

A Closer Look at Printed Circuit Board Milling

Andrew Sterian
Grand Valley State University

Abstract

Several schools have recently incorporated printed circuit board (PCB) mills into EE/EET education as a means of producing PCB's in-house. Compared to breadboards or protoboards, the obvious benefit of the milling technology is the ability to build circuits with more robustness and permanence. Compared to commercial board manufacturers, PCB milling offers quicker turn-around and lower per-board cost. Previous papers have described how PCB mills have been integrated into the curriculum, but there has not been any detailed comparison of the tradeoffs between in-house milling and commercial manufacturing. In this paper we draw upon our experiences in using both circuit board milling and commercial PCB manufacturing for several student design projects. We compare costs, circuit board quality and ease of assembly, faculty time investment, support issues, and most importantly, how the two approaches to PCB construction fundamentally affect engineering instruction. We have identified several disadvantages of PCB milling, including high initial cost, more complex board layout and routing, and consequent limits on circuit complexity. We have also been surprised by some unexpected advantages of PCB milling, such as higher modularity in project designs and greater risk-taking. Our goal is to provide the benefit of experience to schools considering an investment in PCB milling.

1. Introduction

Students have two broad options for constructing electronic circuits: i) prototyping methods, including breadboards, wirewrap, and protoboards, and ii) circuit boards with soldered components. Circuit boards are most suitable for projects that are too complex for breadboards or require some degree of permanence, reliability, compactness, or required shape. The design of a circuit board (placement, layout, and routing) also affords the students exposure to modern CAD tools that are prevalent in the industry but are not applicable to the prototyping approaches.

In our junior level course on digital design, we incorporate a semester-long project culminating in a circuit board realization of a student team's circuit design. We initially used a commercial board manufacturer that was capable of producing inexpensive bare boards. Recently we purchased a circuit board mill in order to produce bare boards in-house. In this paper we compare these two methods of board fabrication. We provide a closer look at the details of circuit board milling and how it differs from commercial fabrication. We also discuss the impact that a circuit board mill has had on student learning, both in our courses and in other projects.

We are assuming that the circuit boards that need to be fabricated are at most two-layer boards, have low-to-moderate density, comprise mainly through-hole components, will be populated and soldered by hand, and are needed in only small quantities. We expect that these assumptions are consistent with the needs of most EE/EET departments for instructional use.

2. Overview of Commercial Board Fabrication

Having circuit boards fabricated by a commercial manufacturer is a fairly simple process that requires the electronic submission of Gerber (i.e., RS274X) files for artwork, a drill drawing, an Excellon drill file, and any special instructions. Modern CAD programs are capable of generating these files automatically from a board layout. Other files and file formats may be involved (for example, solder mask or silkscreen) depending upon the board manufacturer, and some manufacturers directly accept saved workspaces from popular CAD programs. Prices vary depending upon the required turnaround time, the number of board layers, the board area, the number and variety of drills, whether or not silkscreen and solder mask are included, and so on. Feature widths of 0.008" are commonly supported.

As an example, we had two 3"x4" two-layer boards fabricated for approximately \$65 with a two-day turnaround. These boards were through-hole plated but did not have silkscreen or solder mask. Several circuit board manufacturers^{1,2,3} serving the low-volume prototype market are listed in the bibliography.

3. Overview of Circuit Board Milling

A circuit board mill is essentially a three-axis plotter in that the "plotting head" is capable of fine-grain movement in the X-Y plane, and may also be lifted up or pressed down. Since this is a mill, the "plotting head" is actually a high-speed motor that is fitted with either a router bit or a drill bit. For creating copper traces, a router bit is used. The router is pressed down onto the surface of copper-clad board and is then moved in the X-Y plane. As it moves, the router removes copper thus creating an *isolation channel*. A copper trace is created (or, more precisely, left behind) when two isolation channels are milled on either side. See Figure 3.1 for an example of a trace created by milling and a photograph of a typical circuit board mill. The latter shows the three main components of a milling system: the mill, the controlling computer, and a vacuum machine.

The holes in the circuit board are created by replacing the router bit with a drill bit. The drill bit is moved (with the milling head up) to the desired X-Y position, then the head is pressed down briefly then released. A variety of drill bit sizes are used depending upon the needs of the circuit board.

The milling process begins with the same files described in Section 2. That is, Gerber files for artwork and an Excellon file for drill information. Specialized *isolation software* (included with the mill) converts these files into a set of instructions that indicate where the mill should travel in order to leave behind copper in the desired places (and also where the mill head is to drill holes). A separate *mill controller program* (also included with the mill) reads these instructions and sends commands to the mill over a serial connection to actually move the head, start the motor, etc.

In theory, then, there is no difference in the CAD process with respect to the interface between the designer and the manufacturer or mill; the CAD files are the same. In practice, the milling process does reach down to the design level (as we discuss below). A designer will need to be aware from the beginning whether the board will be manufactured commercially or milled.




Figure 3.1: The photograph on the left shows the bottom layer of a typical milled board. The photograph on the right shows our milling system. The mill (left) is controlled by a PC computer via a standard serial port connection. The vacuum system is partially visible under the table. Not shown are the supplies (boards, drill bits, etc.) that occupy additional space to the left of the mill.

4. Details of milling

In this section we present the details of the circuit board mill and the milling process. Our own mill is a Protomat/C30s manufactured by LPKF⁴. This is an entry-level mill with an 8"x13" bed. Higher-end products include automatic tool changers (see Section 4.5 below), higher-power spindle motors (which enable finer milling features using specialized tools), and larger milling beds. T-Tech⁵ is another leading manufacturer of circuit board mills and seems to have a product that is comparable in both capability and price. Recently, several lower-cost milling machines have become available from other manufacturers. The reader is encouraged to compare carefully the specifications of all alternatives and observe a working mill first-hand, if possible. Issues such as accuracy, repeatability, and reliability should be carefully considered.

4.1 Costs

The cost of our Protomat/C30s unit alone was \$9100. Additional costs included a starter package of supplies at a cost of \$400, a direct measuring microscope (see Section 4.4) for \$100, a vacuum system with automatic switch (see Section 4.10) for \$820, and the cost of shipping and insurance which cost \$123. Note that we did not purchase the option of training, which cost an additional \$1980 for two people and two days of training (as recommended). On-site training for two people over two days is \$2500. This training cost may compensate for the time investment required to successfully operate the mill (see Section 4.14).

The starting package of supplies that shipped with the unit was enough to mill a few boards with a limited set of drill sizes. We found it necessary to purchase approximately \$400 of additional drill bits in a variety of sizes for better matching to actual drill diameters on the circuit board (see Section 4.8). An exhaustive purchase of available drill sizes would cost approximately \$1500 but this is unlikely to be necessary for educational use.

Currently, each set of ten drills in a given size costs \$48, a set of ten router bits costs \$170, and a package of ten 9"x12" double-sided copper-clad boards with backings costs \$110. Both copper-clad boards, backings, and router bits are frequent consumables. A single router bit is rated for 45 metres of travel in the X-Y plane, and our typical boards have required between 5 and 10 metres of travel. Thus, a router bit must be replaced approximately every 5 to 10 boards. Drills, too, have a finite lifetime, and they may also break, but we have yet to replace a single drill even though we have already replaced 5 router bits.

Additional supplies include contour routers for cutting the board, creating cutouts and special shapes, etc. and end mills for removing large areas of copper (see Section 4.7). Contour routers cost \$90 for a package of 10 but last much longer than router bits (we are still on our first contour router bit). End mills are expensive, anywhere from \$142 to \$345 for a package of 10, depending on their diameter, and are quickly worn out.

The cost of our mill included one year of warranty service including telephone support and repair (but note that shipping costs for the unit, both to and from the manufacturer, are paid for by the customer). The yearly cost of extending this warranty is 10% of the purchase price of the mill, or \$910 in our case.

4.2 Through-hole plating

Perhaps the most significant difference between circuit board milling and commercial fabrication is the issue of through-hole plating. A two-layer board with plated through-holes has two major benefits: vias do not need to be manually soldered and device pins act as vias. The first benefit represents a savings in soldering time and effort, but the latter benefit can fundamentally affect the design of the board. When a device pin is not a via, it must be soldered on the solder side of the board (unless the device package lends itself to component-side soldering, e.g., resistors, capacitors, etc.). For single-layer boards, this is not an issue. For two-layer boards, this restriction can greatly increase the number of vias that are required and thereby decrease the routability of the board. The end result is that boards without through-hole plating are limited in complexity for a given size.

Figure 4.1 below shows an example of the same circuit routed with and without through-hole plating (using the built-in autorouter of the Eagle CAD program⁷). This board represents a partial circuit for a Motorola 68HC11 microcontroller (52-pin PLCC socket) in the expanded mode configuration with a 32Kx8 memory device and an 8-bit latch to perform address/data bus demultiplexing (for reference, see Figure 2-23 of the M68HC11 reference manual). Several port pins are also brought out to connectors. The board is approximately 2.2" square. The traces are 10 mils wide and all feature spacings are 8 mils (e.g., trace to trace spacing, pad to trace spacing, etc.)




Figure 4.1: Results of autorouting an example circuit assuming a plated through-hole process (left side) and no through-hole plating (right side). There are 11 vias in the board on the left side and each via is 40 mils in diameter. The board on the right side has 51 vias and each via is 50 mils in diameter since 40 mil vias are too difficult to solder manually. Both boards use 10 mil traces and 8 mil feature separations.

With through-hole plating, the board routes on two layers with 11 vias of 40 mil diameter (left side of Figure 4.1). Without through-hole plating, the board is only 76% routable on the bottom layer only. Using both layers and prohibiting the use of device pins as vias, the board is routed with 51 vias of 50 mil diameter (we have found that 40 mil vias are too difficult to solder manually). If these vias are manually soldered (by inserting a wire and soldering both top and bottom sides), it represents an additional 102 solder connections above the 122 required for the component pins. Thus there is much more work required to construct the board and there are higher chances of soldering errors.

Clearly, the lack of a through-hole plating process should be considered at the design stage. A dense board layout may not be achievable without a high number of vias. Also, the CAD program being used must be able to support autorouting without the assumption of through-hole plating. For example, the Eagle CAD program we use can not be directly instructed to route on two layers but not route directly to device pins on the top layer. Fortunately, this program allows user scripts to be written to directly manipulate the CAD data. We had to write such a script to place “restriction circles” above each device pin on the top layer to forbid the autorouter from making a connection there and forcing it to use a via to switch layers. While workable, this represents an extra step in the design process and generally increased complexity (e.g., greater autorouting time).

Our mill manufacturer (LPKF) offers three options for implementing a through-hole plating process. We have not purchased any of these process components thus only briefly discuss their capa-

bilities. One method presses copper bails into each hole and then solders them. This method is intended for a small number of holes and is fairly slow (two bails per minute). Another method uses conductive paste that is injected into each hole (using the same milling machine) and then removed to leave a paste residue. This residue then hardens into a thin conductive layer after baking at 160°C for 45 minutes. This method is intended to be faster and to plate more holes than the copper bail process. Finally, an electroplating process using several chemical baths is offered. This process can be used to plate any number of holes but requires up to 2 hours of processing time. The process also requires some fairly expensive chemicals, as well as proper disposal of these chemicals.

4.3 Isolation

Once the board layout files are generated (usually Gerber files for artwork, Excellon files for drills) they must be converted to instructions for the mill. A specialized program (provided by the mill manufacturer) is used for this purpose. It computes the paths for the mill head in order to create isolation channels with the router bit. At this stage, there are some options that can affect the construction of the final board.

The CAD program has presumably enforced a minimum feature spacing between, for example, pads and traces, traces and traces, and so on. The actual width of the isolation channels, however, is a function of the router bit and its depth of penetration into the board (see Section 4.4). The width of the isolation channel is fixed during milling. Achieving a wider isolation channel must be accomplished by multiple passes over the same path, slightly offset.

For example, we often use a nominal feature spacing of 10 mils for trace-to-trace spacings but 12 mils between pads and other features. We have found that this makes soldering easier and reduces the chances of solder bridges between pads and the surrounding copper. We must therefore use an isolation channel width of 10 mils. To achieve the 12 mil spacing around pads, two isolation passes must be made around the pad, with the second one overlapping the first by 8 mils. The isolation software supports this feature but it is important that the desired settings be co-ordinated with the actual router bit depth. Setting the bit too deep will lead to traces that are overly narrow and, in the extreme case, disappear altogether.

The width of the channel created by the router bit is a parameter that must be specified to the isolation program. This width can be changed, within limits, but it is important that the isolation program's concept of channel width and true channel width (see Section 4.4 below) be in agreement.

4.4 Trace-width accuracy

The actual width of an isolation channel is determined by the router bit's depth of penetration into the copper layer of the circuit board. The tip of the router bit is conical, thus the deeper the penetration, the wider the channel. The mill's head must be adjusted so that the router bit creates an isolation channel of the desired width. This width must agree with the width specified as a parameter to the isolation program, as discussed in Section 4.3.

The isolation channel width can be measured by routing a test pattern (a spiral) and then optically

inspecting the width of the resulting channel. We purchased a direct measuring microscope (from the mill manufacturer) that can be used to magnify the traces and also includes a graticule calibrated in increments of 0.1 mils. After the test pattern is milled, the isolation channels are inspected and the mill head is raised or lowered accordingly (raising it narrows the channels, lowering it widens the channels). This adjustment is performed in increments of approximately $4\mu\text{m}$ as provided by an adjustment wheel on the mill head. It is possible to achieve very accurate channel widths with this process. However, for maximum accuracy this process should be repeated whenever the bit is removed and reinserted (as occurs when switching from milling the bottom side of the board to milling the top side as hole drilling occurs between these two phases).

It is important that the actual channel width be set properly. If it is too large, it is possible that the traces will be extremely narrow or disappear altogether. If the channel width is too narrow, the traces and pads will be barely separated from the surrounding copper making soldering difficult and greatly increasing the chances of solder bridges.

4.5 Milling time

The time it takes to mill a board is affected by the number of traces, whether or not multiple passes are performed to achieve wider isolation channels, the size of the board, whether or not spike removal is required, the number of holes to be drilled, and the number of different drill sizes used. Unless an automatic tool changer is installed, each change from the router bit to a drill bit, or from one drill bit size to another, requires time as the head is moved to the tool change position, the operator must remove one bit, insert another, then the head returns back to the work area of the mill. This process is not terribly long (about a minute) but it can be tedious, especially if there is a large variety of drill sizes. Without an automatic tool changer, then, an operator must essentially be in attendance during the drilling phase of the process to remove then insert the required drill bits.

Additional one-time costs include taping the copper-clad board to the milling bed (to prevent the sides and corners from curling up), un-taping then re-taping the board after flipping it over (for two-sided boards), and setting the corners of the work area prior to milling (since multiple boards may be milled from one copper-clad board).

In practice, our simplest boards have been milled in 10 minutes, while our most complex boards required nearly 1 hour. If a large number of boards are to be milled, this may represent a significant investment in time on behalf of the operator.

It is possible, however, to sequence multiple, independent boards to be milled together. Then, each milling bit is applied to all of the projects before the bit must be changed, the work area need only be set once, the board need only be taped down once, etc. This requires less intervention by an operator than if the multiple boards were milled separately.

4.6 Spike removal

The nature of the milling process means that there will be small pieces of copper that are left on the board. If small enough, these pieces may eventually dislodge and cause short circuits. The iso-

lation software has an option (called “spike removal”) that performs extra milling in order to remove these spikes. A comparison of a board milled with and without spike removal is shown in Figure 4.2.




Figure 4.2: The photograph on the left shows a milled board with spike removal (see text). The photograph on the right is the same board without spike removal. One of the most prominent spikes, present in the rightmost picture but not in the leftmost one, is indicated with an arrow. For scale, the drilled holes are 31.5 mils in diameter. Traces are 8 mils wide, isolation channels are 10 mils wide, and the isolation around pads is 12 mils.

We have noticed, however, that allowing spike removal can significantly increase milling time due to the many additional extra milling movements required and associated lowering and raising of the milling head. This frequent raising and lowering of the milling head may also place stress on the milling machine.

As an example, the board shown in Figure 4.1 milled with spike removal would require 6.615 metres of milling travel compared to the 5.348 metres of travel for the same board without spike removal (almost 25% more). In addition, this does not take into account the large increase in milling head up/down movements.

4.7 Rubout

As seen in Figure 4.2, a conventionally milled board still contains a fair amount of copper that is not electrically connected. This copper can usually be left on the board with no ill effects. However, the designer may wish to remove this copper either to make soldering easier (decrease the probability of solder bridges) or for sensitive RF circuits. In these cases, another step in the milling process known as rubout must be performed. The isolation program may be instructed to effect a rubout phase, in which an end mill moves along the top of the board, removing copper in wide swaths.

We did not purchase any end mills as they are fairly expensive and were not necessary for our projects. The mill manufacturer makes available several photographs that illustrate the milling process both with and without rubout⁸.

4.8 Drill sizing

The issue of using the right drill for a given hole in the board involves both the CAD program and the isolation program. In our environment, both programs attempt to be helpful by defining ranges of values for which a desired drill size successfully maps to an available drill size. For example, a 32 mil desired drill size may be mapped to a 31.5 mil drill bit.

This appears to be a reasonable process but can lead to some unexpected problems. For example, if the CAD program is told that drill bits are available in 0.1mm increments^a and up to 10% oversize is allowed, then a device with pins that require a 1.11mm drill requires only an 8.1% oversize to reach the next widest drill, 1.2mm. The CAD program then instructs the isolation program that a 1.2mm drill is desired. However, the isolation program (as part of its configuration) is told that drills are truly available only in 0.2mm increments: 1.1mm, 1.3mm, 1.5mm, etc. It defines the “capture range” for each drill as 0.1mm away from nominal, thus a 1.3mm drill is used for hole sizes anywhere from 1.2mm through 1.4mm. The end result is that a 1.11mm hole is actually drilled with a 1.3mm drill. This may completely drill out the pad surrounding the hole, making it unsolderable.

While somewhat contrived, the above example does illustrate problems we have actually encountered in which the final drill size is quite unexpectedly too large or too small. We have found several rules to be very useful to prevent such problems: i) the list of drill sizes specified to the CAD program and the isolation program should exactly match the drills available on the mill, ii) the CAD program should be allowed to oversize drills by up to 10% and undersize them by 2%, iii) a Gerber viewer should be used (with the same drill information loaded) to view the job prior to milling.

Note that our isolation program came pre-configured with a wide variety of drill sizes, which did not actually ship with the milling machine. It is important, then, to address point (i) above and ensure that the isolation program is properly configured. Point (ii) above is a rule of thumb that seems to work well. Undersizing drills by 2% catches situations where the desired drill is 0.813mm and the closest drill size is 0.8mm, or just 1.5% smaller, for example. This does not impair the ability to insert the pins in these holes and allows a smaller variety of drill sizes to be stocked.

Regarding point (iii) above, the isolation program itself acts as a workable Gerber viewer. It is very easy to visually identify a drill that has been made too large as the surrounding pad appears very thin or completely obscured. Some time invested at this stage of the design process is very worthwhile.

Note, however, that the list of available drill sizes (or “drill rack”) may be different depending

a. Our mill is of European manufacture thus drill bit sizes are specified in millimetres.

upon whether a board is to be milled or is to be sent out for commercial manufacture. This is another point in the design process where the manufacturing process must be considered and the board layout may be affected. For example, it may be necessary to oversize some pads if a limited range of drill widths are available and the closest available match would be too large for the existing pad size. In addition, through-hole plating reduces the effective diameter of a hole by up to 3 mils or so, and this further complicates proper drill settings.

4.9 Registration

For two layer boards, it is crucial that the milling on the top and bottom layers be aligned to each other. This issue of *registration* is solved by using two removable retention pins that are embedded in two plastic strips which are in turn embedded in a longitudinal groove in the milling bed (see Figure 4.3). These retention pins penetrate the copper-clad board. A new copper-clad board must first have two 3mm holes drilled at exactly the position of these retention pins prior to milling.



Figure 4.3: These photographs illustrate the retention pin mechanism used to ensure top-bottom registration of the circuit board during milling. The bed of the mill has a longitudinal groove into which two (red) plastic strips are inserted. A hole is drilled in each of these strips and a retention pin is inserted. The photograph on the left illustrates this arrangement. A circuit board must have holes drilled to accommodate the retention pins. The board is placed over these pins when milling, thus ensuring registration (see photograph on the right).

When milling, the board must always be placed over these retention pins. After the bottom side is milled, the board is turned over and once again placed over these retention pins prior to milling the top side. A crucial alignment then is to ensure that the imaginary line formed by the two retention pins is exactly the longitudinal axis of the mill, otherwise misregistration occurs.

There are two ways in which this alignment may be violated. If the longitudinal axis of the mill is

parallel to the retention pin line but offset (i.e., not coincident), the milling software can be adjusted to compensate for this offset. A procedure is detailed in the documentation whereby a hole is drilled from one side, the board is flipped over, and a circle is milled around the hole from the other side. The misregistration between the hole and the circle is used to compute the necessary offset. This is a fairly simple procedure and need only be performed infrequently when the retention pins or their plastic strips are moved, or the machine drifts out of alignment.

A more serious form of misalignment occurs when the imaginary retention pin line is no longer parallel to the longitudinal axis of the mill. Adjusting for this form of misalignment is a sensitive process that needs to be performed by the manufacturer.

4.10 Vacuum

The milling machine must have a vacuum attached in order to remove the board material as it is milled out. The vacuum hose attaches at the back of the milling machine and built-in hoses reach to the milling head, where the vacuum is required. A useful attachment for this vacuum is a switch that automatically applies or removes power to the vacuum depending upon whether or not the milling motor is on. This ensures that no milling occurs without the vacuum on (which may damage the board or even the mill), nor does the vacuum need to be remain on at all times.

In our experience, the vacuum is quite noisy, more so than the mill itself, thus it is nice to not have it on all the time. The extra safety of not being able to mill if we forget to turn the vacuum on makes this automatic switch a good investment.

4.11 Multi-layer boards

The mill is capable of milling two-layer boards without any special equipment. Boards with more layers must be milled as multiple two-layer boards. Additional laminating equipment, which can be purchased from the mill manufacturer, can then be used to form a single, multi-layer board.

4.12 Computer requirements

The mill is under the control of a milling program that runs on a PC. We use an old 90 MHz PC running Windows 98 as an operating system without difficulty. Note, however, that the demands of the milling control program are high as the communication with the mill (over a standard serial port) is very timing-sensitive. For example, the program requires that the serial FIFO buffers on the UART be disabled! For this reason, we assume that during milling, no other programs or tasks should be running on the computer. This computer, then, should probably not be a central file server, print server, firewall, or other multi-use system.

4.13 Power supply

The electrical connectivity of the mill and associated equipment must be considered. At the very least, the mill requires a power connection, as does the vacuum and the controlling computer. Together, these three devices can draw a large amount of current (nearing 15A in our estimation). We had the experience of connecting all three devices to a single outlet through a power strip and

observing that the milling motor would noticeably decrease in speed when the vacuum was turned on. Spreading the power connections over separate 15A circuits eliminated the problem.

The conclusion is that a fair amount of power needs to be available to the milling setup and trying to operate the system from a single outlet is ill-advised. In purchasing a mill, this requirement (which will also most likely dictate the physical placement of the mill) must be considered.

4.14 Documentation

The documentation that arrived with our mill was terse at best. There was a large amount of information provided but more in the form of reference material rather than tutorial. It is our impression that the manufacturer expects all mill operators to take their training course and not try to figure things out on their own. It is possible to do so (as we did) but one can expect a fair amount of trial and error before the milling process is streamlined. The documentation that accompanied both the isolation program and the milling control program was equally terse.

Fortunately, it is difficult to truly cause great damage to the mill unless some fairly obvious warnings are not heeded. For example, milling without vacuum or running the milling head motor without a bit inserted may damage the mill, but there are clear warnings regarding these errors. There is little harm, then, in experimentation (save the loss of some milling bits and perhaps a board or two) and this may be, perhaps, the best way to become a successful board maker.

5. Comparison of fabrication techniques

For the EE/EET department considering the purchase of a circuit board mill, we believe it is useful to compare the processes of milling against commercial fabrication.

5.1 Per-board costs

A major benefit of in-house milling is that milled boards are very inexpensive, if fixed costs are not considered (i.e., cost of the mill itself, cost of maintenance, etc.). For example, assuming a 3"x3" board that requires 7.5 metres of mill head travel, we estimate the cost of this board (in actual board usage and router bit wear) to be \$14.50. This does not include wear on drill bits, cost of vacuum bags, etc. as these have negligible per-board costs. The cost for milling a board, then, is much lower than the approximately \$65 that it costs to have a board commercially fabricated (of prototype quality). The end result is that there is much less pressure on students to not make mistakes. If a board design has been so badly implemented, the cost for milling a new board is only a few dollars and some time for the mill operator (perhaps even the student).

This decreased emphasis on achieving a 100% correct board design has two benefits. Firstly, it allows students to spend less time checking the board design and spend more time on the schematic and other aspects of the design. Secondly, circuit boards that do have errors provide a concrete learning opportunity for students: they appreciate first-hand the necessity for careful, correct work as the cost of post-fabrication fixes is much larger than the cost of making sure it is right the first time. In our experience, no number of admonitions to work carefully, check all pinouts, spot-check routed connections, etc. is as good a motivation to do good work as is a board that needs

several cuts and jumpers in order to fix sloppy work. If the board is truly unsalvageable, the cost of milling another is small.

Another benefit of low per-board cost is that students appear more likely to take risks. When we used a commercial fabrication process, we clearly told the students that they had only one opportunity to get it right since we simply couldn't afford to have multiple boards fabricated. Students, then, were very cautious, and understandably so. With in-house milling, some students chose riskier (and more interesting!) projects because the cost of failure was simply some redesign and another milling pass.

5.2 Breadboarding support

A benefit of having an in-house mill is the ability to create inexpensive circuit boards for use in a variety of courses. For example, in our course on Dynamic Systems (a course only taken by Mechanical Engineering students), the LM675 high-current op-amp is used in a motor control laboratory, but the pins of this device are not on 0.1" centers. This makes the device not suitable for breadboarding in the laboratory. It was a simple matter to design and mill adapter boards that provided an interface between the pins of the LM675 and a single-row connector with 0.1" centers that could be inserted into a breadboard.

5.3 Modular circuit design

Another benefit of low per-board cost is that a student project can be partitioned over multiple boards without additional expense. One recent student project used three small circuit boards to implement an RPM meter for a motor shaft. This project had one circuit board with a reflective photosensor mounted to it, along with the detector logic. Another circuit board had an LCD display and connector mounted to it (for remote display). The third circuit board comprised the main logic for the system. Multi-board systems such as these would be discouraged when each board incurs fabrication cost with a commercial manufacturer.

Another student project involved the use of a high-quality analog-to-digital converter interacting with a microcontroller-based design. The latter part of the design was complete and well understood but the students were not confident in the analog-to-digital converter and associated circuitry. The students decided to build a circuit board with a connector instead of the A/D converter circuitry and later build a separate board with the A/D converter that could be plugged in to the connector. This allowed them to experiment with the A/D design and replace it with a different design if necessary. This modular approach to their design allowed them to proceed with both hardware and software development even when they were not yet completely confident in one part of their design.

Dividing a circuit into separate, connecting parts (as in the above example) is a method of managing riskier projects (which we mentioned in Section 5.1 are more likely to be pursued when per-board costs are low). The cost of using commercial fabrication for this approach, however, would have been nearly double the cost of a single board. With in-house milling, the additional cost of the modularized design was negligible.

5.4 Turnaround time

From the time that CAD files are sent to a commercial manufacturer to the time boards are received can be as short as 2 working days. The time to mill in-house, however, is approximately 1 hour. This reduced turnaround time opens up some interesting possibilities in instruction. We are considering a single 3-hour laboratory in which students design a (simple!) circuit, perform the board layout and routing, then all designs are milled as a group. After soldering, students have a physical realization of their design, along with an appreciation for the design-layout-mill process. This type of laboratory would still be possible with commercial manufacture but it would require a break of a few days in the laboratory until the boards were received, plus the cost may be prohibitive.

5.5 Cost analysis

Using the approximations that a milled board costs \$14.50 and a commercially manufactured board costs \$65, it would require the milling of over 200 boards for the savings to recover the initial cost of the mill. Whether or not a circuit board mill is a wise financial investment clearly depends upon the board making needs and habits of a particular department, as well as the nature of the funding for the mill.

The largest recurring costs of a circuit board mill are likely to be a yearly warranty contract (\$910), copper-clad boards and backings (\$110 for a package of 10), and router bits (\$170 for a package of 10).

5.6 Board density

The through-hole plating of commercially fabricated boards has proven to be a big advantage for high-density boards, as discussed in Section 4.2. Before investing in a circuit board mill, it is important to consider the boardmaking needs of the department (i.e., will high-density boards be commonplace?) and determine whether an additional investment should be made in a plating process, or whether commercial manufacture is the best alternative.

5.7 Ease of soldering

We have found commercially fabricated boards to be slightly easier to solder. This is owing to two factors: these boards are tin plated and they do not have the large areas of copper remaining on milled boards. Tin plating can be applied to milled boards prior to soldering for fairly low cost, but this represents an extra process step and additional time expense. The large copper areas on a milled board present opportunities for solder bridges. These copper areas can be removed with the rubout process (see Section 4.7) but this also represents extra processing, time, and in this case, cost due to the wear on the end mills.

5.8 Availability

A benefit of commercial board manufacturers is that they are always available, and even if one is not, others are. A circuit board mill may need repair, which may require shipment of the unit to

and from the manufacturer. This can lead to a delay of days or even weeks (in addition to the cost of shipping, which is not insignificant for a heavy yet delicate unit). If Murphy's Law applies, the mill will fail at the worst possible time and student projects waiting for completion will be stalled.

If a circuit board mill is a central resource for a course, the instructor must consider carefully the consequences of mill failure at a critical time in the course.

5.9 Silkscreens and solder mask

The processes of applying silkscreen and solder mask are separate from milling. Some commercial manufacturers offer these options. If these additional process steps are required, commercial manufacture is the preferred solution (unless in-house facilities for applying silkscreen and solder mask exist). Note that a form of solder mask is possible using a process that mills out holes from a special foil⁹.

5.10 Project support

A circuit board mill can support departments outside of EE/EET and find other uses supporting student projects where commercial milling costs may be too high. Our mill has been used to construct circuits for extracurricular student activities, such as the Solar Boat competition and the Firefighting Robot competition. In this role, the circuit board mill enhances the quality of student education outside of the classroom and increases the level of service and recognition of the EE/EET department.

6. Conclusions

We have presented several details of circuit board milling that will hopefully be of use to the department considering an investment in circuit board milling. This investment, we have argued, should be in an improved learning experience for the students, not in cost, as hundreds of boards need to be milled for solely the purchase price of the mill to be recovered. The major benefits of the circuit board mill are low per-board cost (encouraging modular project designs, higher-risk projects, and supporting extracurricular activities), and turnaround time on the order of hours instead of days. The drawbacks of circuit board milling are lower allowable board density (unless through-hole plating equipment is also purchased), high initial cost, lack of silkscreen and solder mask (and slightly more difficult soldering), and service interruptions during repair.

Finally, the degree of time investment by faculty and laboratory support personnel must be considered. Circuit board milling requires knowledge and experience with the operation and interaction of three separate software programs (CAD program, isolation program, milling control program) and a new hardware device (the mill itself, which may be a very unfamiliar form of equipment for EE/EET faculty, such as the author of this paper). A circuit board mill is by no means "plug and play". The department purchasing a mill must consider carefully where the expertise for operating the mill will reside and who will be responsible for operating the mill. Faculty time commitments must be balanced against continuity of knowledge, the time it takes to train students on the software and hardware, the potential for damaging the equipment, etc.

In our case, the investment in the circuit board mill has been a positive experience. Despite the drawbacks listed above, we firmly believe that our students have derived benefit from in-house milling capabilities and that the investment (in time and money) was well worthwhile.

7. Acknowledgments

I would like to thank David Carney for providing several interesting insights regarding circuit board milling at the University of Waterloo.

Bibliography

1. URL: <<http://www.apcircuits.com>>, Alberta Printed Circuits, a circuit board manufacturer
2. URL: <<http://www.4pcb.com>> Advanced Circuits, a circuit board manufacturer
3. URL: <<http://run.to/pcb>> OLIMEX, a circuit board manufacturer
4. URL: <<http://www.lpkf.de>>, LPKF, a circuit board mill manufacturer
5. URL: <<http://www.t-tech.com>>, T-Tech, a circuit board mill manufacturer
6. Braun, C.G. "The CSM Electronics Prototyping Facility," *Proceedings of the 1997 ASEE Annual Conference*.
7. URL: <<http://www.cadsoft.com>>, CadSoft, authors of the Eagle schematic/board CAD program
8. URL: <<http://lpkf.herrlich-ramuschkat.de/artikel/view/examples.html>>, LPKF photographs showing example milling processes
9. URL: <<http://lpkf.herrlich-ramuschkat.de/artikel/view/viewArtikel.cfm?artGruppeID=12&ID=13>>, LPKF Easy-Solder method describing the creation of solder masks using a milling process

ANDREW STERIAN

Andrew Sterian is currently an Assistant Professor in the Padnos School of Engineering at Grand Valley State University. He received his B.A.Sc. in Electrical Engineering from the University of Waterloo, Canada and the M.S.E. and Ph.D. in Electrical Engineering from the University of Michigan, Ann Arbor. He has taught courses in signal processing, digital systems, and microcontrollers.