailing spaces to be left out of the record.

- IPC-D-356 was unclear on how records should be terminated, resulting in a variety of approaches by file writers. IPC-D-356A mandates that records be terminated by a CR/LF pair.

The result is that 356A files are much more powerful than 356 files, but are easier to read than 356 files. IPC-D-356A files are in many respects easier to write than 356 files, although the sorting requirements of 356A may add complexity to a writer. On balance, however, 356A files are also easier to write than 356 files.

7.5 IPC Artwork

7.5.1 IPC-D-350 (Artwork)

IPC-D-350 is the IPC's idea of Gerber or DPF files. It is a precursor to GenCAM (IPC-2511A). It apparently never caught on. I do not know whether it is based on Gerber or DPF.

7.6 RS-274D G-Code

RS-247D G-Code is often used to control CNC mills and drilling machines, particularly where engravers are used to mill traces and drill holes. Programs are available (and rather easy to write)

7.7 XYRS Formats:

It is strange, perhaps, but there is no standard pick-and-place file format. The closest thing to a standard is a CSV (Comma Separated Value) file format suitable for import or export from popular spreadsheet programs.

The first line of the file lists the field names. This is not, strictly speaking, part of the CSV format, but it is a common convention, which is dnoe to make the data easier to import into other programs.

Note that text fields are in quotes, whereas numeric fields are not. One reason for this is that a text field might contain a comma. Another reason is that by putting all the fields in quotes it tells the importing program that the field is to be treated "as-is", without alteration. Some programs honour this convention, while others do not. (Microsoft Excel, for example, has a tendency to be terribly clever about the fields it imports.)

Note also that quotes within a quoted field must be represented as doubled-up quotes, ala C-string style. When quotes appear in a quoted field, the quotes are doubled to indicate that they do not mark the end of the field. When the CSV field is read, the double-up quotes should be converted back to single instances. Be warned, though, that even though this is the convention, not all programs do this.

You will frequently see CSV files in which the text fields are not surrounded by quotes. In such a case the programmer has decided that it is impossible for the text to contain a comma.

You will sometimes see CSV files in which the fields are separated by semicolons instead of commas. The use of the semicolon in CSV files is particularly prevalent in Europe, since they may use a comma as a decimal separator. A North American might write 3.14159 while a European might write 3,14159 to represent the same value.

An alternative to the comma-separated-value file is the tab-delimited file. This puts a tab character (character value: decimal 9, hex 0x09) between each field. The advantage of tab-delimiting is that you do not have to worry about commas or quotes. (Tabs almost never appear in a human-readable (plain text) data field, except perhaps due to a data-entry mistake that wasn't filtered out.) The disadvantage of delimiting this way is that the tab character is treated differently by different programs. A CSV file can be loaded into any text editor program, and it will be obvious how each record (one line of fields) is divided. But different text editors treat tabs in different ways. In most cases the tab will look exactly like one or more spaces.

Table 15: Pick-and-Place Fields

Field	Code	Description
LocationX	X	X-coordinate of the centroid, usually in inches accurate to a decimil; e.g., "512.21."
LocationY	Y	Y-coordinate of the centroid, usually in inches accurate to a decimil; e.g., "543.75."
Rotation	R	Rotation of the part, in degrees, usually accurate to a tenth of a degree, right-hand cartesian coordinate system; e.g., "90.0."
Layer	S	Side of the board; e.g., "Top" or "Bottom."
RefDes	D	Reference designator of the component: e.g., "C10", "R235", or "U8."
Footprint	F	IPC Footprint name: e.g., "CAPC1005N."
Part No.	P	Part number.

Each record is one line, but a record separator may consist of a line-feed, or a carriage-return and line-feed pair. Also, fields may contain embedded line-breaks so a record may span more than one line. Fields are separated with commas. Leading and trailing space-characters adjacent to comma field separators are ignored. Space characters can be spaces or tabs. Fields with embedded commas must be delimited with double-quote characters. Fields that contain double quote characters must be surrounded by double-quotes, and the embedded double-quotes must each be represented by a pair of consecutive double quotes. A field that contains embedded line-breaks must be surrounded by double-quotes.[55] Fields with leading or trailing spaces must be delimited with double-quote characters.[56] Fields may always be delimited with double quotes. The first record in a CSV file *may* be a header record containing column (field) names. There is no mechanism, however, for automatically determining whether the first record is a header row, so in the general case, this will have to be provided by an outside process (such as prompting the user). The header row is encoded just like any other CSV record in accordance with the rules above.

Gerbv displays well-formed pick-and-place files. Pick-place files should contain XYRS (X, Y, Rotation, Placement Side) information and be held in a comma-separated ASCII (.csv) format.

Common pick-and-place field values are listed in *Tab. 15(138)*. Note that some look for (X)(Y)(R)(S)(D) ordering of the fields. Others look for (D)(S)(X)(Y)(R) ordering. We should put an option in pcbnew to generate either. Also note that there is no footprint data here, but I have seen some Centroid files that contain the (F) field. The fields are described in detail as follows:

RefDes (D): This is the reference designator that matches the BOM and PCB (silk, assembly drawing) designations.

Layer (S): This is either the word "Top" or "Bottom." This is not necessarily the CAD layer designator. "Top" is for a

55. In Microsoft Excel, leading spaces between the comma used as a field separator and the double quote will sometimes cause fields to be read in as unquoted fields, even though the first non-space character is a double-quote. To avoid this quirk, simply remove all leading space after the field-separator comma and before the double-quote character in your CSV export files.

56. Some applications will insist on helping you by removing leading and trailing spaces from all fields regardless of whether the CSV used quotes to preserve them. They may also insist on removing leading zeros from all fields regardless of whether you need them or not. One such application is Microsoft Excel.

part located on the top of the board; "Bottom" for parts on the bottom side of the board. "Top" is often referred to as the "Component" or "Front" side of the board; "Bottom" is often referred to as the "Solder" or "Back" side of the board, by fabrication and assembly shops.

LocationX (X): The X-coordinate describes the part's x-offset from the board origin. The location values require that the part origin be centred in the part. The board XY origin of 0,0 is in the lower left corner of the board. The 0,0 origin for the bottom of the board is the lower left corner, looking at the top of the board, through the board. Preferred units are in inches (0.0000").

LocationY (Y): The Y-coordinate describes the part's y-offset from the board origin. The location values require that the part origin be centred in the part. The board XY origin of 0,0 is in the lower left corner of the board. The 0,0 origin for the bottom of the board is the lower left corner, looking at the top of the board, through the board. Preferred units are in inches (0.0000").

Rotation (R): Rotation goes counter-clockwise for all parts on top and clock-wise for parts on the bottom. For bottom side parts, it is looking through the board, still from the perspective of looking at the top of the board.

Footprint (F):

Make sure that the part libraries correctly place the part origin in the centre. Dual inline chips have pin one on the upper left when at a rotation of zero degrees. The board origin is in the lower left. The part origin for odd shaped parts is centred in the entire part footprint including the leads, not just the package. Quad packs and BGAs have pin one at the upper left with a rotation of zero. Some PLCC and leadless chip carriers have pin one in the centre of one side. In such a case, pin one would be centred on the upper side when set at zero degrees rotation. Two pin and single inline components are oriented at zero degree rotation with pin one to the left. LEDs have the cathode to the left and the anode to the right. Diodes have the cathode to the left and the anode to the right. Electrolytics and polarized capacitors have the positive terminal to the left and the negative terminal to the right. Non-polarized two-pin passives are oriented horizontally when at a rotation of zero. Three-pin and odd shaped parts have pin one in the upper left like dual inline packages.

When components are on the bottom side of the board, the reference point is still 0,0 on the lower left of the board based on looking through the board from the top of the board. Pin one of the part, with a rotation of zero degrees is mirrorred horizontally from the parts on the top of the board. So, as viewed from the top of the board, pin one is usually on the upper right. Rotation goes counter-clockwise for all parts on the top and clockwise for parts on the bottom. This is from the perspective of looking down at the top (component side) of the board.

Only SMT parts should be listed in the Centroid file. The basic format for the Centroid file is a comma delimited (.csv) file with data in the following order: "RefDes", "Layer", "LocationX", "LocationY", "Rotation." Following is an example listing of a Centroid file laid out in this format.

```
RefDes   , Layer  , LocationX , LocationY , Rotation
C1       , Top    , 0.5750    , 0.2055    , 0
D1       , Bottom , 1.5500    , 1.8255    , 270
LED1     , Top    , 0.1902    , 2.0900    , 0
R1       , Top    , 0.3650    , 1.9750    , 90
R2       , Top    , 0.4755    , 1.1352    , 180
U1       , Top    , 0.6352    , 0.4451    , 270
```

Centroid file data is typically verified against silk-screen data by importing the .csv centroid file and then using the verification features of a program such as GC-PowerPlace, UniCAM or CircuitCAM. The verification features typically display centroid and pin-1 reference from the XY-data overlaid on top of the silkscreen for that position. So it is rather important that the silk-screen always identify the pin-1 position (on 180-degree rotation sensitive parts) and reference designator.

7.8 BOM Formats

The BOM tells the assembler what parts are required to be placed (and which are not) as well as where they should go. Assemblers can accept a BOM in a number of formats: .xls, .xlsx, .csv, and tab-delimited. All of these formats can be read and created in Microsoft Excel. The BOM for a kitted order should include:

1. Line/Item Number
2. Quantity Per (number of instances of a single part number)
3. Reference Designator
4. Part Number (Digi-Key part numbers are okay)
5. Part Description (QFN32, SOIC, 0805, etc.)
6. Type (SMT, Thru-Hole, Fine-Pitch, or BGA/Leadless)

For turn-key orders, add:

7. Manufacturer's Name
8. Manufacturer's Part Number
9. Distributor's Part Number (Digi-Key part numbers are okay)

Please highlight in red any components we will not be placing (DNS), or any parts that are not included in the kit.

A Bill of Materials (BOM) preferably in electronic format should contain the following:

1. manufacturer's Part Number and/or Supplier Part #;
2. component Description;
3. quantity;
4. reference Designator;
5. if possible, include Distributor Part # and Manufacturer #;
6. the BOM should indicate any DNI (Do Not Install) locations.

Required fields:

1. Part Number
2. Reference Designations

Preferred fields:

1. Description
2. Quantity

Optional text fields:

1. Part Marking
2. Package
3. Population Type
4. Customer P/N
5. Manufacturer
6. Manufacturer P/N
7. Vendor
8. Vendor P/N
9. Bin Location
10. Stock Code

11. Software Version
12. Part Revision
13. URL

Optional number fields:

1. Cost
2. Height (mils)

Optional boolean fields (Y/N, Yes/No, T/F, True/False):

1. Polarity
2. Socketed

Custom BOM fields are unlimited.

7.9 CAD Formats

7.9.1 GenCAD Format

GenCAD is a format developed by *Mitron* [GenCAD] that is an obvious precursor to the IPC-2511A GenCAM format. In fact, Mitron donated the GenCAD specification to IPC in 1996. [Jones, 2001] GenCAD is available in the following revisions:

- GenCAD Version 1.4, November 1997
- GenCAD Version 1.3, March 30, 1994 and December 5, 1994.
- GenCAD Version 1.2, August 12, 1993.
- GenCAD Version 1.2 DRAFT, April 23, 1993.
- GenCAD Version 1.1, n.d.
- GenCAD Version 1.0, n.d.

Not too many fabricators announce that they are capable of accepting any version of *GenCAD* format. It appears that a number of bare-board testers (likely sold by GenRad) are capable of reading this format; however, I have not seen that fact advertised. By far the most preferred method of providing netlist information for artwork verification, bare-board test and repair is the IPC-D-356 (Lavenir Format 2) IPC-D-356/NTD (Lavenir Format 4), IPC-D-356A and IPC-D-356B formats (see *Sec. 7.4(136)*).

> In the early 1990's the Mitron Corporation, a leading provided of CAM software tools and now a subsidiary of GenRad Inc., developed a proprietary data transfer format called GenCAD. All Mitron tools were modified to use GenCAD as the means of obtaining data from CAD systems. In order [sic] to increase integration capabilities, Mitron partnered with Router Solutions, Inc. (RSI) to provide CAD translation tools from multiple CAD systems to the GenCAD file format. This approach, coupled with open publication of the format schema, resulted in the GenCAD file format becoming the most widely used structured CAD format in the assembly, test and inspection industry today.
>
> However, the GenCAD format was designed almost ten years ago when object-oriented programming and object/data extensibility was in its relative infancy. It was not intended to support true, hierarchical manufacturing data such as schematics, drawings, assemblies and fixtures. Attempting to extend the format to support these features while maintaining legacy support for its 2D roots would result in format degradation. ... [Jones, 2001]

On the other hand, almost all assembly shops support GenCAD 1.4, and many also support GenCAM. This is likely because GenCAD does not suffer any of the limitations when it was used for board manufacture, because most of its limitations have to do with the inability to represent internal via fabrication structures such as buried vias, microvias, back-drilled vias, blind vias, etc., particulary when filling, via capping and plating over, VTP and SCV is involved. But, for board assembly, only the outer layers matter. GenCAM, being related, is supported also by about half of the assembly software vendors (supported by UniCAM, but not by CircuitCAM). So GenCAD still has its niche. Also, because the last version of the specification was 1.4, released in 1997, it can be, like latin, considered a dead language. This means mature and interoperable support over many years.

GenCAD has been fully supported in pcbnew for the following reasons:

1. Being its precursor, the format is very similar to *GenCAM* (IPC-2511A): so much so that supporting *GenCAD* in conjunction with *GenCAM* is straighforward.
2. A roughed out exporter for *GenCAD* was already present in pcbnew.
3. Although few board-testers will require it (fabricators now require IPC-D-356 for bare-board testing), most assemblers support it as the common denominator CAD file format for short production runs.
4. It is certainly more desirable than ODB++ or another proprietary file format belonging to some other CAD or CAM program. See *Apx. D(150)*.
5. Unfortunately, IPC-2581 (258X) did not catch on with assemblers to the extent that it did with board fabricators. Most board fabrication CAM systems not only support ODB++, but also support GenCAM and 258X. In contrast, assembly shop software supports GenCAD, GenCAM and ODB++, but not 258X. Go figure.

Because the purpose of the GenCAD file format is completely concerned with the outside of the board and not the inner layers, its lack of representing vias internal to the board such as blind and buried vias is not an issue. Its purpose should be, however, for assembly shops and not for bare-board-testing.

7.10 CAM Formats

GenCAD. GenCAD [GenCAD] is an EBNF formatted CAD format that was a precursor to the GenCAM format. The GenCAD format is described in *Sec.* **??***(??)*.

GenCAM (IPC-2511A). GenCAM [GenCAM] is an EBNF formatted CAM format with strong similarities to GenCAD. The GenCAM format is described in *Sec. 7.10.1(140)*.

GenCAM-XML (GenX) (IPC-2511B). GenCAM-XML [GenX] (or sometimes simply referred to as GenX) is an XML formatted CAM format largely equivalent to GenCAM, but formatted with XML instead of EBNF. The GenX format is described in *Sec. 7.10.2(141)*.

ODB++. ODB++ is a proprietary CAM database format developed by Valor[57] that includes a directory of primarily XML coded files. The ODB++ format is described in *Sec. 7.10.3(141)*.

258X (IPC-2581). 258X (IPC-2581) [258X] is a vendor-neutral XML-formatted CAM database format that draw from both GenX and ODB++. The 258X format is described in *Sec. 7.10.4(141)*.

Of the above, GenCAD and GenCAM are EBNF formats. They have the familiar `$SECTION` and `$ENDSECTION` layout. GenX and 258X are XML formats. GenX is equivalent to GenCAM, but is coded in XML (with some structural modifications) instead of EBNF.

7.10.1 GenCAM (IPC-2511A)

> In March 2000 the Association Connecting Electronics Industries (IPC), an ANSI accredited standards organization, release the IPC-2510 series of documents which comprise a new manufacturing data transfer specification known as GenCAM. The standard was founded on the primary

57. See *Apx. D(150)*.

> strengths of the GenCAD format, donated by GenRad Inc. to the IPC standards committee in 1996, but re-designed to support all PCB and PCA manufacturing requirements. The IPC's Data Transfer Solutions committee, cosnisting of representatives from over 40 organizations, developed the GenCAM standard over a four-year period as the solution to the proprietary data transfer issue. These organizations covered all areas of the manufacturing arena, including CAD vendors, sub-contract manufacturers and end-users. [Jones, 2001]

GenCAM 1.0 (IPC-2511A) [GenCAM] is the IPC standardized version of the *GenCAD 1.4* format. It has a good number of enhancements, however, and the syntax is largely expanded and incompatible with *GenCAD*. The *GenCAM 1.0* file format is specified in IPC-2511A which is freely available from *http://webstds.ipc.org/2511/2511Apub.pdf*. The *GenCAM* file format uses a `.gcm` file extension.

GenCAM is another attempt from the IPC to replace manufacturing outputs for board fabrication and assembly. It follows the principle discovered in the electronics industry that placing data into multiple files leads to errors. Therefore, the intent of *GenCAM* is that it contain all manufacturing outputs (the database) in a single file.

GenCAM appears to be roughtly based on *GenCAD*. [GenCAD] Also note that both *Mitron Corporation* and *GenRad (General Radio)* were involved in the specification, where *Mitron* was the original creator of the *GenCAD* format, [GenCAD] and *GenRad* bought *Mitron* and made it a subsidiary in 1996.

7.10.2 GenCAM-XML (IPC-2511B)

GenCAM 2.0, *GenCAM-XML* or simply *GenX*, s a format defined by IPC-2511B [GenX] that builds on the *GenCAM* format of IPC-2511A but proscribes the use of *XML* for syntax. The format is incompatible with IPC-2511A (*GenCAM*). The *GenCAM 2.0* file format is specified in IPC-2511B [GenX] which is freely available from *http://webstds.ipc.org/2511/2511Bpub.pdf*. The *GenCAM 2.0* file format uses a `.gcm` file extension.

GenCAM-XML is another attempt from the IPC to replace manufacturing outputs for board fabrication and assembly. It follows the principle discovered in the electronics industry that placing data into multiple files leads to errors. Therefore, the intent of *GenCAM-XML* is that it contain all manufacturing outputs (the database) in a single file.

The reason for *GenCAM-XML* versus *GenCAM* is the rise popularity of the *XML* file format at the time of its specification. Note that *SGML* has been used by government for over twenty-years to maintain document information.

7.10.3 ODB++

7.10.4 258X (IPC-2581)

258X[58] is a format defined by IPC-2581 that builds on both *GenCAM* and *ODB++*. *ODB++* is a format developed by *Valor* (now part of *Mentor Graphics*) primarily for their CAM systems and products. *ODB++*, while popular for CAM systems is unfortunately a proprietary format. Not too many CAD systems were developed to generate *ODB++* because of the lack of vendor neutrality in the format. IPC attempted to alleviate this difficulty by working with *Valor* to specify the *258X* format.

Currently all systems that support *ODB++* (including *Valor's* own CAM products) also support *258X*. As far as adoption of *258X*, it appears to be slow. The IPC launched a website devoted to the vendor-neutral format, *http://web.ipc.org/*, that still exists and provides the latest specifications.

Some support has been added to pcbnew to support export of this format. The format is available by menu under the `File-> Export-> 258X` menu item.

7.11 Exchange Formats

7.11.1 IDF (Intermediate Data Format)

7.11.2 DXF (Drawing Exchange Format)

DXF (Drawing Exchange Format) is a CAD file format developed and maintained by Autodesk [DXF]. KiCad has historically supported DXF as a plotting format. Although some board fabricators accepted (and still accept) PCB artwork in DXF form, it is rare and dwarfed by the universal acceptance of the Gerber RS-274X format, and even the DPF format, for artwork. Nevertheless, DXF format is an important format for exchanging data between CAD systems, even if it is not terribly useful as a manufacturing output format. The popular open-source QCad program uses DXF as its native file format, and InkScape is also capable of reading and writing DXF files.

There are some deficiencies with the historical KiCad plotting to DXF files:

1. The current DXF plotter does not support drill marks or zone cutouts even though the DXF file format is quite capable of supporting clipping shapes.
2. The current DXF plotter plots text as the individual line segments making up the text. AutoCAD supports custom stroke fonts using .SHX files. QCad uses .CXF files for the same purpose. It is quite easy to convert the KiCad modified Hershey's and NewStroke stroke fonts to a .CXF file. Fonts are discussed in more detail in *Sec. 7.12(142)*.
3. The current DXF plotter plots in 2D only. The DXF file format is capable of basic 3D representations: extrusions and rotations. With support for IDF export and import, it should be quite possible to export extruded keep-outs and package elevation and extrusions. Also, it should also be possible to extrude traces and copper patterns as well as placing the patterns on the z-axis. Being able to export 3D data would help with visualization of the IDF import and export as free packages such as QCad do not support import of IDF; even though most commercial grade CAD packages do.

IDF (Intermediate Data Format) is a file format developed by Mentor Graphics that provides for the exchange of basic electronic and mechanical information between an ECAD and MCAD. There are three readily available versions of the specification:

1. *IDF (Intermediate Data Format) Version 2.0*, January 5, 1993. [Kehmeier, 1993] The earliest readily available version of the specification.
2. *IDF (Intermediate Data Format) Version 3.0*, October 32, 1996. [Kehmeier, 1996] Just a slightly improved release over version 2.0.
3. *IDF (Intermediate Data Format) Version 4.0*, July 16, 1998. [Kehmeier and Makowski, 1998] This version was the last version of the specification ever released. It represents a radical departure from version 3.0 in terms of both file format and scope of represented information.

Many CAD systems (such as AutoCAD and SolidWorks) support the exchange of information on PWA (Printed Wiring Assemblies) using the various versions of IDF file format.

Support has been added for exporting IDF version 2.0 or 3.0, as well as importing these versions. The IDF version 4.0 file format is not supported for import. Some rudimentary support has been provided for export. The reasons for not supporting version 4.0 import are:

58. Also sometimes referred to as "Offspring."

1. Any MCAD or CAM system that supports IDF version 4.0 import or export, also supports IDF version 3.0.

2. Version 4.0 is far too rich to be usable for its intended task, particularly for import.

7.11.3 SPEECTRA DSN

7.12 Character Fonts

Historically, pcbnew used a variation of the popular Hershey's "small" single stroke fonts to create text on copper and technical layers. More recently, the NewStroke font has been used. There are a number of difficulties and trade-offs associated with these approaches.

One of the purposes of the original Hershey's stroke fonts was to provide legible fonts for use with pen plotters. The original mechanical plotters were not too good a plotting arcs, so the Hershey's stroke fonts use straight line segments.

In recent years minimal single stroke fonts have a renewed interest due to the fact that they reduce both the time and space effort for display.

Recently KiCad has begun using the NewStroke font. This font has several difficulties for pcbnew:

1. There is a far greater number of average segments per character. This results in much larger Gerber files. The issue with size is not one of disk capacity. Very long Gerber files with many apertures, particularly when not organized into layers, is much more complex to handle. Manual intervention from a human being (hand editing) becomes much more time consuming. Some fabricator charge more when the Gerber file size gets larger.

2. The rounder the symbol, the greater the number of segments. NewStroke tries to get smoother looking symbols with a drastic increase in the number of segments.

3. Rounder symbols do not increase readability for pcbnew. In fact, it reduces readability. Legibility is affected by the height, width and thickness of lines. The design of legend on high-density boards means that character height, width, line width and separate are always at a minimum. Stronger lines along the edges of the character and more open interior characters are more legible on silk screen. A more open character is more legible because there is less change of bleeding filling an area inside the symbol.

On the other hand, when printing schematics with eeschema, the characters are never represented with screen printing and are normally viewed electronically, or laser printed at incredible resolution (incredible when compared to silk screen). The only issue with the NewStroke font for eeschema is that the larger the number of segments, the longer a canvas takes to render. There are some users and developers that would like the modified Hershey's fonts back, both for pcbnew and for eeschema.

The foregoing tends to indicate that a better approach would be to provide loadable fonts for KiCad and to permit the designer to select the font she would like to use for each occasion. Currently, KiCad compiles in an array of characters in which the font metrics and line segments are encoded. There is certainly no reason why the default and optional fonts can not be read in from a file when the program loads or when a font is to be selected. The default fonts can still be compiled in. The first two fonts should be the old modified Hershey's font and the NewStroke font.

(R) 41 (KiCad) *KiCad will be enhanced to support loadable single-stroke fonts.*

Rendering and Plotting of fonts. As stated previously, the purpose of the original Hershey's font was to represent a small character with a minimal number of straight lines. There is some point beyond which the number of line segments used to approximate an consume more time to plot and render than it would have taken to plot the arc directly. Some single-stroke fonts, such as the AutoCAD and QCad standard fonts for DXF, use arcs instead of line segment approximations of arcs. Also, when attempting to represent Bezier curves or splines, a more efficient approximation can be made with several arcs than can be made with a greater number of line segments. Also, the circular approximations can be calculated off-line (i.e., during the creation of the font rather than at the point that it is about to be displayed). With this approach, even the outlines of TrueType fonts can be reasonably approximated.

Therefore, to handle the other end of the spectrum, where fonts are made as curvy as possible, say, for high-quality schematic prints, it would be best to plot these font faces using arcs instead of straight line segments.

(R) 42 (KiCad) *KiCad will be enhanced to display and plot circular arc segments in characters as well as straight line segments.*

Single-stroke fonts are scaled by scaling the x and/or y coordinate axes before plotting the strokes. Single-stroke fonts are typically mono-spaced, but may be proportionally spaced. When scaling is specified as an absolute height and width of the character, it refers to a specific face in the font (e.g., the capital letter 'R'). There are four parameters that specify the scaling of the font: stroke width, character height, character width, and slant. Slanting is achieved by adding a value proportional to the y coordinate to the x coordinate before plotting the strokes. For slanting to work, arcs must be changed into ellipses (by setting the major to minor axis ratio to the slope of the slant) or approximated with straight lines before slanting.

Import an export of character fonts. An investigation into the available formats for stroke fonts reveals that Adobe Type2 or TrueType fonts are incapable of representing stroke fonts. The only apparent open specification of stroke fonts it that provided by QCad in its `.cxf` font file format. QCad has 35 font faces, an AutoCAD DXF compatible "standard" font, as well as a very extensive "normal" and "unicode" font. The "unicode" font, in particular, supports a very wide range of unicode symbols supporting the full asian unicode font faces, and is an excellent font to support KiCad's broad i18n support.

(N) 43 (KiCad) *To permit a wider range of available loadable fonts, KiCad should be enhanced to support direct loading of QCad fonts.*

Font conversions. For the purpose of generating high print quality schematics, it would be good to place any TrueType font at the disposal of eeschema for print. The ttf2cxf tool developed by the principle developer of QCad provides a clean mechanism for covering TrueType fonts to a line segment version of the font in QCad format utilizing the FreeType library. The only deficiency in the approach is that straight line segments are always used to approximate splines. It is possible to approximate Bezier curves with circular arc segments, and it is possible to approximate splines with Bezier curves.

(W) 44 (KiCad) *To permit a wider range of available loadable fonts, KiCad should be enhanced to support stroke versions of TrueType fonts. The ttf2cxf tool can be enhanced to approximate splines with circular arcs instead of line segments. This enhanced version of the ttf2cxf can be incorporated into the common text handling code for KiCad.*

Embedded fonts. As experience has shown with the PDF (Portable Document Format), the only way to ensure that design files are self-contained is to embed any font (beyond a simple set of built-in fonts) into the file itself. Otherwise, the rendering of text can be compromised by the unavailability of the custom font face.

(N) 45 (KiCad) *To insulate the design files from changes to the set of available or custom loadable fonts, KiCad should be enhanced to be able to embed fonts in the pcbnew and eeschema file formats.*

Font selections to meet process parameters. Font size and shape is determined by stroke width, character height, character width, slant, and inter-character spacing. Process parameters for silk screen and copper characters typically specify stoke width, character height and character spacing, but not character width. Character width is determined by applying a maximum aspect ratio, $\alpha_{max} = \frac{H}{W}$, to the character. The maximum aspect ratio is dependent upon the character font. A typical maximum aspect ratio is 2.0. Historically, pcbnew calculated the maximum aspect ratio internally and would not permit the use to exceed the hard-coded values.

(R) 46 (pcbnew) *The user will be allowed to set the maximum aspect ratio for a given font.*

Historically, pcbnew has had only one place for setting default character stroke, height and width. One place for the entire board. Because the needs of characters on silk-screen, copper, and in drawings are separate and distinct, process minimums must be split by process layer. Because a default font selection must be provided for each layer, a single default character stroke, height and width is insufficient.

(R) 47 (pcbnew) *Selection of font and the default font parameters (stroke, height, width, slant) will be added to the design parameters of each layer that has process controls for characters (silk and copper). The default font parameters and selection will apply to all other layers (comments, drawings, etc.). To avoid always designing to the limit of manufacturing capabilities, even where less stringent design parameters would do, the default font parameters need to be separate from the minimums.*

(R) 48 (pcbnew) *To be backward compatible (so that silk and copper designed with previous versions of KiCad can be recreated from the board file), all of the previous fonts will be provided. Three built-in fonts will be provided:*

1. *The original KiCad modified Hershey's font. This will be the default for legend and copper layers.*
2. *The QCad "unicode" font. This will be the default for schematics, comments and drawings layers.*
3. *The NewStroke font. For those that like it (or need it because it was used previously).*
4. *The QCad "iso" font. For those that require it.*

Another issue with character fonts is whether to render characters as their separate single-stroke segments or to simply specify the parameters of the text. Note that RS-274D and DPF formats both permit specifying text as text instead of drawing out the characters.

7.13 Graphics Formats

pcbnew has historically been able to plot, print, or export, to several graphics formats. Some additional plot, print, and export, formats have been added. One of the things that pcbnew has been lacking is the ability to *import* drawing formats. Importing drawing formats is important for the following reasons:

- pcbnews ability to edit drawings is crude. It would be advantageous to be able to use a full-blown drawing program to edit and then import complex drawings into pcbnew.
- Extending pcbnew to support sophisticated editing ability for drawings would be too much duplication of effort.
- Sophisticated open-source drawing programs exist on all platforms. The user should be able to use their drawing program of choice.
- pcbnew is capable (or should be capable) of supporting all (or most) 2-D drawing primitives.
- Sophisticated drawing programs would assist well with importing complex drawings into pcbnew for fabrication prints, assembly drawings, fabrication detail, and assembly detail.

Two popular programs that can be used for drawing on the Linux platform are InkScape and XFig. InkScape uses SVG (Scalable Vector Graphics) as its native file format. XFig uses its own file format, but a conversion tool exists to convert XFig drawings into many forms not supported by other programs. QCad uses DXF [DXF] as its native file format. It is also capable of writing HPGL files.

DXF: pcbnew has historically been capable of exporting or plotting to DXF files. QCad is a 2D CAD program that uses DXF as its native file format. InkScape is also capable of reading and writing DXF files. DXF format is discussed in more detail in *Sec. 7.13.1(143)*.

SVG: pcbnew has historically been capable of exporting or plotting to SVG files. InkScape is a program that uses SVG as its native file format. SVG format is discussed in more detail in *Sec. 7.13.2(144)*.

XFig: pcbnew has not historically been capable of exporting or plotting to XFig files. XFig is a drawing program that uses XFig as its native file format. InkScape is capable of reading XFig files (but is not capable of writing them. XFig format is discussed in more detail in *Sec. 7.13.3(144)*.

Gerber (RS-274X): pcbnew has historically been capable of plotting to Gerber RS-274X files. gerbview and gerbv are viewer and editing programs that use RS-274X as their native file format. Gerber format is discussed in more detail in *Sec. 7.13.4(144)*.

Barco DPF: pcbnew has not historically been capable of plotting to Barco DPF files. Barco format is discussed in more detail in *Sec. 7.13.5(144)*.

HPGL: pcbnew has historically been capable of printing to or plotting to HPGL files. Qcad can also output HPGL files. InkScape is also capable of reading and writing HPGL files. HPGL format is discussed in more detail in *Sec. 7.13.6(144)*.

Post Script: pcbnew has historically been capable of printing to or plotting PS files (but strangely not PDF files). InkScape is capable of reading and writing Post Script and Portable Document Format files. PS format is discussed in more detail in *Sec. 7.13.7(144)*.

7.13.1 DXF

Historically, pcbnew has supported DXF [DXF] as a plot format. DXF is a basic graphical language, so an import function for DXF can attempt the primitive graphics (line segments and arcs basically). The import issues are:

- The internal representation and plotting (e.g., Gerber output) of ellipses and splines is an issue. pcbnew and eeschema have an internal representation for Bezier curves, and ovals but not for ellipses. It is possible to approximate an ellipse

with an oval. It is also possible to approximate a spline with a Bezier curve. It is also possible to approximate both with line segments or arcs for plotting (e.g., Gerber plots).

- The internal representation and plotting of text is also an issue. DXF stroke fonts are a different issue and are basically an issue with the export of text from KiCad as well. When importing DXF, text is given a font face. When importing and exporting it is typical to use only the "standard" AutoCAD stroke-font face. Some of the requirements in *Sec. 7.12(142)* can meet the objective of properly representing fonts for text imported from QCad DXF files.

Because it is a mechanical CAD format, the issues associated with DXF export and considered in more detail in *Sec. 7.11.2(141)*.

(W) 49 (KiCad) *It would be good to be able to import DXF formatted files into KiCad for use by both pcbnew and eeschema.*

In general, however, DXF is far too rich a format to permit reasonable import into KiCad. The restrictions that would be applied to a DXF file to allow it to be successfully imported are too onerous to be practical.

7.13.2 SVG

Historically, pcbnew has supported SVG as a print format, but not as a plot format. To put it more precisely, pcbnew "displayed" the board to the SVG file instead of printing, plotting or exporting. SVG print is really only useful in its historical form for tasks that could equally well be served by screen-shots. To support SVG export we want to actually plot to the file. An import function for SVG could only attempt the more primitive graphics (line and arc segments basically). In general, SVG is far too rich a format to permit reasonable import into KiCad. The restrictions that would have to be applied to an SVG file to allow it to be successfully imported are likely too onerous to be practical.

(W) 50 (KiCad) *Create an SVG plotter in KiCad to support plotting to SVG files. This is in addition to the current SVG print capability.*

7.13.3 XFig

Historically, pcbnew has not supported XFig. Interestingly enough, XFig is one of the few formats that InkScape cannot write. All of the figures in this document are created with XFig. This is because XFig is well supported for LaTeX, and can generate all of the picture formats necessary for generating any type of document with latex (dvi, eps, pdf, jpg, png, etc.).

Several things will be added to pcbnew in support of XFig.

- An XFig plotter will be created to support plotting to XFig files.
- A graphical importer will be created to support importing XFig files in a similar manner to other graphical files.

Font Handling: XFig is rather simplistic when it comes to font handling; however, its fonts are sufficient for generating PS, PDF and LaTeX documents. It provides the basic set of built-in PDF fonts (basic Adobe fonts), as well as the TeX Computer Modern fonts. This provide some challenges for successfully exporting and importing XFig. When exporting to XFig there are several choices regarding text:

1. Export characters as a group of arcs and line segments rendering the KiCad font.
2. Export characters using one of the built-in XFig fonts that most closely matches the metrics of the KiCad font.
3. Export characters using one of the built-in XFig fonts, but let the user choose the mapping.

Scaling: XFig only supports line widths in multiples of 1/80th of an inch (about 12.5 mils). The smallest silk-screen line width that can be supported is approximately 6 mils. Therefore, to be able to represent reasonably small line widths, an XFig drawing might need to be scaled to twice the size of the KiCad representation.

Import representation. Likely the most direct approach to supporting the import of graphics is to represent imported XFig drawings and text as a MODULE inside of pcbnew. This serves to permit the entire imported drawing to be moved about as a unit, and yet allows the use of the module editor to alter any of its internal components. Also, because MODULEs have an associated file name (to specify from which library the module was loaded), it is possible to associate the file name with the XFig source file so that multiple imports can be handled. In this case the feature in the module editor that allows a module to be reverted to its library definition can also be used to reload the XFig source file. This nicely suits the purpose of being able to edit complex drawings with another drawing program, and then import them into KiCad. A scheme can be devised for mapping pcbnew layers to and from XFig layers.

(R) 51 (pcbnew) *pcbnew will be enhanced to support importing of XFig drawings and text.*

7.13.4 Gerber RS-274X

Historically, pcbnew has supported plotting to RS-274X. It might be possible to entertain an import function for Gerbers, however, there are no good drawing programs for Gerber format.

7.13.5 Barco DPF

Historically, pcbnew has not supported plotting to DPF. A plotter has been added to pcbnew that plots DPF in a similar fashion to Gerbers. DPF is well-supported by board fabricator CAM systems; however, unfortunately, most assembler CAM systems only support Gerber.

(W) 52 (gerbview) *gerbview can be enhanced to read DPF files.*

7.13.6 HPGL

Historically, KiCad has supported plotting to HPGL formatted files. Some board fabricators accept fabrication artwork in HPGL. Many printers or print drivers support HPGL. But most assemblers do not support HPGL for artwork: only Gerber RS-274X. These is no reason to support import of HPGL graphics for KiCad.

7.13.7 PostScript

Historically, KiCad has supported plotting to PostScipt files. This is in support of various home printing techniques for developing boards. There are several techniques where the application of toner can be used in a process for board fabrication. Many drawing programs do not support the import of PS formatted files, except directly as embedded postscript. There is little cause to support the import of PS files into KiCad with one exception. There is a need to be able to import company or certification logos for transfer onto silk screen or to place in the sheet reference frame. This is discussed separately in *Sec. 7.14(144)*.

7.14 Logos

There is a strong desire to have the ability to import company logos or certification marks into KiCad for the following purposes:

- Placement of company or association logos into the sheet reference frame.

- Placement of company or association logos or certification marks onto silk screen or copper layers for transfer to the fabricated board's surface.

The major challenges with importing logos into KiCad are as follows:

- The representation of the logo in external form is an issue. Although company logos might be available in a number of forms (e.g., .jpg, .png) for presentation on company websites, or in postscript and PDF documents. They might be available in indexed or RGB bitmap formats or in vector formats such as SVG.
- The representation of the logo in internal form is an issue. For plotting on artwork, it is preferred that the information be represented as a collection of basic filled closed polygons, line segments and arcs.

Logos are usually created once and then imported into KiCad and placed. Creating an eeschema library component (symbol), or a pcbnew module is one of the best ways of incorporating a logo or certification mark into a schematic or board. Because the representation of the logo or certification mark is in a native KiCad format, off-line tools can be used in the creation of the component or module.

KiCad has a bitmap2component directory that includes a dialogue and program that can be used to assist in converting bitmaps to linear closed polygons using the potrace library. This approach is quite viable for off-line generation of these modules; however, the potrace utility is capable of generating both SVG and XFig vector graphics formats. Therefore, the import capability for XFig described in *Req. 51(144)* could be used to import a logo or certification mark. The potrace utility is well suited to conversion of bitmaps into a vector graphics format that could then be converted into a component or module for KiCad. The actual conversion between the vector graphics format and KiCad native components or modules could be performed off-line with scripts or as part of a graphics import function such as that described for XFig in *Sec. 7.13.3(144)*.

A Description of Vias

Machine drilled blind or buried vias can be constructed using a number of techniques.

Sub-Composite Vias (SCV): sub-composite buried vias can be created by making a PTH through a sub-composite laminate. Typically it is better (lower cost) if these SCV cross from the outer layer of one balanced core to another. Where m is the frontmost middle core layer and $m+1$ is the backmost middle core layer, buried vias can be made from layer $m-n$ to layer $m+n+1$ where $n < m$ and $n = 2, 4, ...$ So, for a 16 layer card, $m = 8$, and buried vias can be made between layer pairs: (8,9), (6,11), (4,13), (2,15). Depth-control-drilling can be used to make buried vias between different layers; however, depth-control-drilled vias can only be plated to a maximum aspect ratio. Disregarding that fact, DCD vias must still be drilled from an outer balanced core layer, but can be created to any inner layer. Therefore for a 16 layer card, vias can be made between the following additional pairs: $(2, 3 \leq n \leq 14)$, $(4, 5 <= n <= 12)$, $(6, 7 \leq n \leq 10)$.

Notes about sub-composite buildup:

1. Any core in the stack-up can be a sub-composite. That is, it can be through drilled and holes can be plated.
2. Any two sub-composites can be joined by a prepreg layer into a larger sub-composite. Such a sub-composite can be through-drilled or depth-control-drilled (or both) and holes plated. Depth-control-drilling must not violate maximum aspect ratios for plating.
3. Any sub-composite can have an HDI prepreg or prepreg+foil layer attached to make a larger sub-composite. The HDI prepreg or prepreg+foil layer can be laser drilled for microvias, depth-control-drilled, or through drilled and holes plated. Laser drilling and depth-control-drilling must not violate maximum aspect ratios for plating.

Stackup should be from an inner core outward with balanced thickness and buildup. Sub-composites should be from an inner core outward, also with balanced stack-up. Anywhere in the board where an HDI prepreg or prepreg+foil layer exits, the trace quality of the foil will be inferior to that of a core due to plating variations. As the complexity increases, cost increases dramatically. For example, each sub-composite stage requires 2 hours of press time.

Not all combinations of blind/buried/micro vias are possible. Situations that are not possible are as follows:

1. A buried via that spans only a prepreg layer is not possible. However, depending on the resulting aspect ratio, these may be able to be depth-control drilled from one side or the other, resulting in a stub on one side of the sub-composite or the other. They may also be through-drilled, resulting in two stubs on either side of the sub-composite. The stubs could possibly be back-drilled.
2. A buried via that spans only an HDI prepreg or prepreg+foil layer cannot be created as a PTH and must be laser-drilled (HDI prepreg) or depth-control-drilled (prepreg+foil). If maximum aspect ratios would be violated, through-hole drilling is required which may or may not result in a stub.

This has implications for via stacking as well. When stub creation is nescessary to create a blind or buried via, the stub cannot conflict with another via (blind or buried) or via stub at the same location. pcbnew is currently not able to resolve these conflicts.

Tryng to simplify things. Vias are currently defined to have two layers, a top and bottom layer. Through vias are given a top and bottom layer equal to the front and back of the board. This needs to be changed.

New approach: Vias have two functional layers defined, top and bottom. Layers between the top and bottom function layers can also be functional if the layer connects to a track, a pad or a zone.

Vias also have two presence layers defined, top and bottom. These are initially set to front and back to calculate presence. If there is a unique sub-composite of the board that extends between the top functional layer and the bottom functional layer, the presence layers can be set to the corresponding functional layer. Otherwise, to find the optimum combination follow this procedure:

1. If neither presence layer is equal to a functional layer, remove a sub-composite component from both sides of the board (i.e. HDI prepreg, prepreg+foil, or core+prepreg).
2. If either presence layer would be within the functional layers, back up and quit removing sub-composite components.
3. If both presence layers equal the corresponding functional layer, set via to PTH sub-composite buildup and we are done for this via.
4. If the depth between the top presence layer and the bottom functional layer does not violate maximum depth-control-drilling aspect ratios, set the bottom presence layer to the bottom functional layer, set the via to DCD sub-composite buildup, front-side-drilled and we are done for this via.
5. If the depth between the bottom presence layer and the top functional layer does not violate maximum depth-control-drilling aspect ratios, set the top present layer to the top functional layer, set the via to DCD sub-composite buildup, back-side-drilled, and we are done for this via.
6. Set the via to PTH sub-composite buildup, and we are done for this via.

Microvias are different:

1. If the microvia would span layers other than HDI prepreg or prepreg+foil, refuse the via.
2. If neither presence layer is equal to a functional layer, remove a sub-composite component from both sides of the board (i.e. HDI prepreg, prepreg+foi, or core+prepreg).
3. If either presence layer would be within the functional layers, refuse the via.
4. If both presence layers equal the corresponding functional layer, stop removing sub-composite components. If the intervening layers are HDI prepreg or prepreg+foil, we are done for this via.

A difference between the top presence layer and the top functional layer represents a stub. Likewise, a difference between the bottom functional layer and the bottom presence layer represents a stub. Drilling between the top presence layer and the bottom functional layer represents a depth-control-drilling opportunity. Likewise, drilling between the bottom presence layer and the top functional layer also represents a depth-control-drilling opportunity. When depth-control-drilling is permitted.

B To Do List

1. Expanded microvia definition. Microvias have been defined as going from an external layer to the next internal layer only. This is not correct. Microvias are typically laser drilled and can normally extend from an outer layer through to one or two internal layers. In fact, the depth is not controlled by the number of layers at all, but is controlled by the maximum aspect ration of the laser drilling process and the depth of the external layers. Also, as part of the lamination buildup, it is possible to laser drill internal layers before laminating outer layers. Therefore, microvias can extend between any number of internal layers as long as the maximum and minimum aspect ratios are within bounds.

Microvias, being laser drilled, depend upon the drilled-thru dielectric. HDI layers that contain microvias need to be made of laser-drillable dielectric. Also, laser-drilled dielectric layers are prepreg and not core. DFM rules for microvias should include checking whether layers involved cross an HDI dielectric layer.

2. Restrict blind/buried via definition. Blind or buried vias are machine first drilled as part of the buildup of laminated layers. However, blind vias must extend from an outer layer to an inner layer, however, because they are made as part of the lamination build-up, the blind via must extend across a core inside the board. Buried vias, because they are machine drilled as part of the buildup process, must extend from the outside copper layer of a core to the outside layer of the same or another core. Often, this is further restricted to crossing only one or two cores. Currently, pcbnew permits any layer combination of blind or buried vias.

Also, buried vias should normally be filled. This is because when the board heats, the trapped gasses within the buried via will expand effecting a separation pressure on the laminate.

Blind or buried vias that violate the rules of crossing cores in a sub-composite can still be fabricated using depth-control-drilling.

3. Restrict depth-control-drilled via definition. Depth-control-drilled via are either first drilled after the laminate is completed (making a blind via), or drilled as part of the build process (making a buried via). Buried DCD vias, therefore, are subject to the restriction that they pass from a layer that is an outer layer during the buildup process, to an internal layer. Both blind and buried DCD vias are subject to a restriction on maximum and minimum aspect ratios. The thickness of the dielectric between layers needs to be known to check the via for DFM rules. Currently DCD via are unrestricted in layer pair.

4. Restrict back-drilled via definition. Back-drilled vias can be back-drilled from both sides; however, only one side is preferred. The preferred side is the bottom of the board. It is necessary that the last dielectric layer into which the back-drill penetrates is at least 5mil thick. Currently back-drilled vias are unrestricted in the layer pair.

5. Convert to multiple view Model-Controller-View pattern. WinEDA_PcbFrame's are views of a BOARD. Currently pcbnew ties the board to one frame and display panel. This is not necessary. A BOARD could be have multiple views. Convert the model to have a BOARD associated with a list of WinEDA_PcbFrames. Each WinEDA_PcbFrame can be an open window. There are a number of globals used for track editing that need to be moved inside the board. This will allow mutliple window to be open viewing the same board. Also, there is no reason why multiple window cannot be open for separate boards. This could also permit moving (cut and paste) and copying between boards. Also, this would permit the module editor to simply be another view. Currently the module editor is modal. With this change, the module editor could be made non-modal.

This has benefits for cvpcb and gerbview as well. In paricular, the module editor (well, viewer) in cvpcb is modal. A footprint that is being display must be closed before further actions on the cvpcb main window are possible. This is annoying. With a full multiple MCV (Model-Controller-View) model, multiple footprints could be displayed with no need to close them before taking further actions in cvpcb. The same is true for the pcbnew module editor.

***6. Add differential pairs to* eeschema.** Differential pairs can most easily be handled with a slight modification to the bus concept. Buses currently collect net names that have the same prefix and a numerical ending corresponding to the bus index. Differential pairs are similar in that they have a collection of (two) net names that differ in their termination. Terminations for differential pair net names in schematics normally end with _P/_N or +/- or some times /_B. All of these varations of upper and lower case should be accomodated. Bus names for normal busses have the prefix and then [f..n], where f and n are the start and ending bus member indices. Differential pairs should be the same, where [+-] or [PN] or [B] are used instead (with lower-case variants). Then it would also be possible to have busses of differential pairs BUS[0..n][+-] or differential pairs of busses BUS[+-][0..n].

7. Teach* pcbnew *about busses and differential pairs. Currently pcbnew ignores net names altogether and treats them as arbitrary strings. When pcbnew reads the netlist, it should detect net names that only differ in a number or in a +/-, P/N, or /B. It should collect bus members into their own netclass and differential pairs into their own netclass (when they are not already assigned to a netclass other than the default netclass). The names of these netclasses should be the net name without the member suffix. The netclasses should be initially assigned the default netclass values.

8. Teach* eeschema *about various signal types. The signal types and ERC rules of eeschema are severely lacking. There should be signal types such as those found in FGPA I/O constraints. Signal types should include both standard and voltage. Power types should have a voltage associated with the power signal. This would greatly enhance ERC checks.

9. Teach* pcbnew *about vairous signal types.

10. Export the Readme.txt file.

11. Export ODB++ data.

12. Export IPC-2581 data. IPC-2581 is an XML formatted database similar to ODB++ and GenCAM (IPC-2511). Most CAM systems (even Valor's) can import both ODB++ and IPC-2581 databases. There is a significant amount of publicly acessible information on IPC-2581 format. To access Valor's ODB++ format, Valor requires an NDA, making it largely unusable for opensource EDA programs.

13. Export IPC-D-356/NTD data. Many CAM tools (and therefore fabricators) can import IPC-D-356 and NTD data for feeding information to flying probe testers as well a for creating test harnesses for bed-of-nails testers. It is also used by assembly CAM tools to reverse engineer and verify footprints from Gerbers. GenCAD and GenCAM (IPC-2511) serve these purposes as well; however, it never hurts to have many widely adopted formats for output. Also, many fabricators list Gerber + IPC-D-356 + NTD as the minimum database input requirements.

There was sufficient information on the web about the IPC-D-356 and NTD file format (Lavenir Format 4) and the precursors

to IPC-D-356A to be able to write an export function. IPC-D-356B has very little freely available information (IPC requires that documents be purchased). An option was added to the fabrication outputs menu to generate IPC-D-356A and NTD files.

14. Plot Barco/Ucamco DPF format. Many CAM tools (and therefore fabricators) can import and export Barco/Ucamco DPF format [DPFv5] as well as Gerber. Barco DPF format has the advantage that it is directed at raster scanners rather than the traditional Gerber machine and therefore can accept negative polarity overlays (i.e. REVERSE apertures). [DPFv7] This allows features such as clearance holes to be added to the end of a plot (like postscript). Currently HPGL, Gerber and DXF files are incapable of this. The only current plotter that is capable of creating clearance holes in the final stage is the PS plotter. Where clearance holes are necessary to avoid direct drill on groud plane (without thermals), DPF format is superior to Gerber.

Sufficient information was available from the web and Ucamco to create a plotter for DPF. The plotter supports v5, v6 and v7 constructs; however, the plotter does not calculate true-objects. The plotter does include netlist information. This makes Barco DPF equivalent to Gerber + IPC-D-356A + NTD, however, the information on a layer-by-layer basis is contained in a single file, resulting in less data mismatching.

15. Ground plane stitching. Need more control over whether vias connect to zones or not. Currently all via connect to all zones of the same net (as long as netcode != 0). This can cause problems with attempting to control ground loops and return path currents. For example, when a CMT differential pair change layers, the return current attempts cross ground planes as well. One way to permit this return current to flow between planes is to place ground plane stitching vias on either side of the signal vias. However, it is difficult in pcbnew to generate one of these vias without using segments from a pad to tie the vias to a net. When a via is created from a layer that has no net connection (i.e. no track), check whether the initial layer of the via has a ground plane and whether the via hits in the ground plane, if so, give the via the netcode of the ground plane. Then, by setting the layer pair appropriately, the V command can be used to place stitching vias.

16. EMC isolation using vias. Very high-speed board designes use plated board edges as well as isolation vias to surround internal traces with isolating copper. Currently pcbnew will not allow an isolated via with no netcode and vias with no netcode that are placed close to each other are considered to violate DRC rules.

C eeschema Wish List

This section addresses some eeschema wish-list items in conjunction with the work-flows from eeschema to pcbnew.

C.1 Junctors

Trying to keep connected nets separate in pcbnew when they have been joined in eeschema is pain. Either 0Ω resistors have to be stuffed between the nets, or some other tricks (draw segments on copper layers) have to be pulled. Examples are joining AGND and VSS at a power supply, pulling voltage sense lines from regulators (pcbnew thinks these are redundant tracks and removes them), or controlling the stub lengths on buses or stub-terminated logic (SSTL).

What is proposed is for eeschema to keep nets that are connected using a junction separate when generating its netlist and then adding a set of junctions and their connected nets to the end of the netlist for pcbnew to use. Within pcbnew the traces adjoining a pad will be labelled with the net associated with the pad and will only join with the other nets belonging to the same junction at a point off a pad (e.g., at a via or at a different layer of a pin, or at a track end belonging to the other nets).

Therefore, in the separated ground case, one would label the ground pin on the power supply with, say, PGND, and then attach the GND and AGND nets to PGND using a junction. Then on the board layout when drawing a trace starting from the power supply pin, it will be labelled PGND; drawing a trace from an AGND pin would be labelled AGND; a trace from GND pin, GND. The traces (or zones) could be joined at a via, or at the hanging end of one of the traces. pcbnew would know (well, be taught to know) that these are connected nets because it has the junction information and would not complain in DRC (in fact it will complain when they are not connected).

SPEECTRA DSN has a facility in the file format for specifying this to the freerouter. They call it a *virtual pin*. It is even possible to specify maximum trace lengths from a pad to a *virtual pin* so that the autorouter knows where to join these nets. In the example above, the PGND pin to the *virtual pin* (same as junction), could be limited to be small (say, maximum 5mm), and then the autorouter would pull all of the traces back to the *virtual pin* at the power supply PGND pin that joins GND and AGND. Copper planes can work too.

Another need for this facility is to control the stub lengths on SSTL and other buses. With the junctions in pcbnew, the displayed information can include stub-length, propagation delay, mismatch, etc., because the program can then know that there is a *virtual pin* where the nets join. This can already be done to some extent, because where three traces join is always a junction, but the net names will also be correct, and a trace length report can be generated with the appropriate net names.

It would even be nice to have eeschema be able to specify characteristics of these subnets such as maximum propagation delay, mismatch, etc.

C.2 Buses and Pairs

Historically eeschema has avoided specifying anything about buses or pairs. When passing a netlist to pcbnew bus and pair information is lost. In general, bus members need to be (and are usually for clarity) named as SOMENAMEn, where n is the number of the member, either going from 0 to n, or 1 to n or some such. When these lines are brought into a bus structure on the schematic, eeschema knows that the nets are related, but fails to relay this information to pcbnew. In a similar case, differential pairs are often named SOMENAME_N and SOMENAME_P or some such. Alhough eeschema does not currently support this, a pair structure similar to the bus structure could be used for associating pairs in a differential pair.

Now, to some degree pcbnew could try to associate bus members together simply based on naming (looking for net names that are identical except in the trailing digits), or even differential pairs (looking for net names that are the same except for a _N or _P ending); however, particularly in the case of buses, it is easy to fool the pure naming approach: for example, are the nets VCC25 and VCC33 bus members?

What is proposed is for eeschema to output a list of buses and pairs at the end of the netlist. This would definitively identify to pcbnew which nets belong to a bus and which belong to a pair. This would allow pcbnew to group these nets into separate net classes by bus and pair. Also eeschema would be enhanced to have a differential pair structure on the schematic that looks just like a two-member bus structure. Then eeschema could enforce the naming conventions for specifying pairs in a differential pair bus.

In conjunction with the foregoing discussion of junctions, when bus junctions are formed, they can be added to the net list as a bus junction instead of just n individual junctions. It would even by nice to have some way for eeschema to specify the restrictions placed on bussed subnets extending from the bus junction.

C.3 Signal Families

eeschema allows the schematic designer to specify the basic type of pin for the purpose of performing basic ERC checks. However, normally much more information about the electric requirements of a pin are available at the time that the schematic is drawn. For example, it may be known that a pair of pins form an terminated or unterminated LVDS pair, or that a given net is a terminating resistor stub from an SSTL logic connection. The logic voltage thresholds and common modes are known. Certainly much more information is available over just “input” or “output” or “tri-state”. If this information can be gathered at the time that the schematic is generated, then appropriate rules for layout can be easily generated.

For example, I have a pin that is a 1.8V SSTL logic output pin and it is connected through a junction (or passive) to four SSTL logic input pins. It is then easy to determine (from the logic standard and the frequency of the signals) what the maximum acceptable propagation delays and mismatches there are between the input pin and the junction, and between the junction and the terminating pins.

What is proposed is to enhance eeschema to capture the frequency, waveform, voltage and logic standard of pins and to place this information into the netlist. At least the voltages of logic families supported by popular FPGA should be selectable.

C.4 Programmable Pins

The effects of the programmability of pins on FPGA and other programmable logic devices on the schematic design process is well known. eeschema and its library structures provide no help in this regard. Determining the binding of logical to physical pins on an FPGA and optimizing both the board layout and the internal connections within the FPGA is a very difficult process without the help of a CAD system. In this use case, KiCad is a CUD system (Computer Unassisted Design).

Because an non-optimal netlist (or pin-out) within an FPGA is often far less critical than a non-optimal board layout (the delays involved on the board are far greater than the delays across the die), a typical design cycle would be to design the firmware for the FPGA and generate a netlist and pinout without external constraints to get a first approximation. It would be best if the signal names used for external connections in the FPGA would map to the net names on the board. It would be good if the pin assignments made within the FGPA design suite could be imported into the schematic. This would ease changing the pin

functions on the schematic to match the FPGA synthesis, and then attempting to route the board with these pin-outs. From even cursory examination, it will likely be obvious to the board designer that a change in pin-out (swapping or rearranging pin assignments) would benefit the layout. Because experimentation with routing is an easy thing to do from pcbnew, there should be an ability to back-annotate net changes for pins from pcbnew to eeschema and even to the FPGA design suite (as a set of external constraints). Another synthesis run and re-import into eeschema and another netlist to pcbnew and the cycle is complete.

Of all the components in this cycle, eeschema is the one that cares least what is going on during the cycle. The schematic doesn't care that net MYGROOVYSIGNAL is connected to pin A6. It just needs to be connected to the right pin from the perspective of the firmware. Therefore, to tighten the loop, eeschema should be removed. The FPGA design suite outputs of pin assignments should be imported directly into pcbnew. When changes are made within pcbnew they should be exported back out to the FPGA design suite in a constraint file. Eventually, and when necessary, it should be possible to back-annotate the pin assignments to nets in eeschema.

It should be possible to describe a neutral intermediate format as a simple list of pairings of net names and pin numbers. This simple list can be read by an external script that places it into the form of a constraints file for FPGA synthesis. A similar script could read the FPGA synthesis results and generate a new simple net name pin number pairing list that could then be read by pcbnew. The same simple pairing list could be read by eeschema for back-annotation of nets and pin numbers.

C.5 Components Options

In the design of schematics and boards, it often behoves the process to option the component stuffing on the same board design to achieve separate products with a single debug-test cycle. In this case the board is designed with lands that may or may not be stuffed with a component depending on the option "load".

A simple example is a power supply designed to power digital circuitry where it is not known whether the amount of capacitance will be sufficient to meet voltage ripple and regulation targets. It is quite easy to specify several capacitors in parallel and then only stuff the lesser amount of capacitance. Once power-supply regulation in the finished PWA can be tested, it can be determined whether to stuff additional components.

Another simple example is where an FPGA has several family members with compatible pin-outs that have varying amounts of internal resources. Depending on the other items stuff or plugged into the board, the FPGA model number can be varied to provide the best cost-performance trade-offs. This itself might reduce power requirements and permit the removal of regulators and the stuffing of jumpers (zero-ohm resistors) to bypass the circuitry of the missing regulators. Another similar example is the population of SDRAM, where the capacity of the SDRAM can be smaller (leaving some address leads to, for example, an FPGA, unused).

On some circuits, entire sets of components might be optional. For example, a JTAG or SPI connector and associated components might be necessary for prototyping the board, but can be left unpopulated for production boards (because another interface is available, or ICT can be used to program the device). Some circuits might have special features (such as on-board non-volatile memory) that can be left unpopulated on a lower-cost product.

D The Evil Empire

For the unwary, I have added this section describing the history of the evil empire. I say "evil", because with every step the formats and licensing terms of things issued by these entities gets ever more close and ever more polarized to one vendor.

D.1 Gerber

D.2 Barco

D.3 Ucamco

D.4 Valor

D.5 Mentor Graphics

E IPC-D-356 and NTD Formats

The IPC-D-356 netlist format is an 80 character per line format as listed in *Tab. 16(151)*. The header format is listed in *Tab. 17(152)*.

More information on IPC-D-356:

18.3.2.2 Netlist Formats. A netlist is a set of connection points joined together to form networks. Each point represents a contact point on the surface of a bare board. A connection point can be a drilled hole or a surface mount pad. All points belonging to one network should be connected to each other through PCB layer circuitry and/or power and ground planes.

18.3.2.2.1 IPC-D-356. IPC-D-356 is an ANSI-accepted standard that has become the most widely used standard for transferring netlist information. The IPC-D-356 format is used to transfer netlist information within the PCB design and fabrication community. This information can be used to verify the integrity of the design by netlist extraction from the Gerber graphics with the IPC-D-356 CAD reference. The information is also used within both the bare board and the assembled board test domains.

Most data in the IPC-D-356 file (see Fig. 18.13) consists of electrical test records of two types:

- Drill records (starting with "317")
- Surface mount pad records (starting with "327")

There are also other general parameter records.

```
C This is a comment field
P UNITS CUST
317A01AUXNEG    A10 -1 D0380PA00X+068250Y+002250X0620Y0620
317UN2CAP62PCA0 A10 -2 D0360PA00X+068260Y+001250X0620Y0000
317UN2CAP62PCA0 A9  -1 A01X+066330Y+001600X0800Y0250
327A-1AUXPOS    A9  -2 A01X+066330Y+002500X0800Y0250
327N/C          A9  -3 A01X+066330Y+002500X0800Y0250
327N/C          A9  -4 A01X+066330Y+003000X0800Y0250
327N/C          A9  -5 A01X+064170Y+003000X0800Y0250
327N/C          A9  -6 A01X+064170Y+002500X0800Y0250
```

FIGURE 18.13 A section of an IPC-D-356 file. It contains two drill records ("317") and six surface mount records ("327"). The four bottom solder mask (SMT) points should not be connected, hence their "*net name*" fields contain "N/C" indicating "*Not connected.*"

Each IPC-D-356 record consists of one line and (for historic reasons) is of a fixed length with a maximum of 80 characters.

The information extracted from translation of IPC-D-356 files to an internal netlist might include solder mask coverage and midpoint flags, along with the dimensions and location of connection points and their groupings into networks. Information in component and pin identification fields is not necessarily extracted for comparison purposes.

There are two amendments to IPC-D-356: IPC-D-356A and IPC-D-356B. Most of the extra information included in these formats is important for bare-board electrical testing and, unless buried passives are used, the information is generally not neede for CAD-to-CAM data consistency verification. If buried passives are used, it is better to transfer IPC-D-356A after verifying readability by the supplier.

A Net is a set of contiguous points connected by a conductive path. All the points in a net are electrically connected through traces and plated drills on the

Table 16: IPC-D-356 File Format

Columns	Coding	Comments
1	**P**	parameter
	C	comment
	3	test record
2	**1**	feature and through hole
	2	feature only (SMT, etc.)
	5	tooling feature (plate through holes)
	6	tooling feature (non plate through holes)
3	**7**	this file is using the standard TEST record format
4-17	*string*	netname or node number
	N/C	not connected, single point net, isolated point
18-20		blank for finished PCBs
21-26	*string*	reference designator, e.g. U32 or blank
	VIA	feature is a via
27	**-**	a dash to separate the ref designator from pin id
28-31	*string*	the pin identifier (pin number)
32	**M**	mid net point, otherwise blank or end of net
33	**D**	drilled hole, otherwise blank
34-37	**0000**	drilled hole diameter in 0.0001" or 0.001mm
38	**P**	plate-through hole
	U	unplated through hole
39	**A**	access code
40-41	**00**	test point accessible from both sides (PTH)
	01	test point accessible from side 1 (primary side)
	1n	test point accessible from side "n", where "n" is usually the last outer layer
42	**X**	x-location
43	+	or blank, x-location is positive
	-	x-location is negative
44-49	**nnnnnn**	six digits representing the x-coordinate, leading zeroes may be blank
50	**Y**	y-location
51	+	or blank, y-location is positive
	-	y-location is negative
52-57	**nnnnnn**	six digits representing the y-coordinate, leading zeroes may be blank
58	**X**	x-dimension feature size
59-62	**nnnn**	x-dimension
63	**Y**	y-dimension feature size
64-67	**nnnn**	y-dimension
68	**R**	counter-clockwise rotation of feature
69-71	**nnn**	3-digits represent rotation (000-360)
72		unassigned and blank
73	**S**	soldermask information, or blank
74	**0**	no solder mask
	1	specified primary side mask
	2	specified secondary side mask
	3	specifies both sides solder mask
75	**A**	alternate test record exists (or blank)
76-80	**nnnnn**	five-digit alphanumeric id representing the alternate test record

Table 17: IPC-D-356 File Header

Columns	Coding	Comments
1	**P**	parameter
2-3		blank
4-7		parameter name
8-9		blank
10-72		value
73-80		blank
4-7	**JOB**	the name of the job
10-72	*string*	the name of the job
4-7	**FORM**	`FIXED` or `VARIABLE`
10-72	`FIXED`	always `FIXED`
	`VARIABLE`	variable record format
4-7	**CODE**	indicate a switch to native language character set for comment records
10-72	`01`	code to use
4-7	**DIM**	data information module - for netlist data always `N`
10-72	`N`	
4-7	**UNITS**	units of measurement
10-72	`CUST`	same as `CUST 0`
	`CUST 0`	inches and degrees
	`CUST 1`	millimeters and degrees
	`CUST 2`	inches and radians
4-7	**TITLE**	title of data defined in this file
10-72	*string*	title value
4-7	**NUM**	part number of the data
10-72	*string*	par number value
4-7	**REV**	revision of the data
10-72	*string*	revision value
4-7	**LANG**	data exchange format (usually `SDEF`)
10-72	`SDEF`	data exchange format value
4-7	**SCALE**	scale factor of the data
10-72	`nnnnnnnn`	e.g. for 0.0001" `00010000`
19	`n`	
4-7	**TOL**	tolerance of the data
4-7	**LAYER**	layer specification for the data
4-7	**AREA**	desc. of board's rect boundaries
32-50	`(X+nnnnnn,Y+nnnnnn)`	xmin,ymin
51-70	`(X+nnnnnn,Y+nnnnnn)`	xmax,ymax
1-3	**999**	end of file

printed circuit board. What is a net? A Netlist is a list of nets stored in a data file which define the conductivity inter-connection scheme of a bare circuit board.

What is a Netlist? A netlist is a list of nets stored in data file which define the conductivity interconnection scheme of a bare circuit board. Comparison of customer netlist with the graphical data is imperative to ensure design integrity is being maintained during editing stages. Netlist provided by customer should be bare board netlist. Accepted Formats: IPC-D-356 IPC-D-356A Mentor Netlist Information

Sample IPC-D-356 Test File. The format in supplied IPC-D-356 and IPC-D-356A should complies with the description of the IPC-D-356 specification.

```
C Sample IPC-D-356A Test File
P JOB P023003A.ipc
P UNITS CUST
P TITLE P023003A.ipc
P NUM 23003
P REV 300
P VER IPC-D-356A
P IMAGE PRIMARY
367 - D1970UA00X-001590Y+055374X1970Y0000 S0
367 - D1970UA00X-001583Y+001438X1970Y0000 S0
327NET1 - M A01X-002235Y+037065X0150Y0000 S1
327NET1 - M A01X-002235Y+037235X0150Y0000 S1
317NET1 - M A01X-002150Y+037150X0250Y0000 S1
017NET1 - MD0160PA00X-002150Y+037150
327NET1 - M A02X-002150Y+037150X0135Y0000 S0
317NET1 - M A02X-002150Y+037150X0250Y0000 S2
017NET1 - MD0160PA00X-002150Y+037150
...
327NET2 - M A02X+040700Y+035450X0135Y0000 S0
317NET2 - M A02X+040700Y+035450X0250Y0000 S2
017NET2 - MD0160PA00X+040700Y+035450
327NET2 - A01X+040735Y+031350X0380Y0310 S0
327NET2 - M A01X+040818Y+035425X0150Y0000 S1
327NET2 - M A01X+040820Y+035450X0150Y0000 S1
...
078 X28671 X28748Y5845 X29898 X30245Y5498 Y5198
078 X30498Y4945 X31053 X31305Y5198 Y6597
078 X31503Y6795 X36998 X37245Y6548 Y4398
078 X37395Y4248 Y2753 X37048Y2405 X36503
078 X35805Y3103 Y3703 X35553Y3955 X33998
078 X33698Y3655 X32447 X31698Y2905 X29248
078 X28748Y2405 X26253 X26055Y2603 Y5317 X26918
078 X27083Y5483 Y5718 X26918Y5883 X26175 Y6018
078 X27153Y6995 X28298
999
```

F Photo Chemical Milling

PCM (Photo Chemical Machining) is the process of removing metal using chemicals. The photo chemical machining process works by masking off the areas where metal is not to be removed by accurately printing onto the surface, a photo resist material. Or by exposing the image of the required parts onto a light-sensitive film laminated on to the material. The material stock is placed into a chemical bath and an electric current applied to erode unmasked areas of the material.

The photo chemical etching process is able to cut very accurately (typically less than 15 microns tolerance) both in shape and depth. The photo chemical etching process also imparts no heat or distortion to the workpiece making [it] ideal for thin sheetmetals and delicate parts.

Photo chemical machining has low setup costs making it well suited to low volume runs or prototype parts. Another major advantage is the ability to partially etch or to include fold lines, holes, slots, mesh, text, and logos without significantly increasing the process time or cost.

Photo chemical machining is commonly used to cut finely detailed sheetmetal parts in almost any material. Some typical applications for photo chemical machining are: Solder paste stencil cutting. PCB Sheilding. Finely detailed, precision sheetmetal parts. Electrical contacts. Shims.

G IPC-7351 Naming Convention for Land Patterns

G.1 Surface Mount Land Patterns

Amplifiers ..... **AMP**_partno.
Ball Grid Array (1.27mm) ..... **BGA127P**+colsXrows-qty.
Ball Grid Array (1.50mm) ..... **BGA150P**+colsXrows-qty.
Ball Grid Array (1.00mm) ..... **BGA100P**+colsXrows-qty.
Ball Grid Array (0.80mm) ..... **BGA80P**+colsXrows-qty.
Ball Grid Array (0.75mm) ..... **BGA75P**+colsXrows-qty.
Ball Grid Array (0.65mm) ..... **BGA65P**+colsXrows-qty.
Ball Grid Array (0.50mm) ..... **BGA50P**+colsXrows-qty.
BGA staggered pins (1.27mm) .. **SBGA127P**+colsXrows-qty.
Battery ..... **BAT**_partno.
Capacitor, chip ..... **CAPC**+size
Capacitor, tantalum ..... **CAPT**+size
Capacitor, aluminium electrolytic **CAPAE**+dia+Wheight+H
Capacitor, variable ..... **CAPV**_partno.
Capacitor network, chip ..... **CAPN**_partno.
Capacitor, misc. ..... **CAP**_partno.
Ceramic flat pack ..... **CFP127P**+lead span-qty.
Column grid array ..... **CGA**+colsXrows-qty.
Crystal ..... **XTAL**_partno.
Diode, molded ..... **DIOSMA**, **B** & **C**
Diode, MELF ..... **DIOMELF**+size
Diode, misc. ..... **DIOB**_partno.
Diode, bridge rect. ..... **FB**_partno.
Ferrite bead ..... **FB**_partno.
Fiducial ..... **FID**pad sizeXmask size
Filter ..... **FIL**_partno.
Fuse ..... **FUSE**_partno.
Fuse, resettable ..... **FUSER**_partno.
Inductor, chip ..... **INDC**size
Inductor, molded ..... **INDM**size
Inductor, precision wire wound ..... **INDP**size
Inductor, misc. ..... **IND**_partno.
Inductor network, chip ..... **INDN**_partno.
Keypad ..... **KEYPAD**_partno.
LED ..... **LED**_partno.
LED, chip ..... **LED**size
Liquid crystal display ..... **LCD**_partno.
Microphone ..... **MIC**_partno.
Opto isolator ..... **OPTO**_partno.
Oscillator ..... **OSC**_partno.
PLCC, square ..... **PLCC-**qty.
PLCC, rectangular ..... **PLCCR-**qty.
PLCC, sq. socket ..... **PLCCS-**qty.
PLCC, rect. socket ..... **PLCCRS-**qty.
QFP (1.00mm) ..... **QFP100P**span1Xspan2-qty.
QFP (0.80mm) ..... **QFP80P**span1Xspan2-qty.
QFP (0.65mm) ..... **QFP65P**span1Xspan2-qty.
SQFP (0.50mm) ..... **SQFP50P**span1Xspan2-qty.
SQFP (0.40mm) ..... **SQFP40P**span1Xspan2-qty.
SQFP (0.30mm) ..... **SQFP30P**span1Xspan2-qty.
TQFP (0.80mm,≤ 1.60mm) .... **TQFP80P**span1Xspan2-qty.
TQFP (0.65mm,≤ 1.60mm) .... **TQFP65P**span1Xspan2-qty.
TQFP (0.50mm,≤ 1.60mm) ... **TSQFP50P**span1Xspan2-qty.
TQFP (0.40mm,≤ 1.60mm) ... **TSQFP40P**span1Xspan2-qty.
TQFP (0.30mm,≤ 1.60mm) ... **TSQFP30P**span1Xspan2-qty.
CQFP (1.27mm) ..... **CQFP127P**span-qty.
CQFP (0.80mm) ..... **CQFP80P**span-qty.
CQFP (0.635mm) ..... **CQFP635P**span-qty.
QFN (0.80mm) ..... **QFN80P-**widthXlen-qty.
QFN (0.65mm) ..... **QFN65P-**widthXlen-qty.
QFN (0.50mm) ..... **QFN50P-**widthXlen-qty.
QFN (0.40mm) ..... **QFN40P-**widthXlen-qty.
LCC ..... **LCC-**qty.

Relay ..... **RELAY**_partno.
Resistor, chip ..... **RESC**size
Resistor, molded ..... **RESM**size
Resistor, MELF ..... **RESMELF**size
Resistor network, chip ..... **RESN**_partno.
SOIC-J (1.27mm) .....
SOIC (1.27mm) .....
SOP (1.27mm) .....
SOP (1.00mm) .....
SOP (0.80mm) .....
SOP (0.65mm) .....
SOP (0.635mm) .....
SSOP (0.50mm) .....
SSOP (0.40mm) .....
SSOP (0.30mm) .....
TSOP (1.27mm) .....
TSOP (1.00mm) .....
TSOP (0.80mm) .....
TSOP (0.65mm) .....
TSSOP (0.55mm) .....
TSSOP (0.50mm) .....
TSSOP (0.40mm) .....
TSSOP (0.30mm) .....
.....
.....
.....
.....

References

258X. Generic requirements for printed board assembly products manufacturing description data and transfer methodology. Std. IPC-2581 with Amendment 1, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, March 2004. URL http://webstds.ipc.org/2581/-IPC-2581witham1pub.pdf.

BDTST. Sectional requirements for implementation of bare board product electrical testing data description [BDTST]. Std. IPC-2515A, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, November 2000. URL http://webstds.ipc.org/2515/2515Apub.pdf.

DPFv5. DPF format description. DPF Version 5, Barco Graphics NV, Electronic Tooling Systems division, Gent, Belgium, November 1998. URL http://www.photoplotstore.com/-pages/DPF_syntax.pdf.

DPFv7. DPF version 7 format description. DPF Version 7, Ucamco Software, Gent, Belgium, March 2009. URL http://-www.ucamco.com.

DXF. DXF reference. Manual AutoCAD 2008, Autodesk, January 2007. URL http://images.autodesk.com/adsk/files/-acad_dxf0.pdf.

EI-2010. Printed wiring board design guide. Guide Release 082710, Endicott Interconnect Technologies, Inc., Endicott, New York, August 2010. URL http://www.-endicottinterconnect.com/.

GenCAD. GenCAD specification 1.4. DPF Revision 1.4/A, Mitron, Portland, Oregon, November 1997. URL http://eletronix.ru/forum/index.php?s=-28e5ea69adde2003ccfb412fe72702c5&act=Attach&type=-post&id=15091.

GenCAM. Generic requirements for implementation of product manufacturing description data and transfer methodology. Std. IPC-2511A, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, March 2000. URL http://webstds.ipc.org/2511/2511Apub.pdf.

GenX. Generic requirements for implementation of product manufacturing description data and transfer XML schema methodology. Std. IPC-2511B, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, January 2002. URL http://webstds.ipc.org/2511/2511Bpub.pdf.

Tom Hausherr. The CAD library of the future. Whitepaper, PCB Libraries, Inc., February 2006. URL http://www.dnu.no/-arkiv1/The%20CAD%20Library%20of%20the%20Future.pdf.

IPC-2524. PWB fabrication data quality rating system. Std. IPC-2524, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, February 1999. URL http://-webstds.ipc.org/2524/IPC-2524.pdf.

IPC-2531. SMEMA standard recipe file format specification. Std. IPC-2531, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, March 1999. URL http://webstds.ipc.org/2531/2531pub_000.pdf.

IPC-2584. Sectional requirements for implementation of printed board fabrication data description. Std. IPC-2584, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, May 2007. URL http://webstds.ipc.org/-2584/IPC-2584.pdf.

IPC-2614. Sectional requirement for board fabrication documentation. Std. IPC-2614, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois. URL http://-www.ipc.org/.

IPC-7525. Stencil design guidelines. Std., Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, May 2000. URL http://www.wssi.com/gerbtool/-WISEGTDOC/IPC-7525(L).pdf.

IPC-7525A. Stencil design guidelines. Std. IPC-7525A Working Draft 1, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, February 2004. URL http://-jpk.jnrp.cn/dqx/xxwz/5xxzy/hybz4-IPC7525A.pdf.

IPC-NC-349. Computer numerical control formattting for driller and routers. Std. IPC-NC-349, Institute for Interconnecting and Packaging Electronic Circuits, Northbrook, Illinois, August 1985. URL http://www.ipc.org/TOC/IPC-NC-349.pdf.

Simon Jones. Open standards-based software tools optimize the PCB manufacturing enterprise. Whitepaper, GenRad, Inc., May 2001. URL http://www.teradyne.com/atd/resource/-docs/d2b/jones_openstandards.pdf.

David Kehmeier. Intermediate data format—mechanical data exchange specification for the design and analysis of printed wiring assemblies. IDF Version 2.0, Revision 3 (Final), Intermedius Design Integration, LLC, San Jose, January 5, 1993. URL http://www.simplifiedsolutions.com/images/-idf_v20_spec.pdf.

David Kehmeier. Intermediate data format—mechanical data exchange specification for the design and analysis of printed wiring assemblies. IDF Version 3.0, Revision 1, Mentor Graphics Corporation, San Jose, October 31, 1996. URL http://-www.simplifiedsolutions.com/images/idf_v30_spec.pdf.

David Kehmeier and Tom Makowski. Intermediate data format—version 4.0—data exchange specification for the mechanical design, electrical layout, and physical analysis of printed circuit assemblies. IDF Version 4.0, Revision A (Pre-implementation Draft), Intermedius Design Integration, LLC, San Jose, July 17, 1998. URL http://www.simplifiedsolutions.com/-images/idf_v40_spec.pdf.

Dae Hyun Kim and Myoung-Jun Kim. An extension of polygon clipping to resolve degenerate cases. In *Computer-Aided Design & Applications*, volume 3, pages 447–456, 2006. doi: 10.3722/cadaps.2006.447-456. URL http://citeseerx.ist.-psu.edu/viewdoc/download?doi=10.1.1.91.9804&rep=-rep1&type=pdf.

METRIX04. Design for manufacturability of rigid multi-layer boards. Manual Rev. B 11904, Merix Corporation, January 2004. URL http://www.merix.com/files/DFM%20Manual-%20rev%20b.pdf.

METRIX06. Controlled impedance designs. Technical report, Merix Corporation, September 2006. URL http://www.-merix.com/files/Controlled%20Impedance%20Designs.pdf.

Anirudh Modi. Hidden line removal using BSP tree and polygon clipping. Practical training report, Indian Institute of Technology, Bombay, 1996. URL http://www.anirudh.net/-practical_training/ptreport.pdf.

PCB-2010. IPC-7351B naming convetion for standard SMT land patterns. Technical Report IPC-7531B, PCB Matrix, Corp., January 2010. URL IPC-7x51%20&%20PCBM%20Land-%20Pattern%20Naming%20Convention.pdf.

Yu Peng, Jun-Hai Yong, Wei-Ming Dong, Hui Zhang, and Jia-Guang Sun. A new algorithm for Boolean operations on general polygons. *Computers & Graphics*, 29:57–70, 2005a. doi: 10.1016/j.cag.2004.11.001. URL http://liama.ia.ac.cn/wiki/_media/user:dong:dong_cg_05.pdf.

Yu Peng, Jun-Hai Yong, Hui Zhang, and Jiaguang Sun. Efficient algorithm for general polygon clipping. In *Proc. 6th Int'l Conf. on Computer-Aided Industrial Design and Conceptual Design 2005*, 2005b. URL http://hal.archives-ouvertes.fr/docs/00/51/76/87/PDF/YuPeng2005b.pdf.

Charles Robinson and Amit Verma. Optimizing test strategies during PCB design for boards with limited ICT access. Technical Report Version 1.0, Teradyne Inc., March 2002. URL http://www.teradyne.com/atd/resource/docs/d2b/optimteststrat_verma.pdf.

RS-274X/C. Gerber RS-274X format user's guide. RS-274X Revision C, Barco Gerber Systems, South Windsor, Connetiticut, USA, September 21, 1998. URL http://www.photoplotstore.com/pages/rs274xc.pdf.

RS-274X/D. Gerber RS-274X format user's guide. RS-274X Revision D, Barco Graphics NV, Gent, Belgium, March 1, 2001. URL http://www.artwork.com/274x/rs274xrevd_e.pdf.

Klamer Schutte. An edge labeling approach to concave polygon clipping, July 1995. URL http://clippoly.sourceforge.net/clip.pdf.

SEHO-2005. Design guidelines for selective wave soldering process. Technical report, SEHO Seitz & Hohnerlein GmbH, Kreuzwertheim, Germany, February 2005. URL http://www.tt-electronic.dk/pdf/Design-Guidelines.pdf.

SMEMA 1.2. SMEMA mechanical interface standard. SMEMA Std. 1.2, Surface Mount Equipment Manufacturers Association. URL http://smema.ipc.org/smema1.2.pdf.

SMEMA 3.1. SMEMA fiducial mark standard. SMEMA Std. 3.1, Surface Mount Equipment Manufacturers Association. URL http://smema.ipc.org/smema3.1.pdf.

SMEMA 4. SMEMA reflow terms and definitions. SMEMA Std. 4.0, Surface Mount Equipment Manufacturers Association. URL http://smema.ipc.org/smema4.pdf.

SMEMA 5. SMEMA screen printing terms and definitions. SMEMA Std. 5.0, Surface Mount Equipment Manufacturers Association. URL http://smema.ipc.org/smema5.pdf.

SMEMA 7. SMEMA fluid dispensing terms and definitions. SMEMA Std. 7.0, Surface Mount Equipment Manufacturers Association. URL http://smema.ipc.org/smema7.pdf.

SPEECTRA. SPEECTRA design language reference. SPEECTRA Product Version 10.0, Cadence Design Systems, Inc., San Jose, May 2000. URL http://www.eltm.ru/store/spdlr.pdf.

SRFFv1. SMEMA standard recipe file format specification. SMEMA Std. SRFF 1.0, Surface Mount Equipment Manufacturers Association. URL http://smema.ipc.org/SRFFv1.pdf.

SS-2007. Printed circuit board design for manufacturability guidelines. Document Number SOP-4-DFM01-B Revision H, Sanmina-SCI Corporation, 2007.

TI-2010. PCB design guidelines for 0.5mm package-on-package applications processors, part I. Application Report SPRABB3—June 2010, Texas Instruments, Dallas, TX, June 2010.

TP-101D. SMTA testability guidelines. SMTA Std. TP-101D 2008, Surface Mount Technology Association, Edina, Minneapolis, 2008. URL http://smema.ipc.org/.

Glen Walther. Tolerance analysis of Ohmega-Ply Resistors in multilayer PCB design. *CircuitTree*, January 2001.

Reiner Zoch and Christian Ott. Effects of an appropriate PCB layout and soldering nozzle design on quality and cost structure in selective soldering processes. Whitepaper, SEHO Systems GmbH, Kreuzwertheim, Germany, January 2009. URL http://www.circuitmart.com/pdf/pcb_layout_soldering_nozzle_design.pdf.

Acronyms

Following is a list of acronyms used throughout this document, with the acronym and its corresponding expansion:

ARE	Automatic Reverse Engineering
ASCII	American Standard Code for Information Interchange 2
BBV	Blind/Buried Via
BGA	Ball Grid Array
BGA	Ball-Grid Array
BOM	Bill Of Materials
BQFP	Quad Flat Package with Bumper Corners
CAD	Computer-Aided Design
CAM	Computer-Aided Manufacture
CBGA	Ceramic Ball Grid Array
CFP	Ceramic Flat Package
CGA	Column Grid Array
CNC	Computer Numerical Controlled
CPL	Component Placement List
CQFP	Ceramic Quad Flat Package
CR	Carriage Return
CSP	Chip Scale Package
CSV	Comma Separated Value
CSV	Comma Separated Values
CSWSC	Conformal Selective Wave Solder Carrier
DFA	Design for Assembly
DFM	Design for Manufacturability
DFM	Design for Manufacturing
DFN	Dual Flat No-lead
DFT	Design for Test
DFT	Design For Testability
DFX	Design for Excellence
DFX	Design For Excellence
DIP	Dual-in-Line Package
DIY	Do It Yourself
DNI	Do Not Install
DNS	Do Not Supply

DOD	Drop On Demand
DPF	Dynamic Process Format
DPI	Dots Per Inch
DRC	Design Rule Check
DXF	Drawing Exchange Format
EBCDIC	Extended Binary Coded Decimal Interchange Code
EBNF	Extended Baccus-Naur Form
ECAD	Electrical CAD
EDA	Electronic Design Automation
EIA	Electronic Industries Association
EMC	Electromagnetic Compatibility
EMI	Electromagnetic Interference
ENEPIG	Electroless Nickel/Electroless Palladium/Immersion Gold
ENIG	Electroless Nickel/Immersion Gold
ERC	Electrical Rule Check
ESR	Equivalent Series Resistance
FBGA	Fine-pitch Ball Grid Array
FEM	Finite-Element Modelling
FHS	Finished Hole Size
FPGA	Field Programmable Gate Array
GS	Gerber Systems Corporation
HASL	Hot Air Solder Leveling
HDI	High-Density Interconnect
HPGL	Hewlett-Packard Graphics Language
IAg	Immersion Silver
IC	Integrated Circuit
ICT	In Circuit Test
ICT	In-Circuit Test
IDF	Intermediate Data Format
IPC	Institute for Interconnecting and Packaging Electronic Circuits
IP	Intellectual Property
ISn	Immersion Tin
JITS	Just In Time Supply
LCC	Quad Leadless Ceramic Chip Carrier
LF	Line Feed
LGA	Land Grid Array
LPI	Liquid Photo-Imageable
MCAD	Mechanical CAD
MCV	Model Controller View
MCV	Model-Controller-View
MLCC	Multi-Layer Chip Capacitor
NC	Numerical Control
NDA	Non-Disclosure Agreement
NDA	Non-Disclosure Agreements
NMSD	Non-Solder Mask Defined
NPTH	Non-Plate-Through Hole
NRE	Non-Recurring Expense
NSMD	Non-Solder Mask Defined
NTD	Network Trace Data
NTD	Network Trace Description
OSP	Organic Solderability Preservative
PBGA	Plastic Ball Grid Array
PCB	Printed Circuit Board
PCI	Peripheral Computer Interconnect
PCM	Photo Chemical Machining
PDF	Portable Document Format
PLCC	Plastic Leaded Chip Carrier
PQFN	Pull-back Quad Flat Package
PSON	Pull-back Small Outline No-lead
PS	Post Script
PTF	Polymer Thick Film
PTH	Plate-Through Hole
PWA	Printed Wiring Assemblies
PWA	Printed Wiring Assembly
PWB	Printed Wiring Board
QFN	Quad Flat No-lead
QFN	Quad Flat Package No-Lead
QFP	Quad Flat Package
SCV	Sub-Composite Via
SFP	Small Formfactor Pluggable
SGML	Standard Generalized Markup Language
SMA	Surface Mount Adhesive
SMD	Solder Mask Defined
SMD	Surface-Mount Device
SMEMA	Surface Mount Equipment Manufacturers Association
SMOBC	Solder Mask Over Bare Copper
SMTA	Surface Mount Technology Association
SMT	Surface-Mount Technology
SOD	Small Outline Diode
SOIC	Small Outline IC
SOJ	Small Outline J-Lead
SON	Small Outline No-lead
SOP	Small Outline Package
SOT	Small Outline Transistor
SQFP	Shrink Quad Flat Package
SSOP	Shrink Small Outline Package
SVG	Scalable Vector Graphics
SWSC	Selective Wave Solder Carrier
SWS	Selective Wave Solder
TH	Through Hole
TMAG	Testability Management Action Group
TO	Generic DPAK
TQFP	Thin Quad Flat Package
TSOP	Thin Small Outline Package
TSQFP	Thin Shrink Quad Flat Package

TSSOP	Thin Shrink Small Outline Package
UL	Underwriter's Laboratory
UV	Ultraviolet
VBP	Via Beside Pad
VTP	Via Through Pad
WYSIWYG	What You See Is What You Get
XML	Extensible Markup Language
XYRS	XYRS Data

Glossary

Following is a glosary of terms and acronyms used throughout this document, with a brief description for each:

American Standard Code for Information Interchange 2 (ASCII): A character encoding scheme based on the ordering of the English alphabet.

Annular Ring: A ring of copper surrounding a drilled hole where it passes through a pad.

Aspect Ratio: The ratio of an aperture's breadth to depth. For a circular aperture, this is d/h, where d is the diameter of the cylindrical barrel and h is the height of the cylinder.

Automatic Reverse Engineering (ARE): A process used by assembly shops to determine the footprints and placement requirements for a PWA.

Back-drilled Via: A back-drilled via starts out life as a through via. In the mechanical drilling stage, a oversized drill is used to bore out the stub remaining on the unused backside of the via, forming a shorter via barrel. A back-drilled via is illustrated in *Fig. 21(30)*.

Ball-Grid Array (BGA): A component package that uses a grid of solder-balls for connection.

Ball-Grid Array (BGA): A type of packaging of electronic components.

Bill Of Materials (BOM): A list of materials used in the fabrication and assembly of a PWA.

Blind Via: A blind via is a via that traverses from an outer layer to an inner layer of a multilayer PCB. Blind vias can be formed by sub-laminate through drilling, depth-controlled drilling, or laser drilling. When less that or equal to 6mil hole size, a blind via is normally called a microvia. A blind via is illustrated in *Fig. 18(28)*.

Breakout: A situation where, due to positional tolerance, a drilled hole can extend beyond the boundary of a pad: thus breaking out from the pad. Breakout can be mitigated by placing a teardrop at the end of an adjoining trace.

B-Stage Laminate: An epoxy-glass laminate that is cured, but not yet hardened.

Buried Via: A buried via is a via that traverses from one inner layer of a PCB to another inner layer. Buried vias can be formed by sub-laminate through drilling. A buried via is illustrated in *Fig. 19(29)*.

Carriage Return (CR): An ASCII control character that causes a carriage return.

Ceramic Ball Grid Array (CBGA): A type of packaging of a BGA that includes a ceramic case.

Ceramic Quad Flat Package (CQFP): A type of QFP where the body is ceramic instead of plastic.

Chip Scale Package (CSP): A category of packaging of electronic components that includes BGA.

Comma Separated Value (CSV): A file format supported by popular spreadsheet programs.

Computer-Aided Design (CAD): A computer program that aides in the design of electronic or mechanical systems.

Computer-Aided Manufacture (CAM): A computer program for generating manufacturing programs and data from design information.

Computer Numerical Controlled (CNC): A type of NC machine that is controlled using a computer (instead of a puched tape).

Depth-control-drilled (DCD) Via: A depth-control-drilled via is a form of construction of a blind via where the drill machine drills to a controlled depth. A depth-control-drilled via is illustrated in *Fig. 22(31)*.

Design for Excellence (DFX): An approach to quality that incorporates all downstream considerations into design.

Design For Excellence (DFX): An extension of DRC and DFM for total quality.

Design for Manufacturability (DFM): An extension of DRC where additional (all) rules for manufacturability are checked (not just copper clearances).

Design For Testability (DFT): An approach to quality where an electronic design considers testability at the earliest stages.

Design Rule Check (DRC): A set of checks meant to ensure that a design is within a range of capabilities of manufacturing

Design Rule Check (DRC): Checking of design rules applied to the copper features of a PCB.

Do It Yourself (DIY): An approach to manufacturing where everything is done in the garage or basement.

Drawing Exchange Format (DXF): A computer file format developed by Autodesk for the exchange of drawing information (CAD data) between CAD programs [DXF].

Dynamic Process Format (DPF): An intended replacement for Gerber or RS-274X format developed by Barco [DPFv7].

Electrical CAD (ECAD): A CAD system that is tailored toward the design of electronic assemblies.

Electrical Rule Check (ERC): A check of the compatibility of the connections of component pins based on the characteristics of the pin function.

Electroless Nickel/Electroless Palladium/Immersion Gold (ENE A type of surface finish for a PCB where the exposed copper surfaces are first plated with nickel, then palladium, and finally gold, using a series of chemical baths.

Electroless Nickel/Immersion Gold (ENIG): A type of surface finish for a PCB where the exposed copper surfaces are first plated with nickel followed by gold using chemical baths

Electromagnetic Compatibility (EMC): A measure of the compatibility between systems as regards electromagnetic emissions.

Electromagnetic Interference (EMI): The interference caused by electromagnetic waves emanating from a culprit system interfering with a victim system.

Electronic Design Automation (EDA): A system of computer programs automating the electronic design process.

Electronic Industries Association (EIA): A standards generating industry association initially responsible for the standardization of electrical connectors.

Equivalent Series Resistance (ESR): When modelling a capacitor or inductor, the equivalent series resistance is the ESR.

Etch Compensation: The process of altering CAD widths of traces to compensate for the etch factor of the trace. The etching process is more rapid at the top of a trace than at the bottom of the trace. The resulting cross-section is roughly a trapezoid instead of a rectangle. Etch compensation is illustrated in *Fig. 25(34)*.

Extended Baccus-Naur Form (EBNF): A mathematical method of describing computer syntaxes.

Extended Binary Coded Decimal Interchange Code (EBCDIC): An 8-bit character encoding (code page) developed by IBM and used by initially by IBM mainframe and midrange computer operating systems.

Extensible Markup Language (XML): A descriptive syntax based on SGML that is used to describe the content and form of a document.

Fiducial Marks: Sometimes simply called *fiducials*, are marks placed on a production panel or board to assist vision or machine vision systems in establishing a reference plane for the panel, board, or locality on a board.

Field Programmable Gate Array (FPGA): A reprogrammable digital device.

Fillet: A trapezoidal copper feature added to the end of a tracek where it meets a pad to mitigate breakout of a drilled hole from the pad at the angle of the track. Also called a teardrop. A fillet is one of the forms of construction of a teardrop.

Fine-pitch Ball Grid Array (FBGA): A type of packaging of a BGA that has fine-pitched balls.

Finished Hole Size (FHS): The size of a hole specified as the diameter of the hole after finishing: that is after plating and surface finish has been applied.

Finite-Element Modelling (FEM): An approach to modelling electromagnetic waves using an approximation based on a three-dimensional array of discrete elements.

Gerber Systems Corporation (GS): The developer of the Gerber photoplotter format.

Hewlett-Packard Graphics Language (HPGL): A graphics language and file format developed by Hewlett-Packard for pen plotters and still supported on laser-printers.

High-Density Interconnect (HDI): A method of interconnecting high-density PCBs using microvias on the outer layers of the board.

Hole Filling: The filling of holes in a PCB with conductive or non-conductive fillings (such as resin) and possibly planarizing and plating over the resulting construction.

Hot Air Solder Leveling (HASL): A process for applying finish to the exposed copper surfaces of a PCB where the board is dipped in molten solder and then leveled with a blast of hot air.

Immersion Silver (IAg): A printed circuit board surface finish with excellent planarity and high-frequency handling.

Immersion Tin (ISn): A printed circuit board surface finish with excellent planarity and low-friction for press-fit pins.

In Circuit Test (ICT): A method of testing a PWA whereby a test fixture is held by vacuum to a board connected in a circuit.

In-Circuit Test (ICT): A testing procedure for testing an assembled board using a specialized one-side or clam shell fixture of probes. The fixture is retained with a vacuum, requiring that via holes be capped, plugged or filled to afford a seal for the vacuum.

Institute for Interconnecting and Packaging Electronic Circuits An industry association providing standards for PCB fabrication and assembly.

Intermediate Data Format (IDF): A file format developed by Mentor Graphics that provides for the exchange of basic electronic and mechanical information between an ECAD and MCAD. See *Sec. 7.11.1(141)*.

Keyhole: A circular copper feature added to the end of a track where it meets a pad to mitigate breakout of a drilled hole from the pad at the angle of the track Another name for a keyhole is a snowman. Also generically called a teardrop. A keyhole is one of the forms of construction for a teardrop.

Laser-Blind Via: A blind via that is fabricated using laser drilling: also commonly a *Microvia*.

Line Feed (LF): An ASCII control character that causes a line feed.

Liquid Photo-Imageable (LPI): A type of soldermask where a photographic image is used to control the polymerization of the mask.

Machine-Blind Via: A blind via that is fabricated using machine drilling. As the limit of machine drilling is approximately 6mil, it takes a laser to generate a microvia.

Mechanical CAD (MCAD): A CAD system that is tailored toward the design of mechanical or thermodynamic assemblies.

Microvia: A microvia is a small via, less than 6mil in size, normally laser-drilled, but can be chemically ethched or machine drilled, that traverses through one or two outermost HDI layers on a multilayer PCB. A microvia is illustrated in *Fig. 20(29)*.

Model-Controller-View (MCV): A computer program design approach that separates data and tasks into a model, a controller and a view.

Multi-Layer Chip Capacitor (MLCC): A type of SMT capacitor that is built using multiple layers of conductor on a ceramic chip.

Network Trace Description (NTD): An auxiliary file and format to the IPC-D-356 netlist format developed by Veratim that describes the traces on a PCB for board testing. This format was eventually absorbed into IPC356A.

Non-Disclosure Agreement (NDA): A legal agreement whereby the parties agree not to disclose confidential information.

Non-Plate-Through Hole (NPTH): A type of hole in a PCB that is not plated, in contrast to a PTH.

Non-Solder Mask Defined (NSMD): One of two mechanisms for describing a pad: NSMD is a pad that is not defined by solder mask, that is, it has a solder mask opening that is greater than the pad shape. The pad is defined by the copper extents of the pad. This is in contrast to soldermask defined pads which is a technique for describing a pad on a plane using only its solder mask opening. In this case the opening is the exact size of the pad.

Numerical Control (NC): An approach to controlling production machinery from numbers on a punched tape or in an electronic file.

Organic Solderability Preservative (OSP): A type of surface finish for a PCB where the exposed copper surfaces are coated with a very thin organic substance.

Peripheral Computer Interconnect (PCI): A standard for the interconnection of computer add-in cards with computer motherboards.

Photo Chemical Machining (PCM): A process of using photoimagable chemical masks and corrosive chemicals to mill stock.

Plastic Ball Grid Array (PBGA): A type of packaging of a BGA that includes a plastic case.

Plate-Through Hole (PTH): A type of hole in a PCB with a copper-plated barrel.

Plating Thieving: The process of adding neutral copper zones or grid patterns to an external layer of a single or multilayer PCB in an effort to balance the local concentration of active plating chemistry and protect against over-plating in some areas and under-plating in others.

Portable Document Format (PDF): A widely used paperless document format developed by Adobe Systems.

Post Script (PS): A Forth-language based printer language developed by Adobe Systems. The language is also the basis of the PDF.

Prepreg: Epoxy-glass laminate that is in the B-stage: that is, cured but not hardened. The laminate does not harden until heated in a press.

Press Factor: The difference in thickness between a B-stage laminate (prepreg) before and after lamination (compression). The finished thickness of the laminate is dependent upon the copper coverage of the copper layers being pressed into the prepreg. Press factor is illustrated in *Fig. 45(54)*.

Printed Circuit Board (PCB): An electronic circuit fabricated by printing the electrical connections on copper layers. Also referred to as PWB.

Printed Wiring Assembly (PWA): A PWB resplendent with attached electrical and mechanical components, including any retaining fixtures.

Printed Wiring Board (PWB): An electronic circuit fabricated by printing the electrical connections on copper layers. Also referred to as PCB.

Quad Flat Package No-Lead (QFN): A type of packaging of electronic components that features contacts emanating to four quadrants.

Quad Flat Package (QFP): A type of packaging of electronic components that features leads extending in four quadrants.

Quad Leadless Ceramic Chip Carrier (LCC): A type of integrated circuit package.

Resin Thieving: See *Venting*.

scratching: The removal of portions of dark areas already exposed to light by scratching off the emulsion in those areas.

Shrink Quad Flat Package (SQFP): A type of QFP that is smaller than normal.

Shrink Small Outline Package (SSOP): An integrated circuit package.

Small Outline Diode (SOD): An integrated circuit package.

Small Outline IC (SOIC): An integrated circuit package.

Small Outline Package (SOP): An integrated circuit package.

Small Outline Transistor (SOT): An integrated circuit package.

Snowman: A circular copper feature added to the end of a track where it meets a pad to mitigate breakout of a drilled hole from the pad at the angle of the track. Another name for a snowman is a keyhole. Also generically called a teardrop. A snowman is one of the forms of construction of a teardrop.

Solder Mask Defined (SMD): A technique whereby a pad is defined by its solder mask opening rather than its copper outline. SMD pads are necessary where the copper pad is flooded with copper on one or more sides.

Standard Generalized Markup Language (SGML): A document markup language developed by IBM for the US government that is still in use. It is the basis of XML.

Sub-Composite Via (SCV): A method of contructing interconnecting staggered vias by connecting a microvia in an HDI layer with a via buried through a core. The connection uses a special filleted shape.

Surface Mount Adhesive (SMA): An adhesive used to retain SMT components on a PCB during wave soldering or double-sided two-stage reflow.

Surface-Mount Device (SMD): A device that is mounted to the surface of a PCB using solder-paste and reflow processes. The acronym SMT is often used in preference to SMD, because SMD can also mean solder-mask-defined when referring to lands or pads.

Surface Mount Equipment Manufacturers Association (SMEMA A non-profit organization of companies manufacturing equipment or producing software for surface mount board production.

Surface-Mount Technology (SMT): A technology for mounting PCB components flush with the board surface.

Teardrop: A teardrop is a copper feature added to the end of track where it meets a pad to mitigate breakout of a drilled hole from the pad at the angle of the track. Teardrop refers to the technique in general, as well as one of the forms of construction.

Thin Quad Flat Package (TQFP): A type of QFP that is shorter than 1.60mm.

Thin Shrink Small Outline Package (TSSOP): An integrated circuit package.

Thin Small Outline Package (TSOP): An integrated circuit package.

Through Hole (TH): A type of hole in a PCB intended on containing a component lead.

Through Via: A through via is a via that passes completely through a multilayer PCB from top to bottom outer layers. It is the most common construction for vias. A through via is illustrated in *Fig. 17(27)*.

Venting: The process of adding neutral copper zones or grid patterns to an internal layer of a multilayer PCB in an effort to balance the distribution of resin in B-stage epoxy laminate material (prepreg) and protect against excessive bow and twist of the resulting laminated board. Also called "Resin Thieving".

Via Beside Pad (VBP): A via construction that is situated beside a pad. That is, the via is in closer proximity to a pad than would normally be allowed. These vias must normally be filled when larger than a given size or aspect ratio.

Via Plugging: Also known as via capping, via plugging is the process of plugging or capping vias from one side of the board, for the purpose of avoiding solder wicking, particularly for vias under dense BGA fields, or for creating a vacuum seal for ICT.

Via Through Pad (VTP): A type of via construction where the via passes through a pad on the PCB.

Via Through Pad (VTP): A via construction that passes through a pad. These vias must normally be filled and plated over when larger than a given size or aspect ratio.

What You See Is What You Get (WYSIWYG): An approach to graphical user interfaces (particularly word processors) where the displayed items are as close as possible to the final (normally printed) result.

XYRS Data (XYRS): A data file consisting of x-coordinate data (X), y-coordinate data (Y), rotation (R), and side (S), for the placement of SMT components. Also called a CPL (Component Placement List), Pick-and-Place file, or simply XY data.

Index