


PCB Design Tutorial



by David L. Jones

Email: david AT alternatezone DOT com

Revision A - June 29th 2004

The latest version of this tutorial can be found through www.alternatezone.com

Freely distributable for educational and personal use.

Copyright© 2004 David L. Jones

Introduction.....	3
The Old Days.....	3
PCB Packages	3
Standards.....	3
The Schematic.....	4
Imperial and Metric.....	4
Working to Grids	5
Working from the top	6
Tracks.....	6
Pads	7
Vias	8
Polygons	8
Clearances	8
Component Placement & Design	9
Basic Routing	11
Finishing Touches	13
Single Sided Design	13
Double Sided Design	14
Other Layers	14
Silkscreen	14
Solder Mask	14
Mechanical Layer.....	15
Keepout	15
Layer Alignment	15
Netlists	15
Rats Nest.....	16
Design Rule Checking	16
Forward and Back Annotation.....	17
Multi layer Design	17
Power Planes	18
Good Grounding.....	19
Good Bypassing	19
High Frequency Design Techniques	19
Double Sided Loading.....	20
Auto Routing.....	20
Auto Placement	21
Design For Manufacturing.....	21
Panelisation.....	21
Tooling Strips.....	21
Fiducial Marks	22
Thermal Relief.....	22
Soldering.....	22
Basic PCB Manufacture.....	23
Surface Finishies	24
Electrical Testing	24
Signature.....	24
Submitting your design for manufacture.....	25

Introduction

You've designed your circuit, perhaps even bread boarded a working prototype, and now it's time to turn it into a nice Printed Circuit Board (PCB) design. For some designers, the PCB design will be a natural and easy extension of the design process. But for many others the process of designing and laying out a PCB can be a very daunting task. There are even very experienced circuit designers who know very little about PCB design, and as such leave it up to the "expert" specialist PCB designers. Many companies even have their own dedicated PCB design departments. This is not surprising, considering that it often takes a great deal of knowledge and talent to position hundreds of components and thousands of tracks into an intricate (some say artistic) design that meets a whole host of physical and electrical requirements. Proper PCB design is very often an integral part of a design. In many designs (high speed digital, low level analog and RF to name a few) the PCB layout may make or break the operation and electrical performance of the design. It must be remembered that PCB traces have resistance, inductance, and capacitance, just like your circuit does.

This article is presented to hopefully take some of the mystery out of PCB design. It gives some advice and "rules of thumb" on how to design and lay out your PCBs in a professional manner. It is, however, quite difficult to try and "teach" PCB design. There are many basic rules and good practices to follow, but apart from that PCB design is a highly creative and individual process. It is like trying to teach someone how to paint a picture. Everyone will have their own unique style, while some people may have no creative flair at all!

Indeed, many PCB designers like to think of PCB layouts as works of art, to be admired for their beauty and elegance. "If it looks good, it'll work good." is an old catch phrase.

Lets have a go shall we...

The Old Days

Back in the pre-computer CAD days, PCBs were designed and laid out by hand using adhesive tapes and pads on clear drafting film. Many hours were spent slouched over a fluorescent light box, cutting, placing, ripping up, and routing tracks by hand. Bishop Graphics, Letraset, and even Dalo pens will be names that evoke fond, or not so fond memories. Those days are well and truly gone, with computer based PCB design having replaced this method completely in both hobbyist and professional electronics. Computer based CAD programs allow the utmost in flexibility in board design and editing over the traditional techniques. What used to take hours can now be done in seconds.

PCB Packages

There are many PCB design packages available on the market, a few of which are freeware, shareware, or limited component full versions. Protel is the defacto industry standard package in Australia. Professionals use the expensive high end Windows based packages such as 99SE and DXP. Hobbyists use the excellent freeware DOS based Protel AutoTrax program, which was, once upon a time, the high-end package of choice in Australia. Confusingly, there is now another Windows based package also called AutoTrax EDA. This is in no way related to the Protel software.

This article does not focus on the use of any one package, so the information can be applied to almost any PCB package available. There is however, one distinct exception. Using a PCB only package, which does not have schematic capability, greatly limits what you can do with the package in the professional sense. Many of the more advanced techniques to be described later require access to a compatible schematic editor program. This will be explained when required.

Standards

There are industry standards for almost every aspect of PCB design. These standards are controlled by the former Institute for Interconnecting and Packaging Electronic Circuits, who are now known simply as the IPC (www.ipc.org). There is an IPC standard for every aspect of PCB design, manufacture, testing, and anything else that you could ever need. The major document that covers PCB design is IPC-2221, "Generic Standard on

Printed Board Design". This standard superseded the old IPC-D-275 standard (also Military Std 275) which has been used for the last half century.

Local countries also have their own various standards for many aspects of PCB design and manufacture, but by and large the IPC standards are the accepted industry standard around the world.

Printed Circuit Boards are also known (some would say, more correctly known) as Printed Wiring Boards, or simply Printed Boards. But we will settle on the more common term PCB for this article.

The Schematic

Before you even begin to lay out your PCB, you **MUST** have a complete and accurate schematic diagram. Many people jump straight into the PCB design with nothing more than the circuit in their head, or the schematic drawn on loose post-it notes with no pin numbers and no order. This just isn't good enough, if you don't have an accurate schematic then your PCB will most likely end up a mess, and take you twice as long as it should.

"Garbage-in, garbage-out" is an often used quote, and it can apply equally well to PCB design. A PCB design is a manufactured version of your schematic, so it is natural for the PCB design to be influenced by the original schematic. If your schematic is neat, logical and clearly laid out, then it really does make your PCB design job a lot easier. Good practice will have signals flowing from inputs at the left to outputs on the right. With electrically important sections drawn correctly, the way the designer would like them to be laid out on the PCB. Like putting bypass capacitors next to the component they are meant for. Little notes on the schematic that aid in the layout are very useful. For instance, "this pin requires a guard track to signal ground", makes it clear to the person laying out the board what precautions must be taken. Even if it is you who designed the circuit and drew the schematic, notes not only remind yourself when it comes to laying out the board, but they are useful for people reviewing the design.

Your schematic really should be drawn with the PCB design in mind.

It is outside the scope of this article to go into details on good schematic design, as it would require a complete article in its own right.

Imperial and Metric

The first thing to know about PCB design is what measurement units are used and their common terminologies, as they can be awfully confusing!

As any long time PCB designer will tell you, you should always use imperial units (i.e. inches) when designing PCBs. This isn't just for the sake of nostalgia, although that is a major reason! The majority of electronic components were (and still are) manufactured with imperial pin spacing. So this is no time to get stubborn and refuse to use anything but metric units, metric will make laying out of your board a lot harder and a lot messier. If you are young enough to have been raised in the metric age then you had better start learning what inches are all about and how to convert them.

An old saying for PCB design is "thou shall use thous". A tad confusing until you know what a "thou" is.

A "thou" is 1/1000th of an inch, and is universally used and recognised by PCB designers and manufacturers everywhere. So start practicing speaking in terms of "10 thou spacing" and "25 thou grid", you'll sound like a professional in no time!

Now that you understand what a thou is, we'll throw another spanner in the works with the term "mil" (or "mils"). 1 "mil" is the same as 1 thou, and is NOT to be confused with the millimeter (mm), which is often spoken the same as "mil". The term "mil" comes from 1 thou being equal to 1 mili inch. As a general rule avoid the use of "mil" and stick to "thou", it's less confusing when trying to explain PCB dimensions to those metricated non-PCB people.

Some PCB designers will tell you not to use metric millimeters for ANYTHING to do with a PCB design. In the practical world though, you'll have to use both imperial inches (thous) and the metric millimeter (mm). So which units do you use for what? As a general rule, use thous for tracks, pads, spacings and grids, which are most of your basic "design and layout" requirements. Only use mm for "mechanical and manufacturing" type requirements like hole sizes and board dimensions.

You will find that many PCB manufacturers will follow these basic guidelines also, for when they ask you to provide details for a quote to manufacture your board. Most manufacturers use metric size drills, so specifying imperial size holes really is counterproductive and can be prone to errors.

Just to confuse the issue even further, there are many components (new surface mount parts are an example) which have metric pin spacing and dimensions. So you'll often have to design some component footprints using metric grids and pads. Many component datasheets will also have metric dimensions even though the spacing are designed to an imperial grid. If you see a "weird" metric dimension like 1.27mm in a component, you can be pretty sure it actually has a nice round imperial equivalent. In this case 1.27mm is 50 thou.

Yes, PCB design can be confusing!

So whatever it is you have to do in PCB design you'll need to become an expert at imperial to metric conversion, and vice-versa. To make your life easier though, all the major PCB drafting packages have a single "hot key" to convert between imperial and metric units instantly ("Q" on Protel for instance). It will help you greatly if you memorise a few key conversions, like **100 thou (0.1 inch) = 2.54mm, and 200 thou (0.2 inch) = 5.08mm** etc

Values of 100 thou and above are very often expressed in inches instead of thous. So 0.2 inch is more commonly used than 200thou.

1 inch is also commonly known as 1 "pitch". So it is common to hear the phrase "0.1 inch pitch", or more simply "0.1 pitch" with the inches units being assumed. This is often used for pin spacing on components.

100 thou is a basic "reference point" for all aspects of PCB design, and a vast array of common component lead spacing are multiples or fractions of this basic unit. 50 and 200 thou are the most common.

Along with the rest of the world, the IPC standards have all been metricated, and only occasionally refer to imperial units. This hasn't really converted the PCB industry though. Old habits die hard, and imperial still reigns supreme in many areas of practical usage.

Working to Grids

The second major rule of PCB design, and the one most often missed by beginners, is to lay out your board on a fixed grid. This is called a "snap grid", as your cursor, components and tracks will "snap" into fixed grid positions. Not just any size grid mind you, but a fairly coarse one. 100 thou is a standard placement grid for very basic through hole work, with 50 thou being a standard for general tracking work, like running tracks between through-hole pads. For even finer work you may use a 25 thou snap grid or even lower. Many designers will argue over the merits of a 20 thou grid vs a 25 thou grid for instance. In practice, 25 thou is often more useful as it allows you to go exactly half way between 50 thou spaced pads.

Why is a coarse snap grid so important? It's important because it will keep your components neat and symmetrical; aesthetically pleasing if you may. It's not just for aesthetics though - it makes future editing, dragging, movement and alignment of your tracks, components and blocks of components easier as your layout grows in size and complexity.

A bad and amateurish PCB design is instantly recognisable, as many of the tracks will not line up exactly in the center of pads. Little bits of tracks will be "tacked" on to fill in gaps etc. This is the result of not using a snap grid effectively.

Good PCB layout practice would involve you starting out with a coarse grid like 50 thou and using a progressively finer snap grid if your design becomes "tight" on space. Drop to 25 thou and 10 thou for finer routing and placement when needed. This will do 99% of boards. Make sure the finer grid you choose is a nice even division of your standard 100 thou. This means 50, 25, 20, 10, or 5 thou. Don't use anything else, you'll regret it.

A good PCB package will have hotkeys or programmable macro keys to help you switch between different snap grid sizes instantly, as you will need to do this often.

There are two types of grids in a PCB drafting package, a snap grid as discussed, and a "visible" grid. The visible grid is an optional on-screen grid of solid or dashed lines, or dots. This is displayed as a background behind your design and helps you greatly in lining up components and tracks. You can have the snap grid and visible grid set to different units (metric or imperial), and this is often very helpful. Many designers prefer a 100 thou visible grid and rarely vary from that.

Some programs also have what is called an “Electrical” grid. This grid is not visible, but it makes your cursor “snap” onto the center of electrical objects like tracks and pads, when your cursor gets close enough. This is extremely useful for manual routing, editing and moving objects.

One last type of grid is the “Component” grid. This works the same as the snap grid, but it’s for component movement only. This allows you to align components up to a different grid. Make sure you make it a multiple of your Snap grid.

When you start laying out your first board, snap grids can feel a bit “funny”, with your cursor only being able to be moved in steps. Unlike normal paint type packages which everyone is familiar with. But it’s easy to get used to, and your PCB designs will be one step closer to being neat and professional.

Working from the top

PCB design is always done looking from the top of your board, looking through the various layers as if they were transparent. This is how all the PCB packages work. The only time you will look at your board from the bottom is for manufacturing or checking purposes. This “through the board” method means that you will have to get used to reading text on the bottom layers as a mirror image, get used to it!

Tracks

There is no recommended standard for track sizes. What size track you use will depend upon (in order of importance) the electrical requirements of the design, the routing space and clearance you have available, and your own personal preference. Every design will have a different set of electrical requirements which can vary between tracks on the board. All but basic non-critical designs will require a mixture of track sizes. As a general rule though, the bigger the track width, the better. Bigger tracks have lower DC resistance, lower inductance, can be easier and cheaper for the manufacturer to etch, and are easier to inspect and rework.


The lower limit of your track width will depend upon the “track/space” resolution that your PCB manufacturer is capable of. For example, a manufacturer may quote a 10/8 track/space figure. This means that tracks can be no less than 10 thou wide, and the spacing between tracks (or pads, or any part of the copper) can be no less than 8 thou. The figures are almost always quoted in thou’s, with track width first and then spacing.

Real world typical figures are 10/10 and 8/8 for basic boards. The IPC standard recommends 4thou as being a lower limit. Once you get to 6thou tracks and below though, you are getting into the serious end of the business, and you should be consulting your board manufacturer first. The lower the track/space figure, the greater care the manufacturer has to take when aligning and etching the board. They will pass this cost onto you, so make sure that you don’t go any lower than you need to. As a guide, with “home made” PCB manufacturing processes like laser printed transparencies and pre-coated photo resist boards, it is possible to easily get 10/10 and even 8/8 spacing.

Just because a manufacturer can achieve a certain track/spacing, it is no reason to “push the limits” with your design. Use as big a track/spacing as possible unless your design parameters call for something smaller.

As a start, you may like to use say 25 thou for signal tracks, 50 thou for power and ground tracks, and 10-15 thou for going between IC and component pads. Some designers though like the “look” of smaller signal tracks like 10 or 15 thou, while others like all of their tracks to be big and “chunky”. Good design practice is to keep tracks as big as possible, and then to change to a thinner track only when required to meet clearance requirements.

Changing your track from large to small and then back to large again is known as “necking”, or “necking down”. This is often required when you have to go between IC or component pads. This allows you to have nice big low impedance tracks, but still have the flexibility to route between tight spots.



In practice, your track width will be dictated by the current flowing through it, and the maximum temperature rise of the track you are willing to tolerate. Remember that every track will have a certain amount of resistance, so the track will dissipate heat just like a resistor. The wider the track the lower the resistance. The thickness of the copper on your PCB will also play a part, as will any solder coating finish.

The thickness of the copper on the PCB is nominally specified in ounces per square foot, with 1oz copper being the most common. You can order other thicknesses like 0.5oz, 2oz and 4oz. The thicker copper layers are useful for high current, high reliability designs.

The calculations to figure out a required track width based on the current and the maximum temperature rise are a little complex. They can also be quite inaccurate, as the standard is based on a set of non-linear graphs based on measured data from around half a century ago. These are still reproduced in the IPC standard.

A handy track width calculator program can be found at www.ultracal.com/calc.htm, and gives results based on the IPC graphs.

As a rule of thumb, a 10degC temperature rise in your track is a nice safe limit to design around. A handy reference table has been included in this article to give you a list of track widths vs current for a 10degC rise. The DC resistance in milli ohms per inch is also shown. Of course, the bigger the track the better, so don't just blindly stick to the table.

Track Width Reference Table (for 10deg C temp rise). Track Width is in Thous (mils)			
Current (Amps)	Width for 1oz	Width for 2 oz	milli Ohms/Inch
1	10	5	52
2	30	15	17.2
3	50	25	10.3
4	80	40	6.4
5	110	55	4.7
6	150	75	3.4
7	180	90	2.9
8	220	110	2.3
9	260	130	2.0
10	300	150	1.7

Pads

Pad sizes, shapes and dimensions will depend not only upon the component you are using, but also the manufacturing process used to assemble the board, among other things. There are a whole slew of standards and theories behind pad sizes and layouts, and this will be explained later. Suffice it to say at this stage that your PCB package should come with a set of basic component libraries that will get you started. For all but the simplest boards though, you'll have to modify these basic components to suit your purpose. Over time you will build up your own library of components suitable for various requirements.

There is an important parameter known as the pad/hole ratio. This is the ratio of the pad size to the hole size. Each manufacturer will have their own minimum specification for this. As a simple rule of thumb, the pad should be at least 1.8 times the diameter of the hole, or at least 0.5mm larger. This is to allow for alignment tolerances on the drill and the artwork on top and bottom layers. This ratio gets more important the smaller the pad and hole become, and is particularly relevant to vias.

There are some common practices used when it comes to generic component pads. Pads for leaded components like resistors, capacitors and diodes should be round, with around 70 thou diameter being common. Dual In Line (DIL) components like IC's are better suited with oval shaped pads (60 thou high by 90-100 thou wide is common). Pin 1 of the chip should always be a different pad shape, usually rectangular, and with the same dimensions as the other pins.

Most surface mount components use rectangular pads, although surface mount SO package ICs should use oval pads. Again, with pin 1 being rectangular.

Other components that rely on pin numbering, like connectors and SIP resistor packs, should also follow the "rectangular pin 1" rule.

Octagonal pads are seldom used, and should generally be avoided.

As a general rule, use circular or oval pads unless you need to use rectangular.

Vias

Vias connect the tracks from one side of your board to another, by way of a hole in your board. On all but cheap home made and low end commercial prototypes, vias are made with electrically plated holes, called Plated Through Holes (PTH). Plated through holes allow electrical connection between different layers on your board.

What is the difference between a via and a pad? Practically speaking there is no real difference, they are both just electrically plated holes. But there are differences when it comes to PCB design packages. Pads and Vias are, and should be, treated differently. You can globally edit them separately, and do some more advanced things to be discussed later. So don't use a pad in place of a via, and vice-versa.


Holes in vias are usually a fair bit smaller than component pads, with 0.5-0.7mm being typical.

Using a via to connect two layers is commonly called "stitching", as you are effectively electrically stitching both layers together, like threading a needle back and forth through material. Throw the term stitching a few times into a conversation and you'll really sound like a PCB professional!

Polygons

"Polygons" are available on many PCB packages. A polygon automatically fills in (or "floods") a desired area with copper, which "flows" around other pads and tracks. They are very useful for laying down ground planes. Make sure you place polygons after you have placed all of your tracks and pads.

Polygon can either be "solid" fills of copper, or "hatched" copper tracks in a crisscross fashion. Solid fills are preferred, hatched fills are basically a thing of the past.



An example of a "Solid Polygon Fill" (Left), and a "Hatched Polygon Fill" (Right)

Clearances

Electrical clearances are an important requirement for all boards. Too tight a clearance between tracks and pads may lead to "hairline" shorts and other etching problems during the manufacturing process. These can be very hard to fault find once your board is assembled. Once again, don't "push the limits" of your manufacturer unless you have to, stay above their recommended minimum spacing if at all possible.

At least 15 thou is a good clearance limit for basic through hole designs, with 10 thou or 8 thou being used for more dense surface mount layouts. If you go below this, it's a good idea to consult with your PCB manufacturer first.

For 240V mains on PCB's there are various legal requirements, and you'll need to consult the relevant standards if you are doing this sort of work. As a rule of thumb, an absolute minimum of 8mm (315 thou) spacing should be allowed between 240V tracks and isolated signal tracks. Good design practice would dictate that you would have much larger clearances than this anyway.

For non-mains voltages, the IPC standard has a set of tables that define the clearance required for various voltages. A simplified table is shown here. The clearance will vary depending on whether the tracks are on an internal layers or the external surface. They also vary with the operational height of the board above sea level,

due to the thinning of the atmosphere at high altitudes. Conformal coating also improves these figures for a given clearance, and this is often used on military spec PCBs.

Clearances for Electrical Conductors			
Voltage (DC or Peak AC)	Internal	External (<3050m)	External (>3050m)
0-15V	0.05mm	0.1mm	0.1mm
16-30V	0.05mm	0.1mm	0.1mm
31-50V	0.1mm	0.6mm	0.6mm
51-100V	0.1mm	0.6mm	1.5mm
101-150V	0.2mm	0.6mm	3.2mm
151-170V	0.2mm	1.25mm	3.2mm
171-250V	0.2mm	1.25mm	6.4mm
251-300V	0.2mm	1.25mm	12.5mm
301-500V	0.25mm	2.5mm	12.5mm

Component Placement & Design

An old saying is that PCB design is 90% placement and 10% routing. Whilst the actual figures are of no importance, the concept that component placement is by far the most important aspect of laying out a board certainly holds true. Good component placement will make your layout job easier and give the best electrical performance. Bad component placement can turn your routing job into a nightmare and give poor electrical performance. It may even make your board unmanufacturable. So there is a lot to think about when placing components!

Every designer will have their own method of placing components, and if you gave the same circuit (no matter how simple) to 100 different experienced designers you'd get a 100 different PCB layouts every time. So there is no absolute right way to place your components. But there are quite a few basic rules which will help ease your routing, give you the best electrical performance, and simplify large and complex designs.

At this point it is a good idea to give you an idea of the basic steps required to go about laying out a complete board:

- Set your snap grid, visible grid, and default track/pad sizes.
- Throw down all the components onto the board.
- Divide and place your components into functional "building blocks" where possible.
- Identify layout critical tracks on your circuit and route them first.
- Place and route each building block separately, off the board.
- Move completed building blocks into position on your main board.
- Route the remaining signal and power connections between blocks.
- Do a general "tidy up" of the board.
- Do a Design Rule Check.
- Get someone to check it

This is by no means a be-all and end-all check list, it's highly variable depending on many factors. But it is a good general guide to producing a professional first-class layout.

Lets look in more detail at the procedure described above.

We have already looked at the grids and track/pad sizes, these should be the first things that you set up before you start doing anything. No exceptions!

Many people like to jump straight into placing all the components into what they think is the most optimum position on the board, all in one hit. Whilst this can work for small circuits, you don't have much of a hope when you have more complex circuits with hundreds of components spread across many functional circuit blocks. Why?, because it's very easy to run out of "routing space", which is the room to lay down all your tracks. If you fix all your component positions and then try to route everything, you can easily paint yourself into a corner so to speak. Alternatively, if you space the components out too much, you can end up with a large board that does not make efficient use of space.

The hallmark of an inexperienced designer is a board that has every component evenly spaced out, and then has thousands of tracks and vias crisscrossing the board. It might work, but it can be ugly and inefficient, not to mention bigger and more expensive to manufacture.

The best way to start your layout is to get ALL of your components onto the screen first.

If you have a companion schematic package, then the simplest way to do this is to get your PCB program to import your schematic design and select all the components automatically. This will also be discussed later. If all you have is a PCB program, then you'll have to select each component from the library and place them down manually.

With all the components on screen, you should get a good indication of whether or not your parts will easily fit onto the size (and shape) of board that you require. If it looks like it's going to be a tight fit then you know that you will have to work hard to try and keep the component spacing "tight", and the tracking as efficient as possible. If it looks like you have plenty of room then you can be a bit more liberal in your layout. Of course, if it looks like you have buckleys chance of getting your components on the board, you'll have to go back to the drawing board.

Now analyse your schematic and determine which parts of the design can be broken up into "building blocks". Often this is fairly obvious. Say for example you have a complex looking active filter in your circuit. This would typically have a single input line and a single output line, but it will have lots of components and connections as part of the filter. This is a classic "building block" circuit, and one that lends itself well to combining all of these parts together in the same location. So you would grab all of these parts and start to rearrange them into their own little layout off to one side of your board. Don't worry too much about where the actual block goes on your board yet.

You will also need to partition off electrically sensitive parts of your design into bigger blocks. One major example is with mixed digital and analog circuits. Digital and analog just do not mix, and will need to be physically and electrically separated. Another example is with high frequency and high current circuits, they do not mix with low frequency and low current sensitive circuits. More about this later.

As a general rule, your components should be neatly lined up. Having ICs in the same direction, resistors in neat columns, polarised capacitors all around the same way, and connectors on the edge of the board. Don't do this at the expense of having an electrically poor layout, or an overly big board though. Electrical parameters should always take precedence over nicely lined up components.

Symmetry is *really nice* in PCB design, it's aesthetically pleasing and just "looks right". If you have something like two identical building block circuits side by side, and one is laid out slightly differently, it sticks out like a sore thumb.

If you have placed your components wisely, 90% of your work will be done. The last 10% should just be joining the dots so to speak. Well, not quite, but good placement is a good majority of your work done.

Once you are happy with the component placements, you can start to route all the different building blocks separately. When finished, it is then often a simple matter to move and arrange the building blocks into the rest of your design.

The Design Rule Check (DRC) will be covered later, but it is an essential step to ensuring that your board is correct before manufacture. A DRC basically checks for correct connectivity of your tracks, and for correct widths and clearances.

Getting someone to check your board may sound like an overly bureaucratic process, but it really is a vital step. No matter how experienced you are at PCB design, there will *always* be something you overlooked. A fresh pair of eyes and a different mindset will pick up problems you would never see.

If you don't have anyone to check your board over and don't have DRC capability, then you'll have to do it yourself. Get a printout of your schematic and a highlighter pen. Now, compare every single electrical "net"

connection on your board with the schematic, net by net. Highlight each net on the schematic as you complete it. When you are finished, there should be no electrical connections left that aren't highlighted. You can now be fairly confident that your board is electrically correct.

Basic Routing

Now it's time for some basic routing rules. Routing is also known as "tracking".


Routing is the process of laying down tracks to connect components on your board. An electrical connection between two or more pads is known as a "net".

- Keep nets as short as possible. The longer your total track length, the greater it's resistance, capacitance and inductance. All of which can be undesirable factors.
- Tracks should only have angles of 45 degrees. Avoid the use of right angles, and under no circumstances use an angle greater than 90 degrees. This is important to give a professional and neat appearance to your board. PCB packages will have a mode to enforce 45 degree movements, make use of it. There should never be a need to turn it off. Contrary to popular belief, sharp right angle corners on tracks don't produce measurable EMI or other problems. The reasons to avoid right angles are much simpler - it just doesn't look good, and it may have some manufacturing implications.
- Forget nice rounded track corners, they are harder and slower to place and have no real advantage. Stick to 45 degree increments. Rounded track bends belong to the pre-CAD taped artwork era.
- "Snake" your tracks around the board, don't just go "point to point". Point to point tracking may look more efficient to a beginner at first, but there are a few reasons you shouldn't use it. The first is that it's ugly, always an important factor in PCB design! The second is that it is not very space efficient when you want to run more tracks on other layers.
- Enable your Electrical grid, which is sometimes referred to as a "snap to center" or "snap to nearest" option. Let the software find the centers of pads and ends of tracks automatically for you. This is great for when you have pads and tracks which aren't lined up to your current snap grid. If you don't have these options enabled then you may have to keep reducing your snap grid until you find one that fits. Far more trouble than it's worth. There is almost never a reason to have these options disabled.
- Always take your track to the center of the pad, don't make your track and pad "just touch". There are few reasons for this. The first is that it's sloppy and unprofessional. The second is that your program may not think that the track is making electrical connection to the pad. Proper use of a snap grid and electrical grid will avoid problems here.
- Use a single track, not multiple tracks tacked together end to end. It may make no difference to the look of your final board, but it can be a pain for future editing. Often you'll have to extend a track a bit. In this case it's best to delete the old one and place a new one. It may take a few extra seconds, but it's worth it. People looking at your finished board may not know, but YOU'LL know! It's the little touches like this that set good PCB designers apart.
- Make sure your tracks go right through the exact center of pads and components, and not off to one side. Use of the correct snap grid will ensure that you get this right every time. If your track doesn't go through the exact center then you are using the wrong snap grid. Why do you need to do this? It makes your board neater and more symmetrical, and it gives you the most clearance.
- Only take one track between 100 thou pads unless absolutely necessary. Only on large and very dense designs should you consider two tracks between pads. Three tracks between pads is not unheard of, but we are talking seriously fine tolerances here.
- For high currents, use multiple vias when going between layers. This will reduce your track impedance and improve the reliability. This is a general rule whenever you need to decrease the impedance of your track or power plane.
- Don't "drag" tracks to angles other than 45 degrees
- "Neck down" between pads where possible. Eg, a 10 thou track through two 60 thou pads gives a generous 15 thou clearance between track and pad.
- If your power and ground tracks are deemed to be critical, then lay them down first. Also, make your power tracks as BIG as possible.


- Keep power and ground tracks running in close proximity to each other if possible, don't send them in opposite directions around the board. This lowers the loop inductance of your power system, and allows for effective bypassing.
- Keep things symmetrical. Symmetry in tracking and component placement is really nice from a professional aesthetics point of view.
- Don't leave any unconnected copper fills (also called "dead copper"), ground them or take them out.

If you are laying out a non-plated through double sided board, then there are some additional things to watch out for. Non-plated through holes require you to solder a link through the board on both the top and bottom layer.

- Do not place vias under components. Once the component is soldered in place you won't be able to access the joint to solder a feed through. The solder joint for the feed through can also interfere with the component.
- Try and use through hole component legs to connect top tracks to bottom tracks. This minimises the number of vias. Remember that each via adds two solder joints to your board. The more solder joints you have, the less reliable your board becomes. Not to mention that that it takes a lot longer to assemble.





An example of GOOD power routing (Left) and BAD power routing (Right)



An example of GOOD routing (Left) and BAD routing (Right)

Finishing Touches

Once you have finished all your routing, your board isn't done quite yet. There are a few last minute checks and finishing touches you should do.

- If you have thin tracks (<25 thou) then it's nice to add a "chamfer" to any "T" junctions, thus eliminating any 90 degree angles. This makes the track more physically robust, and prevents any potential manufacturing etching problems. But most importantly, it looks nice. 
- Check that you have any required mounting holes on the board. Keep mounting holes well clear of any components or tracks. Allow room for any washers and screws.
- Minimise the number of hole sizes. Extra hole sizes cost you money, as the manufacturer will charge you based on not only the number of holes in your boards, but the number of different hole sizes you have. It takes time for the very high speed drill to spin down, change drill bits, and then spin up again. Check with your manufacturer for these costs, but you can't go wrong by minimising the number of hole sizes.
- Double check for correct hole sizes on all your components. Nothing is more annoying than getting your perfectly laid out board back from the manufacturer, only to find that a component won't fit in the holes! This is a very common problem, don't get caught out.
- Ensure that all your vias are identical, with the same pad and hole sizes. Remember your pad to hole ratio. Errors here can cause "breakouts" in your via pad, where the hole, if shifted slightly can be outside of your pad. With plated through holes this is not always fatal, but without a complete annular ring around your hole, your via will be mechanically unreliable.
- Check that there is adequate physical distance between all your components. Watch out for components with exposed metal that can make electrical contact with other components, or exposed tracks and pads.
- Change your display to "draft" mode, which will display all your tracks and pads as outlines. This will allow you to see your board "warts and all", and will show up any tracks that are tacked on or not ending on pad centers.
- If you wish, add "teardrops" to all your pads and vias. A teardrop is a nice "smoothing out" of the junction between the track and the pad, not surprisingly, shaped like a teardrop. This gives a more robust and reliable track to pad interface, better than the almost right angle between a standard track and pad. Don't add teardrops manually though, it's a waste of time. But if your program supports automatic teardrop placement, feel free to use it. 

Single Sided Design

Single sided design can greatly reduce the cost of your board. If you can fit your design on a single sided board then it is preferable to do so. Look inside many of today's consumer items like TV's and DVD players, and you will almost certainly find some single sided boards. They are still used because they are so cheap to manufacture.

Single sided design however requires some unique techniques which are aren't required once you go to doubled sided and multi-layer design. It is certainly more challenging than a double sided layout. In fact, a single sided board design will be regarded inversely proportional to the number of jumper links used. No jumper links earns the admiration of many peers!

It is all about a balance between board size and the number of jumper links required. Almost every single sided board will require some jumper links, so it is important to minimise these.

Component placement is even more critical on a single sided board, so this is no time to make all your components nice and neatly aligned. Arrange your components so that they give the shortest and most efficient tracking possible. It is like playing a game Chess, if you don't think many moves ahead then you will get yourself in a corner pretty quickly. Having just one track running from one side of your board to the other can ruin your whole layout, as it makes routing any other perpendicular tracks impossible.

Many people will route their board as though it is a double sided board, but only with straight tracks on the top layer. Then when the board is to be manufactured, the top layer tracks are replaced with jumper links. This can

be a rather inefficient way to approach single sided design, and is not recommended. You must be frugal in your placement, and don't be afraid to rip everything up and try again if you see a better way to route something.

With experience, you will be able to tell before you even start, if a design is worth trying to route on a single sided board.

Double Sided Design

Double sided design gives an extra degree of freedom for designing your board. Things that were next to impossible on a single sided board become relatively easy when you add an additional layer.

Many (inexperienced) designers tend to become lazy when laying out double sided boards. They think that component placement doesn't matter all that much, and hundreds of vias can be used to get them out of trouble. They will often lay out components like ICs in neat rows, and then proceed to route everything using a right angle rules. This means that they will route all the tracks on the bottom layer in one direction, and then all the tracks on the top layer perpendicular to the bottom layer. The theory is that if you chop and change between layers enough times you can route almost anything using a "step" type pattern. This technique can be ugly and inefficient, and is a throw back to the old manual tape days. Many basic auto routers work in this way.

Stick to using good component placement techniques and efficient building block routing.

Double sided design can also give you the chance to make use of good ground plane techniques, required for high frequency designs. This will be discussed later.

Other Layers

There are a few other important layers on your PCB beside the copper track layers.

Silkscreen

The "silkscreen" layer is also known as the "component overlay" or "component layer". It is the layer on the top of your board (and bottom if needed) that contains your component outlines, designators (C1, R1 etc), and free text.

This is added to your board using a silkscreening process. White is a standard colour, but other colours are available upon request. You can even mix and match colours on the one board, but that usually costs extra.

When designing your board, make sure that you keep all your component designators the same text size, and oriented in the same direction.

When laying out your own component footprints, where possible, make sure that you add a component overlay that reflects the actual size of your component. This way you will be able to tell at a glance how close you can physically position your components. Ensure that all polarised components are marked, and that pin 1 is identified.

Your silkscreen layer will be the most inaccurately aligned of all your layers, so don't rely on it for any positional accuracy. Ensure that no part of the silkscreen overlaps a bare pad. There is no minimum width requirement for lines on the component overlay, so feel free to use smaller lines and text sizes to fit things in. If parts of the text or lines don't turn out perfectly on your board then it does not affect your design, unlike tracks and pads.

As a general rule, don't put component values on the silkscreen, just the component designator.

Solder Mask

A solder mask is a thin polymer coating on your board which surrounds your pads to help prevent solder from bridging between pins. This is essential for surface mount and fine pitch devices. The solder mask typically covers everything except pads and vias. Your PCB program will automatically remove solder mask from pads and vias. The gap it leaves between the pad and the solder mask is known as the "mask expansion". The mask expansion should usually be set to at least a few thou. Be careful not to make it too big, or there might be no solder mask between very fine pitch devices.

Your solder mask is displayed in your PCB package as a negative image, just like your power plane. Under normal circumstances you don't need to put anything on your solder mask layer. But if you want to leave the solder mask off a certain part of your board, you can place tracks and fills on your solder mask layer. It is often handy to remove a small square of solder mask from the top of your board, where there are no tracks underneath. This leaves a nice bare and visible part of your board to write something with a pen.

Solder masks come in two types, silkscreen, or "photo imageable". Photo imageable masks provide better resolution and alignment, and are preferred over silkscreened. You can get different colour solder masks, but the standard colour is green.

On most standard quality boards, the solder mask is laid directly over the bare copper tracks. This is known as Solder Mask Over Bare Copper, or SMOBC. You can get other coatings over your tracks in addition to the solder mask, but these are usually for fairly exotic applications.

You can have vias covered with solder mask if you wish, this is known as tenting. This is useful for close tolerance designs, to prevent solder from flowing into vias.

Mechanical Layer

The mechanical layer (which may go under other names depending on the package) is used to provide an outline for your board, and other manufacturing instructions. It is not part of your actual PCB design, but is very useful to tell the PCB manufacturer how you want your board assembled. There are no hard and fast rules for this layer, use it however you like, just make sure you tell your PCB manufacturer.

Keepout

The keepout layer generally defines areas on your board that you don't want auto or manually routed. This can include clearance areas around mounting hole pads or high voltage components for instance.

Layer Alignment


When the PCB manufacturer makes your board, there will be alignment tolerances on the artwork film for each layer. This includes track, plane, silkscreen, solder mask, and drilling. If you don't allow for this in your design, and make your tolerances too fine, you can end up in big trouble. Consult the manufacturer for what alignment tolerances they can achieve, and also what alignment tolerance you are paying for!

Netlists

A netlist is essentially a list of connections ("nets") which correspond to your schematic. It also contains the list of components, component designators, component footprints and other information related to your schematic. The netlist file can be generated by your schematic package. Generating a netlist is also called "schematic capture".

Your PCB package can then import this netlist file and do many things. It can automatically load all the required components onto your blank board. It can also assign a "net" name to each of your component pins. With nets assigned to your PCB components, it is now possible to Auto Route, do Design Rule Checking, and display component connectivity. This is the fundamental concept behind modern Schematic and PCB CAD packages.

Rats Nest



A Typical "Rats Nest" Display.

Your job of component placement will be made infinitely easier by having a "rats nest" display enabled. If there is one reason for going to the trouble of drawing up an accurate schematic and importing a netlist, this is surely it. For large designs, a rats nest display is essential.

A rats nest display is one where the program will draw a straight line (not a track) between the pads of components which are connected on the schematic. In effect, it shows the connectivity of your circuit before you start laying out tracks. At the start of your board layout, with all your components placed down randomly, this will appear as a huge and complicated random maze of lines. Hence the name rats nest.

The rats nest may look very daunting at first, but when you move each component the lines will automatically move with them. In this way you can see instantly which components are connected to which, without having to refer back to the schematic and constantly cross reference component designators. Once you have used this feature once, you won't want to live without it. Even when doing simple designs with a few dozen components, you will miss this functionality.

With the rats nest display enabled, it will be almost possible to lay out all of your components optimally in no time, without having to lay down one single track. The rats nest display will effectively show you what your tracks will connect to. The rats nest lines should disappear when you route your tracks between components, so your design will get less and less "complicated looking" as you go along. When all the rats nest lines disappear, your board is fully routed.

Design Rule Checking

Design Rule Checking (DRC) allows you to automatically check your PCB design for connectivity, clearance, and other manufacturing errors. With the large and complex PCBs being designed today, it is impractical to manually check a PCB design. This is where the DRC comes into its own, it is an absolutely essential step in professional PCB design.

Examples of what you can check with a DRC are:

- Circuit connectivity. It checks that every track on your board matches the connectivity of your schematic.
- Electrical clearance. You can check the clearance between tracks, pads, and components.

- Manufacturing tolerances like min/max hole sizes, track widths, via widths, annulus sizes, and short circuits.


A complete DRC is usually performed after you have finished your PCB. Some packages however have the ability to do “real time” (or “online”) DRC checking as you create your board. For instance, it won't let you connect a track to a pad it shouldn't go to, or violate a clearance between track and pad. If you have real-time DRC capability, use it, it's an invaluable tool.

Forward and Back Annotation

Forward Annotation is when you make changes to your existing PCB layout via the schematic editor. The program will take your schematic netlist and component designators, and import them into your PCB design, and making any relevant changes. Some packages will also automatically remove old PCB tracks that are no longer connected. You can do this at any time during your PCB layout. If you update your schematic, then you must forward annotate into your PCB design. You can do edits like this manually, but forward annotation automates the process.

Back Annotation is when you change one of the component designators (eg. “C1” to “C2”) on your PCB and then automatically update this information back into your Schematic. More advanced back annotation features allow you to swap gates on chips, and perform other electrical changes. There should never be much real need to use back annotation.

Multi layer Design



Typical 6 Layer PCB Construction

A multi layer PCB is much more expensive and difficult to manufacture than a single or double sided board, but it really does give you a lot of extra density to route power and signal tracks. By having your signals running on the inside of your board, you can pack your components more tightly on your board to give you a more compact design.

Deciding to go from double sided to multi layer can be a big decision, so make sure that a multi layer board is warranted on the grounds of board size and complexity. You can forget about making multi layer boards yourself, it requires a commercial manufacturer. Most of the hobby board suppliers will not do multi layer boards.

Multi layer boards come in even number of layers. With 4, 6, and 8 layer being the most common. You can go many layers above this, but now you are in the realm of the very specialised.

Technically you can get an odd number of layers manufactured, like a 3 layer board for instance. But it really won't save you any cost over a 4 layer board. In fact a 3 layer board might even be more expensive than a 4 layer board because it calls for a non-standard manufacturing process.

If you decide to go multi layer then make sure you make use all of your layers, there is no point leaving one completely blank.

With a multi layer board, you would typically dedicate one complete layer to a ground plane, and another to your power. With perhaps a few signal tracks thrown on the power layer if you need to. If you have a digital only board, then you'd often dedicate the entire power layer also. If you have room on the top or bottom layer, you can

route any additional power rail tracks on there. Power layers are almost always in the middle of the board, with the ground closer to the top layer.

Once you have your power taken care of on the inner layers, you'll be surprised at the room you now have available for your signal tracks. It really does open up a whole new dimension to routing.

If power planes are vital, and you have a lot of connections to route, then you may have to move from 4 to 6 layers. Six layers will give you four full signal routing layers and two layers dedicated to power. You can really do some advanced routing with 6 layers. Eight layers and above is basically more of the same.

With multi layer design comes the options of using different types of vias to improve your routing density. There are three types of vias - standard, blind, and buried.

Standard vias go through the whole board, and can connect any of the top, bottom or inner layers. These can be wasteful of space on layers which aren't connected.

"Blind" vias go from the outside surface to one of the inner layers only. The hole does not protrude through the other side of the board. The via is in effect "blind" from the other side of the board.

"Buried" vias only connect two or more inner layers, with no hole being visible on the outside of the board. So the hole is completely buried inside your board.

Blind and buried vias cost more to manufacture than standard vias. But they are very useful, and almost mandatory for very high density designs like those involving Ball Grid Array (BGA) components.

Power Planes

It is good practice to use "power planes" to distribute power across your board. Using power planes can drastically reduce the power wiring inductance and impedance to your components. This can be vital for high speed digital design for instance. It is good design practice to use power planes whenever possible. They can even be used on double sided boards, if most of your signal tracks are on the top layer.

A power plane is basically one solid copper layer of board dedicated to either your Ground or Power rails, or both. Power planes go in the middle layers of your board, usually on the layers closest to the outer surfaces. On a 4 layer board with complex power requirements it is common to dedicate one layer to your ground plane, and another layer to your various positive and negative power tracks. Your ground rail is usually your signal reference line, so a ground plane is first preference before a power plane is considered.

Many PCB packages have special Power Plane layers that are designed and laid out in reverse to your other normal tracking layers. On a normal tracking layer, your board is assumed to be blank, and you then lay down tracks which will become your actual copper tracks. On a power plane however, your board is assumed to be covered with copper. Laying down tracks on a power plane actually removes the copper. This concept can take some getting used to.

A simple power plane will not have any "tracks" (or removed copper bits) at all on it, but will just be one solid layer of copper. In which case you don't need to lay down any tracks to remove any copper. However, it is common practice on more complex boards to "split" the power plane by laying down tracks. This may be done to separate an analog and a digital ground, which will reduce the amount of digital ground noise which is coupled into the more sensitive analog circuitry. A typical split power plane would involve a "track" being placed from near your input power connector or main filter capacitors and the opposite edge of the board. Be careful not to accidentally cause a power "loop" on your board by inadvertently connecting the two halves of your plane on the other side of the board.

As a matter of course, you should place "tracks" completely around the outer edge of your board. This will ensure that the power planes do not extend right to the edge of the board. Power planes on the edges of your board can short to not only one another, but also to any guide rails or mounting hardware.

You don't have to use the actual Power Plane layer on your PCB package if you don't want to. You can use a regular signal layer and lay down copper fills and tracks yourself. Power Planes layers though often have some advantages that will vary from one PCB package another.

Good Grounding

Grounding is fundamental to the operation of many circuits. Good or bad grounding techniques can make or break your design. There are several grounding techniques which are always good practices to incorporate into any design.

- Use copper, and lots of it. The more copper you have in your ground path, the lower the impedance. This is highly desirable for many electrical reasons. Use polygon fills and planes where possible.
- Always dedicate one of your planes to ground on multi-layer boards. Make it the layer closest to the top layer.
- Run separate ground paths for critical parts of your circuit, back to the main filter capacitor(s). This is known as “star” grounding, because the ground tracks all run out from a central point, often looking like a star. In fact, try and do this as matter of course, even if your components aren’t critical. Separate ground lines keep current and noise from one component from affecting other components.
- If using a ground plane, utilise “split” plane techniques to give effective star grounding.
- “stitch” required points straight through to your ground plane, don’t use any more track length than you need.
- Use multiple vias to decrease your trace impedance to ground.

Good Bypassing

Active components and points in your circuit which draw significant switching current should always be “bypassed”. This is to “smooth” out your power rail going to a particular device. “Bypassing” is using a capacitor across your power rails as physically and electrically close to the desired component or point in your circuit as possible. A typical bypass capacitor value is 100nF, although other values such as 1uF, 10nF and 1nF are often used to bypass different frequencies. You can even have two or three different value capacitors in parallel.

When bypassing, you CANNOT replace multiple capacitors with one single capacitor, it defeats the entire purpose of bypassing! It is not uncommon for a large design to have hundreds of bypass capacitors.

As a general rule, you should use at least one bypass capacitor per IC or other switching component if possible. Common values of bypass capacitors are 100nF for general purpose use, 10nF or 1nF for higher frequencies, and 1uF or 10uF for low frequencies.

Special low Equivalent Series Resistance (ESR) capacitors are sometimes used on critical designs such as switch mode power supplies.

High Frequency Design Techniques

High frequency design is where you really need to consider the effects of parasitic inductance, capacitance and impedance of your PCB layout. If your signal is too fast, and your track is too long, then the track can take on the properties of a transmission line. If you don’t use proper transmission line techniques in these situations then you can start to get reflections and other signal integrity problems.

A “critical length” track is one in which the propagation time of the signal starts to get close to the length of the track. On standard FR4 copper boards, a signal will travel roughly 6 inches every nano second. A rule of thumb states that you need to get really concerned when your track length approaches half of this figure. But in reality it can actually be much less than this. Remember that digital square wave signals have a harmonic content, so a 100MHz square wave can actually have signal components extending into the GHz region.

In high speed design, the ground plane is fundamental to preserving the integrity of your signals, and also reducing EMI emissions. It allows you to create “controlled impedance” traces, which match your electrical source and load. It also allows you to keep signals coupled “tight” to their return path (ground).

There are many ways to create controlled impedance “transmission” lines on a PCB. But the two most basic and popular ways are called Microstrip and Stripline.

A Microstrip is simply a trace on the top layer, with a ground plane below. The calculation involved to find the characteristic impedance of a Microstrip is relatively complex. It is based on the width and thickness of the

trace, the height above the ground plane, and the relative permittivity of the PCB material. This is why it is important to keep the ground plane as close as possible to (usually) the top layer.

A Stripline is similar to the Microstrip, but it has an additional ground plane on top of the trace. So in this case, the trace would have to be on one of the inner layers. The advantage of stripline over microstrip is that most of the EMI radiation will be contained within the ground planes.

There are many free programs and spreadsheets available that will calculate all the variations of Microstrip and Stripline for you.

Some useful information and rules of thumb for high frequency design are:

- Keep your high frequency signal tracks as short as possible.
- Avoid running critical high frequency signal tracks over any cutout in your ground plane. This causes discontinuity in the signal return path, and can lead to EMI problems. Avoid cutouts in your ground plane wherever possible. A cutout is different to a split plane, which is fine, provided you keep your high frequency signal tracks over the relevant continuous plane.
- Have one decoupling capacitor per power pin.
- If possible, track the IC power pin to the bypass capacitor first, and then to the power plane. This will reduce switching noise on your power plane. For very high frequency designs, taking your power pin directly to the power plane provides lower inductance, which may be more beneficial than lower noise on your plane.
- Be aware that vias will cause discontinuities in the characteristic impedance of a transmission line.
- To minimise crosstalk between two traces above a ground plane, minimise the distance between the plane and trace, and maximise the distance between traces. The coefficient of coupling between two traces is given by $1/(1+(\text{Distance between traces} / \text{height from plane})^2)$
- Smaller diameter vias have lower parasitic inductance, and are thus preferred the higher in frequency you go.
- Do not connect your main power input connector directly to your power planes, take it via your main filter capacitor(s).

Double Sided Loading

Loading components on both sides of a PCB can have many benefits. Indeed, it is becoming an increasingly popular and necessary option when laying out a board. There are two main driving factors behind a decision to go with double sided loading. The first is that of board size. If you require a particular board size, and all your components won't fit on one side, then double sided load is an obvious way to go. The second reason is that it is required to meet certain electrical requirements. Often these days, with dense high speed surface mount devices packed onto a board, there is either no room for the many bypass capacitors required, or they cannot be placed close enough to the device to be effective. Ball Grid Array (BGA) devices are one such component that benefit from having the bypass capacitors on the bottom of the board.

Indeed, it is common to find double sided loaded boards with nothing but bypass capacitors mounted on the back. This allows the bypass capacitor to get as close to the physical device power pin as possible.

Be sure to involve your PCB assembler in discussions during the layout of your board. There are many things you can and can't do with double sided loading.

Auto Routing

"Real PCB designers don't auto route!", is an age old war cry. Whilst many will claim this is true, reality may often kick in, and there certainly are times when you do need to consider the use of an auto router.

Auto routing is the process of getting the PCB software to route the tracks for you. It will even attempt to route your entire board if you let it. Most of the medium to top range PCB packages will do this, and the technology and theory behind autorouting techniques can be mind boggling. Artificial intelligence and neural based technology are some of the marketing buzz words used.

If the PCB program can route the board for you, why not always use it? Doesn't it just automate a mundane process like laying down tracks?

The answers can be complicated and many, but no matter how “smart” an auto router is, it simply cannot replace a good human PCB designer. It is like trying to ask a computer program to paint a picture for you. If you give it enough information it may be able to produce something legible, but it won't be artistic, and certainly won't be a Mona Lisa.

Many people think that auto routers are a tool to help not so experienced PCB designers. In fact, the opposite is true! In the hands of an inexperienced designer, an auto router will produce a complete mess. But in the hands of a very experienced designer, an auto router can produce excellent results much quicker than the human designer could do.

Auto routers come in handy when you have complex boards with not much routing space, on non-critical parts of your layout. Non-critical parts of a board might include low frequency or static control signals to components like LED displays, switches and relays to name a few. Advanced autorouters do come with tools to let you specify exactly how you want electrically important tracks laid out. But by the time you have told it in excruciating detail what to do with every track, you could have laid it out yourself!

Never use an auto router to do your complete board, it will be a mess. But if you let it loose on a very specific non-critical area of your board, you can get some excellent results, sometimes indistinguishable from manual routing. You can even auto route a single connection, and this is sometimes handy when you are having trouble finding routing space in the final phase of your layout.

Unless you are very experienced at PCB design, simply stay away from auto routers. This cannot be stressed enough. Real designers don't auto route!

Auto Placement

Auto Placement tools are available in many higher end PCB packages. Professional PCB designers do not use Auto Placement tools, it's that simple. Don't rely on the Auto Place feature to select the most optimum layout for you. It will never work (unless it's an extremely simple board), regardless of what the program makers claim.

These tools do have one useful function however, they give you an easy way to get your components initially spread across your board.

Design For Manufacturing

Panelisation

If you are looking at getting your board automatically assembled with a pick-and-place machine, then it pays you to get as many boards onto the one “panel” as you can. A panel is simply a large PCB containing many identical copies of your board. It takes time to place a board into position on a pick and place machine, so the more boards you can load at once, the more cost effective your manufacturing will be.

A panel will also contain tooling strips on the top and bottom, to allow for automated handling of the panel. Different manufacturers may have different maximum panel sizes they can produce.

Each individual board can be “routed out” and joined with “breakout tabs”, or simply butted together and scoured out with a “V groove”. A V groove is a score mark placed on your board that allows you to easily “snap” the board along the groove. A breakout tab is a small strip of board perhaps 5-10mm long joining your board to your panel. Small non-plated holes are also drilled along this strip, which allows the board to be snapped or cut out of the panel after assembly.

You will need to consult your board loader to determine optimum panelisation size and requirements.

Tooling Strips

Tooling strips are strips of blank board down the top and bottom side of your board. They contain tooling holes, fiducial marks, and other manufacturing information if required.

Standard tooling holes are required for automated handling of your board. 2.4mm and 3.2mm are two standard hole sizes. Four tooling holes per panel is sufficient, one in each corner.

The tooling strips connect to your board(s) with breakout tabs or V Grooves.

Fiducial Marks

Fiducial marks are visual alignment aids placed on your PCB. They are used by automated pick and place machines to align your board and find reference points. A video camera on the machine can identify the center of fiducial marks and use these points as a reference.

On a panel there should be 3 fiducial marks, known as global fiducials. Bottom left/right and top left corners. They should be at least 5mm away from the board edges. They can be mounted on the tooling strips.

The fiducial mark should be a circular pad on the copper layer of diameter 1.5mm typically. The fiducial should not be covered with solder mask, and the mask should be removed for a clearance of at least 3mm around. The pad can be bare copper or coated like a regular pad.

Two local fiducial (one in opposite corners) is also required next to each large fine pitch surface mount device package on your board.

Thermal Relief

If you solidly connect a surface mount pad to a large copper area, the copper area will act as a very effective heat sink. This will conduct heat away from your pad while soldering. This can encourage dry joints and other soldering related problems. In these situations a thermal relief connection, which comprises several (usually 4) smaller tracks connecting the pad to the copper plane. Thermal relief options can be set automatically in many packages.

Soldering

Soldering considerations need to be taken into account when laying out your board.

There are three basic soldering techniques - hand, wave, and reflow.


Hand soldering is the traditional method typically used for prototypes and small production runs. Major impacts when laying out your board include suitable access for the iron, and thermal relief for pads. Non-plated through double sided boards should allow for ample room to get the soldering iron onto the top side pads.

Wave soldering is a common process used for surface mount and through hole production soldering. It involves passing the entire board over a molten bath of solder. Solder masks are absolutely essential here to prevent bridging. The major thing to watch out for when designing is ensuring that small components are not in the wave solder "shadow" of larger components. The board travels through the wave solder machine in one direction, so there will be a lack of solder trailing behind larger components. Surface mount devices are fixed to the board with an adhesive before wave soldering.


Reflow soldering is the latest technique, and is suitable for all surface mount components. The blank board is first coated with a mask of solder paste over the pads (solder "stencils" are used for this). Then each component is placed, and is sometimes held in place by an adhesive. The entire board is then loaded into an infrared or nitrogen oven and "baked". The solder paste melts (reflows) on the pads and component leads to make the joint. A newer reflow method called pin-in-paste or intrusive reflow is available for through hole devices.

Combinations of wave and reflow soldering can be used for mixed through hole and surface mount boards.


Wave soldering has the advantage of being cheap, but the disadvantage of imposing placement limits on your components. Reflow soldering is more complex and expensive, but it allows for very dense surface mount packing.



SMD Wave Soldering



Through-hole Wave Soldering



SMD Reflow Soldering

Basic PCB Manufacture

A PCB usually consists of a blank fibreglass substrate (“the board”), which is usually 1.6mm thick. Other common thicknesses are 0.8mm and 2.4mm. There are many types PCB substrate material, but by far the most common is a standard woven epoxy glass material known as FR4. This material has standard known properties, typical values of which are shown in the accompanying table.

The most often used parameter is probably the dielectric constant. This figure is important for calculating high speed transmission line parameters and other effects. An FR4 PCB is made up of glass and resin. Glass has a dielectric constant of approximately 6, and the resin has a dielectric constant of approximately 3. So an FR4 PCB can typically have a figure ranging from under 4, to almost 5. If you need an exact figure you will have to consult with your PCB manufacturer.

Typical FR4 Properties:

Dielectric Constant	3.9 to 4.8
Dielectric Breakdown	39kV/mm
Water Absorption	<1.3%
Dissipation Factor	0.022
Thermal Expansion	16-19ppm/degC

NOTE: These values can vary with manufacturers, check with your supplier for exact figures.

Other exotic base materials like Teflon are also available, but are only used for special designs that require a higher grade base material for a specific reason. There are cheaper materials than FR4, like phenolic base and CEM-1. These are hobbyist grade boards, but are also often used in some mass consumer products due to their low cost. They are not suitable for plated through holes or fine tolerance designs.

A blank base material coated with copper is known as a copper clad board.

A multi-layer board is made up of various individual boards separated by Preimpregnated Bonding Layers, also known as “prepreg”. There are different ways to stack these board layers up, and this will dictate what you can do with planes and blind/buried vias. Consult the manufacturer for their recommendations on this.

Surface Finishies

You can get your PCB manufactured with several different types of pad and track surface finish.

Very low cost single and double sided boards without solder mask typically have a basic tin coated finish. Beware of potential shorts between tracks with this method.

Your standard professionally manufactured board will typically have solder mask over bare copper (SMOBC) tracks, and a tin finish on the pads and vias which is Hot Air Leveled (HAL). Hot air leveling helps most surface mount components to sit flat on the board.

For large and critical surface mount components, a gold “flash” finish is used on the pads. This gives an extremely flat surface finish for dense fine pitch devices.

Peelable solder masks are available, and are handy for temporary masking of areas on your board during wave soldering or conformal coating.

Electrical Testing

You can have your finished PCB checked for electrical continuity and shorts at the time of manufacture. This is done with a automated “flying probe” or “bed of nails” test machine. It checks that the continuity of the tracks matches your PCB file. It may cost a fair bit extra, but this is pretty mandatory for multi-layer boards. If you have a manufacturing error on one of your inner layers, it can be very difficult to fix.

Signature

Like any work of art, no board is complete without adding your name or signature to it!

The signature can take any form your like. Some people put their name, initials, or a fancy symbol. Whatever it is, just make sure you add something.

A signature can be placed on any of the copper layers, or on the component overlay.

Submitting your design for manufacture

The first thing to know is which format to send your PCB file in. In Australia the standard format is any version of Protel (AutoTrax, PFW2.8, 99SE, DXP etc). Every manufacturer in Australia will happily take a Protel file. In fact, Protel format is their preferred way to receive a file. Many will also take other proprietary formats as well, but you'll have to check with them first. Supplying the original PCB package file will ensure that what you see on the screen is what you will get when your board is delivered. Unless you have a good reason to do so, don't supply your file in any other format.

Gerber files are the traditional and industry recognised file format, all major manufacturers will accept them. Many PCB designers still insist on generating and supplying Gerber files themselves, in order to have total control over the manufacturing process. In all but a few extreme cases, generating Gerber files is not necessary and a thing of the past. Generating gerbers adds an extra step of complexity to the PCB process where errors can creep in. So avoid the use of gerber files where possible, they can be troublesome unless you know exactly how to generate them correctly.

The manufacturer will ask for a lot of information before they quote, ask them what you need to provide with your file. Here is a basic checklist:

- A reference code and revision for your board. This makes it easy for both parties to track the progress of it.
- Desired manufacturing time, known as the "turn-around". 24 hours will cost a LOT more than 2 weeks!
- Quantity of boards required
- Board thickness (1.6mm, 0.8mm, 2.4mm etc). 1.6mm is standard
- Type of board (FR4, Teflon etc). FR4 is standard
- Number of layers
- Surface finish (SMOBC, HAL, Gold Flash etc). SMOBC and HAL is standard.
- What colour you want your solder mask and component overlay.
- Copper weight (1oz, 2oz etc). 1oz is standard.
- Whether or not you want electrical testing. This is pretty necessary for multi-layer boards.
- The Track/Space clearance of your board
- How your board dimensions are defined. Eg, on the mechanical layer.
- Whether you want boards "panelised" or supplied individually cut.

Many manufacturers will have "prototype" services where they will fit as many of your boards onto a standard "panel" as they can, all for one fixed price.

In most cases you will be charged a "tooling" cost. This is the cost of printing the photo masks for your board, and also setting up their machines. This is usually a one-off cost, so if you get the same board manufactured again, you won't have to pay the tooling charge.

That's the end of this tutorial.

Do you believe that is all there is to know about PCB design?

If you answered no, then you'd be right!

Good PCB design takes lots of experience, so go get started on your next board.

Happy routing!

David L. Jones ?