

-- SheetCNC.co.uk -- Clever Cuts

A Guide to Using a Home-Built CNC Mill/Router

Please remember that all the information in this document is copyright. This means that you should not make a copy of this document, nor send any electronic or paper copy of this document to any third party, without written permission of SheetCNC.co.uk.

Owners of a SheetCNC.co.uk product are permitted to make printed and electronic copies of this document for their own backup and working purposes. Copies may not be sold or transferred.

Every copy must include this notice.

Thank you.

Contents

Warning.....	3
Introduction.....	4
Temporary Spoilboard.....	4
CAD.....	5
Concepts.....	5
Design Choices.....	5
Procedure.....	6
Further Learning.....	8
CAM.....	10
Concepts.....	10
Cutting Choices.....	10
Setting Up SheetCAM for SheetCNC.....	10
Creating G-Code from the CAD Drawing.....	14
Explanation.....	22
Other SheetCAM Actions.....	25
Cutting the Part.....	26
Procedure.....	26
Chips.....	29
Chip Load.....	29
Trade-Offs.....	30
Materials and Settings.....	32
Materials for First Projects.....	32
First Settings for SheetCNC.....	32
Probing.....	35
Probing The Z-Axis.....	35
Probeless XY Touching-Off.....	38
Machining Tricks.....	39
T-Nut Lap Joints.....	39
Screwed T-Joint.....	39
Quick and Dirty Fastener.....	41
Cable-Tie T-Joint.....	41
Indexing and Orienting Parts.....	42
Adding Bearings.....	42
Precisely Matching Mating Parts.....	43
Gears.....	43
Arcs.....	44
Just One Tab.....	45
Perfect Perforations.....	47
Overcutting.....	47
Printed Circuit Boards.....	48
Aluminium and Brass.....	52
Production Jobs and Repeat Jobs.....	53
Touching Off Repeatably.....	53
Positioning the Workpiece Repeatably.....	53
Clamping the Workpiece Repeatedly in the Same Position.....	53
Workflow.....	54
Troubleshooting.....	56

Warning

This document is provided free in the hope that it may be useful but with no warranty whatsoever.

While this document uses SheetCNC as a worked example, it does not form part of the assembly or usage instructions of SheetCNC. No warranty whatsoever is given about the suitability of the procedures or the third-party equipment and software described in this document.

The focus of this document is on enabling the user to acquire new skills specific to CAD, CAM, and CNC. For sake of brevity and clarity, and so as to focus entirely on the particular techniques described, aspects relating to health and safety are omitted.

Take note that CNC machines can be dangerous if misused. SheetCNC.co.uk cannot know the full circumstances of your homebuilt machine. You are advised, before operating the machine, to heed any safety instructions and warnings supplied with your equipment and to work in accordance with the procedures established by your risk assessment.

Suggested cutting parameters, such as speed, RPM, and depth of cut are for rough guidance only. The reader should start with conservative settings and work up gradually to full-speed cuts, particularly when testing a newly built machine.

Introduction

This document gives guidance in using a 3-axis CNC mill/router such as SheetCNC. It covers CAD drawing, through the CAM process, and on to cutting a simple part on the machine.

We also offer some suggested materials to experiment with, and design tips for creating strong and useful assemblies.

This is not a comprehensive CNC instruction manual nor a tutorial, but we hope it covers sufficient to get you started. SheetCNC owners are welcome to contact us for further guidance. We will always help as much as we reasonably can, by email, at no extra charge.

Temporary Spoilboard


If you are intending to make trial parts on the machine before surfacing its bed, you must still install a temporary spoilboard on the bed. Any sheet of MDF of at least 9mm thickness, securely fastened, should suffice. The temporary spoilboard need only cover a portion of the bed - allowing a large margin around any workpiece you intend to cut. Ensure that you know the location of any temporary fasteners and avoid placing any workpiece on or near those fasteners.

Do not use plywood for the spoilboard.

CAD

First step in producing work on SheetCNC is to create a drawing of the parts. For this a CAD program is required. In this example we use TurboCAD v21. Adapt the example as necessary for your chosen software.

As you may be using a trial version of SheetCAM for your CAM software, we'll limit the example to something utterly trivial, so as not to exceed the limits of the free version of that software.



Concepts

In TurboCAD basic 2D objects - lines, circles, rectangles - are drawn using any combination of snaps, freehand placement, and numeric specification. These objects are then modified and combined as necessary using various tools.

A drawing is made up of layers. Each layer is conceptually like a separate sheet of acetate with objects drawn on it in coloured marker-pen. The complete drawing is comparable to the collection of all those acetates stacked together on an overhead projector for viewing. Each individual object in the drawing has to be on exactly one layer in the drawing. A layer can, of course, have many objects on it. Each drawing can have as many layers as the user wishes.

To help the user to visualise the layers on-screen, the CAD program assigns each layer a default colour. Every object on the layer will - unless specifically set otherwise - be shown in the colour of that layer.

Design Choices

The example design shown above - although intrinsically dull - does raise a couple of points:

- The design includes the outline of the material from which it will be cut (the black outline).

- There are fixing points (in two of the corners) that will be used to fasten the workpiece to the bed, in this case using screws. The material must be held tightly against the tendency of the cutter to both lift the material off the table and push it around on the table. The cutter must, of course, avoid the fasteners during cutting of the part. So we specify the location of the fasteners during the design phase, choosing their locations to keep them clear of the cutter and also clear of the part being produced. In the machining stage we use the machine itself to mark out the locations of the fasteners. Clearly we need to know the size of the cutter before we can position the fasteners - a wider cutter needs more clearance between the part and the fasteners. So, even at this stage, machining choices have to be part of the design process.
- The design of this part makes no concessions to the method of manufacture. In particular it is drawn with perfectly square-cut corners. Sharp corners force a CNC machine to decelerate for the turn and then accelerate back to full cutting speed afterwards. If the part does not need square corners, it would be better to round them off so that the cutter moves through a large smooth arc allowing the speed to be maintained. Such design choices ultimately have consequences for cutter wear, machine performance, and finish quality. See *Chips*.

Procedure

This procedure should work with any new, default installation of TurboCAD v21. It will differ only slightly for other versions.

Open TurboCAD.

Create a new file and immediately save it as a DXF-R14 format, rather than the default TCW format. This will cause the file to be saved always as DXF format in future.

Menu->File->SaveAs...

Click SaveAsType to select DXF.

Click Setup... and choose AutoCAD R14 Drawing, precision 6, other options unchecked.

Save the file. A filename without spaces is preferred.

Set up CAD to use metric dimensions.

Menu->Options->SpaceUnits...

...etc

Turn on Inspector Bar and display of User Coordinate System

Menu->Workspace->InspectorBar (ensure is on)

Menu->Workspace->UserCoordinateSystem (ensure is on)

Menu->Workspace->WorldCoordinateSystem (ensure is off)

Set up snaps to operate on vertices, middle points (of lines), centre points (of circles), and intersections only. Ensure magnetic point is set to be shown.

Menu->Modes->Snaps->...

...etc

Create some useful layers.

Menu->Options->Layers...

Click "New Layer" symbol and name the new layer OutsideThru.

Click on the colour-box for the new layer and select red.

Repeat to create layers InsideThru (green), Pocket3 (blue), Drill (black), Material (black).

Close the Layers window.

Use the rectangle tool to draw a square with sides 100mm.

Menu->Insert->Line->Rectangle.

Click somewhere near to the origin-axes which are displayed in the (empty) drawing at (0,0) to position the first corner of the rectangle there.

Press <Tab> and type '100' for the A dimension. Do not press <Enter>.

Press <Tab> and type '100' for the B dimension.

Press <Enter>.

Tip: Zoom by rolling the mouse-wheel. Pan by zooming out a little, moving the cursor to where you would like to centre the view, then zoom in again. Reset the view by pressing <Alt><UpArrow>.

Move the rectangle to the final location.

Press <F7> to select the last-created object (the rectangle).

Tap key 'D' then, without clicking, move the rectangle's handle until it snaps to the bottom-left corner of the rectangle, then click the left mouse button to lock it there.

Press <Tab> repeatedly until focus is on the 'Pos X' box, and type '20'. Do not press <Enter>.

Press <Tab> again and type '20' into the 'Pos Y' box.

Press <Enter> to commit to moving the rectangle to the specified new coordinate.

Press <Esc> to unselect.

Place the rectangle onto the OutsideThru layer.

Press <SpaceBar> to choose the select tool.

Click somewhere on the rectangle's perimeter to select it.

Press <Ctrl>-F to open the properties dialogue for the rectangle.

On the General tab set the Layer to OutsideThru and set Print Style to 'By Layer'.

Click OK.

Use the rectangle tool to draw another square with 20mm sides, somewhere near the large rectangle.

Menu->Insert->Line->Rectangle.

Etc

Select the newly drawn square, and - without moving its handle - move that rectangle to PosX=70 and PosY=70, similarly to above. Put it on layer Pocket3.

Press <F7>...

Etc

Use the circle tool to draw a circle intersecting with the edge of the small square.

Insert->Circle/Ellipse->CenterAndPoint.

Move the cursor close to the top edge of the small square, until the magnetic point snaps to the centre of the edge. Click to set the centre point of the circle there.

Press <Tab> and type the radius '5'.

Press <Enter>

Move the circle to layer InsideThru.

Press <F7>.

Press <Ctrl>-F

Etc

Use the rectangle tool to create a square 140mm on each side. Leaving the square's handle at its centre, move the square to PosX=70 and PosY=70, as above. Put it on layer Material.

Press <F7>.

Etc

Unselect everything.

Press <Esc>.

Insert two points to specify fasteners for the workpiece.

Insert->Point->Star

Click twice, to roughly position the two points somewhere near the bottom-left and top-right corners of the largest rectangle.

Press <F7> to select the last point.

Use <Tab> Etc to move the point to PosX=5 PosY=5.

Press <Ctrl>-F to open the properties of the point.

Change the point style to 'AutoCAD' and the size to 6mm then close the properties dialogue.

Press <F7> again to select the previous point.

Repeat as above, but move this one to PosX=135 PosY=135

Unselect everything.

Press <Esc>.

Save the file, again as DXF-R14 format.

<Ctrl>-S

That's everything. The drawing is created and saved.

Further Learning

To get the most out of your CAD package so as to create complex shapes on SheetCNC, there are several other essential 2D tools to master. These include...

- 2D Add and 2D Subtract. *Modify->2DBooleanOperations->...* (but see the warning in *Tricks*.)
- Explode object into parts. *Format->Explode*.
- Join polyline. *Modify->JoinPolyline*.
- Chamfer and line tools. *Modify->...*
- Arcs, polygons, etc. *Insert->...*
- Text. *Insert->Text*. ('Flexible' text can be used to define pockets for engraving.)

Blocks, Groups and Selections

In TurboCAD, as in most CAD programs, there are three key methods for subdividing a drawing into parts so that a complex drawing can be managed more easily.

Blocks: *Format->CreateBlock*. A block joins multiple objects into a single entity and creates a record of that entity so that it can be repeatedly inserted into the drawing. Blocks act as a single entity for purposes of arrangement in the drawing. The collective properties of the objects in the block (colour, layer, etc) cannot be modified en masse in the drawing. The properties of an individual member of the block can be modified only by editing the block itself (*View->Blocks*). Changes to the recorded block automatically change all the previously inserted instances in the drawing. Blocks can be nested. Blocks are effectively transparent to SheetCAM: the constituent objects of a block are interpreted just as if they were part of the drawing, with the serious exception

- a bug - that mirrored blocks are read as if un-mirrored.

Groups: Objects can be grouped so that they act as a single entity for purposes of arrangement in the drawing. However, groups should be avoided because SheetCAM ignores the contents of a group. If importing to TurboCAD a drawing that contains groups already, they can be disassembled into the constituent parts using *Format->Explode*. The grouping is probably useful, so convert each group to a block by first placing the group into a block, then editing the block and exploding the group. Note that groups can be nested, so it may be necessary subsequently to explode further groups. Take care not to mistakenly explode non-group objects - rectangles will explode to four lines, for example. If unsure whether an object is a group, select it and view its properties using <Ctrl>-F.

Save Selection: *File->SaveAs...Selection*. If a subset of a drawing is selected, then this subset can be saved as a file containing only the selection. This is useful for transferring a subset of parts to the CAM program. Take great care: do not overwrite the main drawing with just a subset of parts - so always choose a new file name when saving a selection. Be aware that immediately after the save-selection operation TurboCAD reverts to working on the *original* file, not the *newly saved* file.

Non-Cutting Layers

CAD users frequently draw lines and other objects that are not intended to be cut but which help the user to align components. These are known as construction lines. For example, a block that contains a drawing of a hole might include a 'X' at the centre of the hole, to help with positioning the hole in the correct place in the drawing. Clearly the 'X' would never actually be cut. Or a drawing of a gearwheel component might include the outline of the gearbox into which it fits, to guide the design process.

The user may also wish to include notes regarding a part in the drawing.

Such elements of the drawing should be put on separate layers. The layers can be hidden as necessary, both in TurboCAD and in SheetCAM.

Other Resources

All of this barely scratches the surface of TurboCAD, or indeed any other serious CAD package. Whenever apparently stymied by a requirement that the CAD package seems unable to satisfy, it pays to search the forums, scour the manual, and google for the solution. Bear in mind that TurboCAD, like most other popular CAD offerings, has been undergoing rapid development for many years. A forum posting saying that a particular feature is missing may in fact be referring to an older and less capable version of the program. It pays to check the date of any posting.

CAM

The saved CAD drawing must be converted to CNC machine commands that LinuxCNC can read. This conversion is carried out by CAM software. The resulting program is colloquially known as 'G-Code', after the letter which precedes the majority of the machine commands.

This example uses SheetCAM v6 to generate the G-Code. Adapt the example as necessary for your chosen software.

Once the G-Code program has been created, it can be run in LinuxCNC to cut material on SheetCAM. The material and the cutter have to be specified during the CAM process. We'll assume that you're cutting 15mm thick MDF. And we'll assume that a 6mm diameter 2-flute HSS cutter is in use. Adjust these as necessary.

Concepts

In SheetCAM each cutting operation which is required (say, 'cut around the outlines') is linked to a specific layer in the input drawing (say, 'OutsideThru'). In this way operations can be grouped and applied to a specific set of objects in the drawing. The beauty of this is that, once the operation-to-layer relationships have been set up in SheetCAM once, the CAD drawing can thereafter be modified and then re-imported into SheetCAM. SheetCAM will respond by applying [almost] all the operations to the revised drawing automatically. It is even possible to import a completely different drawing and, if it has layers of the same name, the operations will be applied in exactly the same way. It is hard to overstate how useful this feature is when iteratively designing a new product.

Other CAM programs may work differently, for example requiring the user to specify each operation on each object in the drawing, in which case the following procedure would need to be heavily revised.

Cutting Choices

The CAM user must decide what cutters ('tools') to use and at what depths and speeds to feed those cutters, for each operation. They must choose whether to use the ideal cutter for each type of cut, or compromise by using a single cutter for all the cuts to eliminate cutter changes.

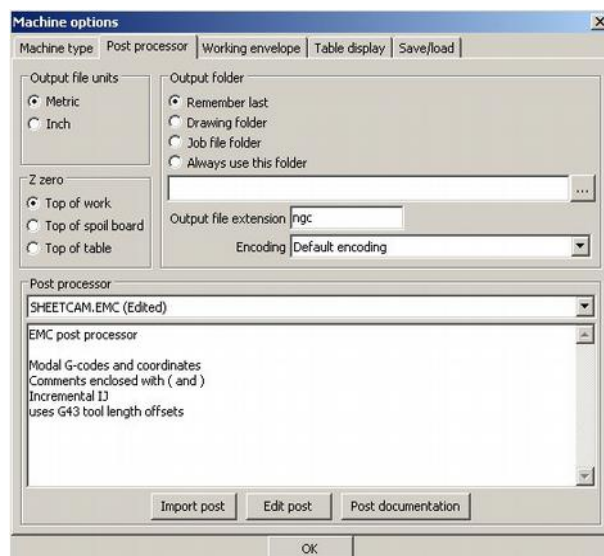
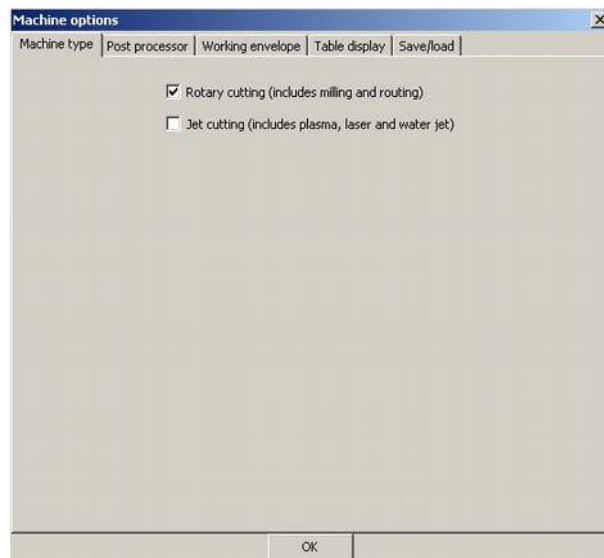
The user must also decide the order of operations. For example, it clearly makes sense to cut the outlines of the parts only after machining all the interiors of the parts, because once the outlines are cut the part will be too insecurely fastened to be worked on. Where cuts on different layers overlap each other there may be an optimum sequence which minimises burring or tearing-out on critical edges.

Setting Up SheetCAM for SheetCNC

A one-off setup process is required for a newly-installed instance of SheetCAM so that it can create files suitable for SheetCNC. We've assumed you are using a LinuxCNC machine controller.

Open SheetCAM.

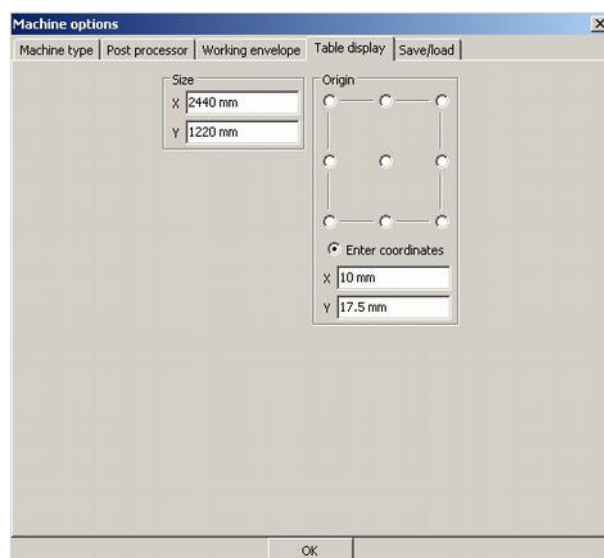
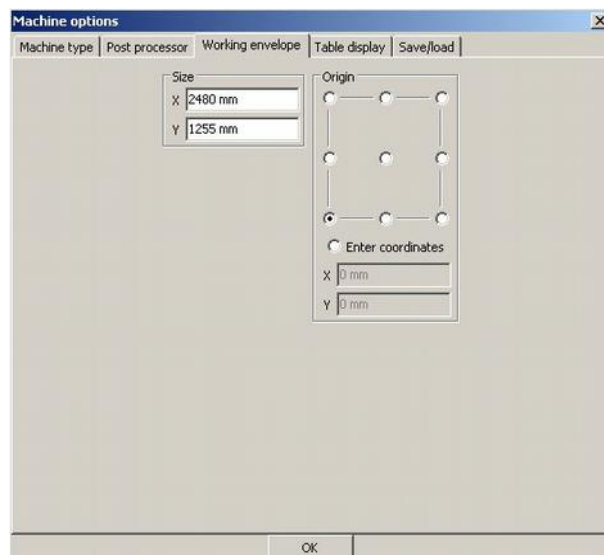
Set up the machine options to match SheetCNC:
Menu->Options->Machine...



From our webpage, download the file SC.SHEETCNC.EMC.scpst*. This is a post-processor which we have created to maximise the precision of SheetCNC's cutting operations. Save the file in an accessible location on your PC, for example in a folder 'SheetCNC' in your home directory.

Click *Import Post*. Navigate to the file SC.SHEETCNC.EMC.scpst and select it. SheetCAM will now use this file to define some details of the G-code which it generates, so that the G-code is optimised for SheetCNC. Note: The file SHEETCAM.EMC, provided with SheetCAM, will also work with SheetCNC but it will not be optimal.

* This assumes that you are using LinuxCNC as the machine controller. If using any other controller then you may need to edit the post-processor file yourself. Please contact us for details of changes to the post-processor that may be required. We make no warranty as to the suitability of the file or the changes we have made for your specific machine and setup. Verify the resulting G-Code before running it on any CNC machine.



Set up the application options. Our preferences are shown below, but yours may differ:
Menu->Options->ApplicationOptions...

Application options

Display | Units | Annoying dialogs | Drawing import | Toolbars | Default directories

Start point size 5 mm

Tab size 100 %

Direction arrow size 100 %

Minimum direction arrow size 2 mm

Direction arrow frequency 1

Circle marker size 4 mm

Colours

Cut path

Feed override

Rapid path

Inside shape

Outside shape

Open paths

Use black scheme

Use white scheme

Change

☒ Reverse mouse wheel

☒ Use bottom left coordinates

☐ Use fast rubberband graphics

☒ Negative angles are clockwise

OK Cancel Help

Application options

Display | Units | Annoying dialogs | Drawing import | Toolbars | Default directories

Linear units mm

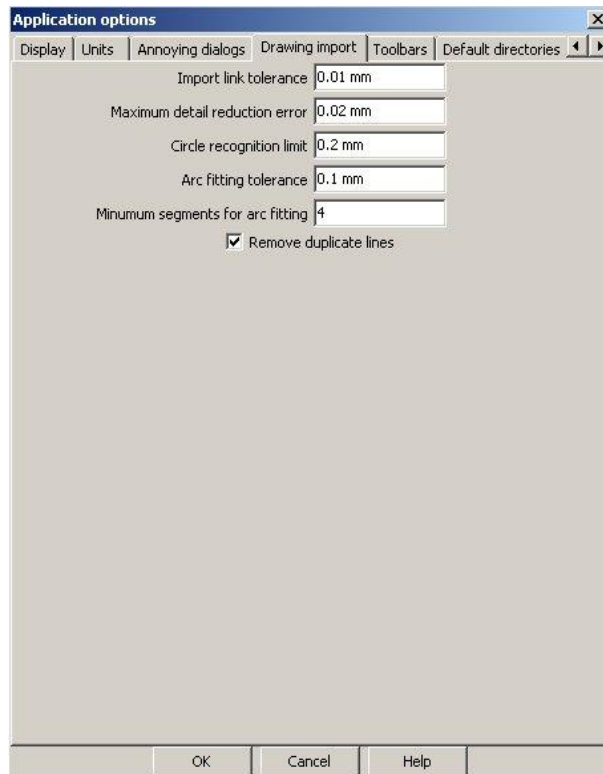
Angular units degrees

Feed rate mm/min

Thread pitch mm

Time s

OK Cancel Help



SheetCNC users may wish to download the SheetCAM ToolSet that we provide on our website, and store it in a suitable location (e.g. a folder SheetCAM in their home directory.)

SheetCAM should now be fully set up.

Creating G-Code from the CAD Drawing

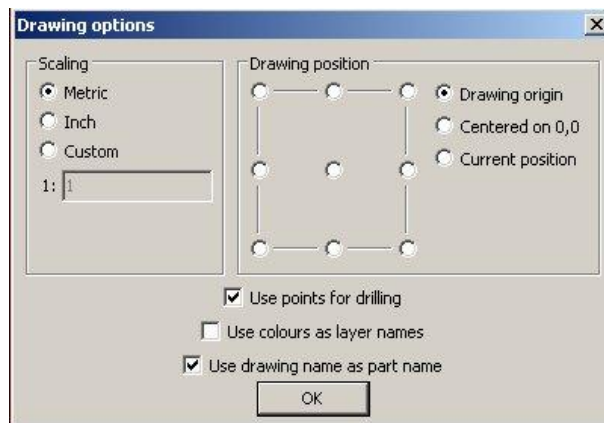
At the end of this section you will have G-Code files ready for cutting the part on SheetCNC. This example assumes the cutter is a 6mm 2-flute HSS end-mill, but if carbide tooling is available then use that for preference and adjust the settings as necessary.

Start the program SheetCAM.

Import your drawing file:

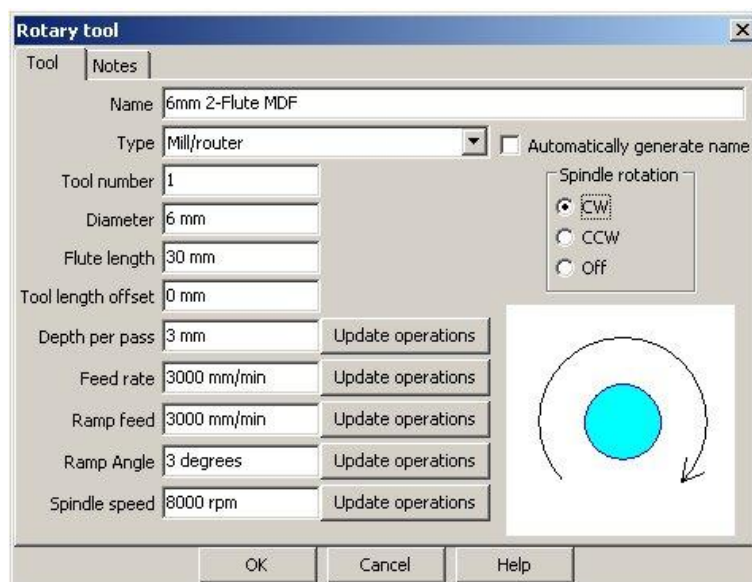
Menu->File->ImportDrawing...

...which opens a Drawing Options pop-up preparatory to importing the file...



Specify the cutter¹:

Menu->Tools->NewRotaryTool...

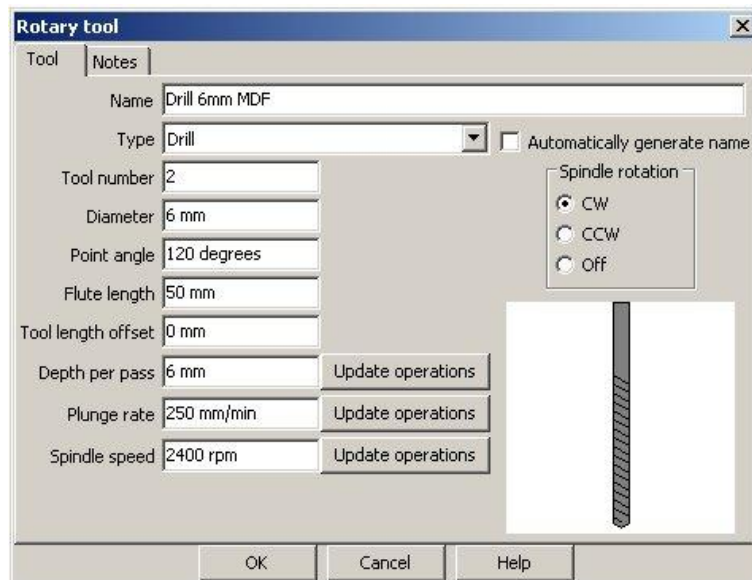


The maximum depth per pass is limited by the poor chip clearance rates of the HSS cutters. Here we specify a conservative 3mm depth per pass to ensure a clean cut.

Specify a drill tool:

Menu->Tools->NewRotaryTool...

¹ SheetCNC users may prefer to skip creating tool specifications. They can instead simply import the ToolSet which they previously downloaded from our website. Click *Menu->OpenToolset...* and navigate to the downloaded file.



Rotary tool

Tool | Notes

Name: Drill 6mm MDF

Type: Drill ☐ Automatically generate name

Tool number: 2

Diameter: 6 mm

Point angle: 120 degrees

Flute length: 50 mm

Tool length offset: 0 mm

Depth per pass: 6 mm

Plunge rate: 250 mm/min

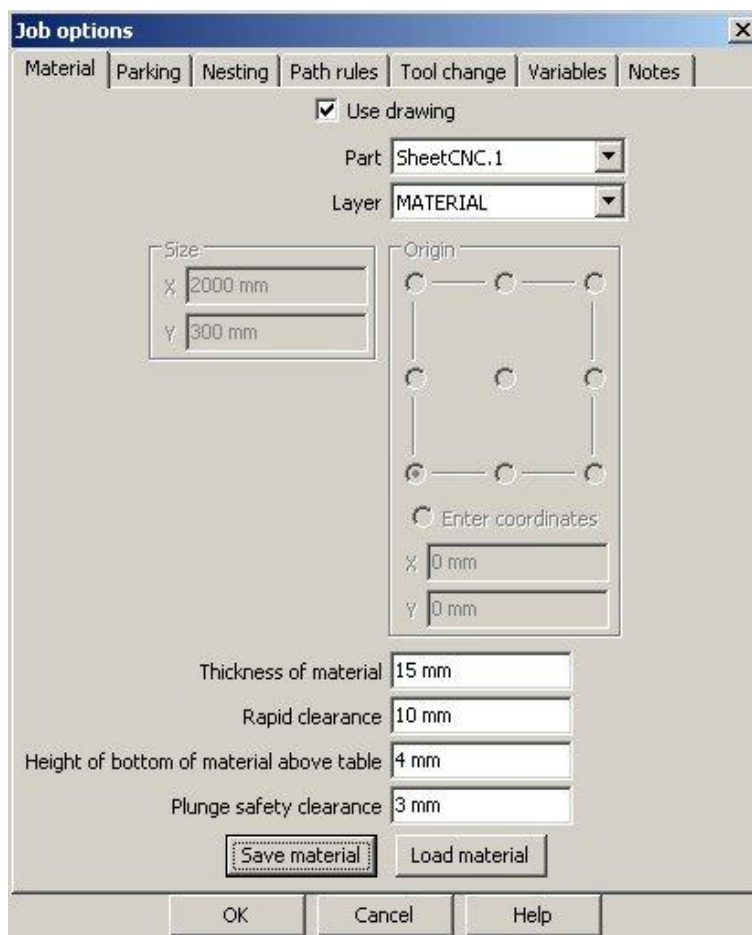
Spindle speed: 2400 rpm

Spindle rotation: ☒ CW ☐ CCW ☐ Off

OK Cancel Help

Specify the material, partly using information already in the drawing. Also specify where the cutter should be parked at the end of the program:

Menu->Options->JobOptions...



Job options

Material | Parking | Nesting | Path rules | Tool change | Variables | Notes

☒ Use drawing

Part: SheetCNC.1

Layer: MATERIAL

Size: X: 2000 mm Y: 300 mm

Origin: ☐ Enter coordinates X: 0 mm Y: 0 mm

Thickness of material: 15 mm

Rapid clearance: 10 mm

Height of bottom of material above table: 4 mm

Plunge safety clearance: 3 mm

OK Cancel Help

There is no need to click 'Save Material'. Just click OK.

Job options

Material | **Parking** | Nesting | Path rules | Tool change | Variables | Notes

Parking position

X: 0 mm ☒ Use X

Y: 0 mm ☒ Use Y

Z: 20 mm ☒ Use Z

OK Cancel Help

Note the Z position, to park the cutter clear of the workpiece.

Specify the first cutting operation. The layer InsideThru is to be machined around the inside of the profile.

Menu->Operation->Contour...

Contour

Basic | **Cut start** | Cut path | Notes

Contour method: Inside Offset

Layer: INSIDETHRU

Tool: T1: 6mm 2-Flute MDF

Start depth: 0 mm

Cut depth: 16 mm

Depth per pass: 3.3 mm

Finish allowance: 0 mm

Feed rate: 3000 mm/min

Ramp feed: 3000 mm/min

Ramp Angle: 3 degrees

Spindle speed: 8000 rpm

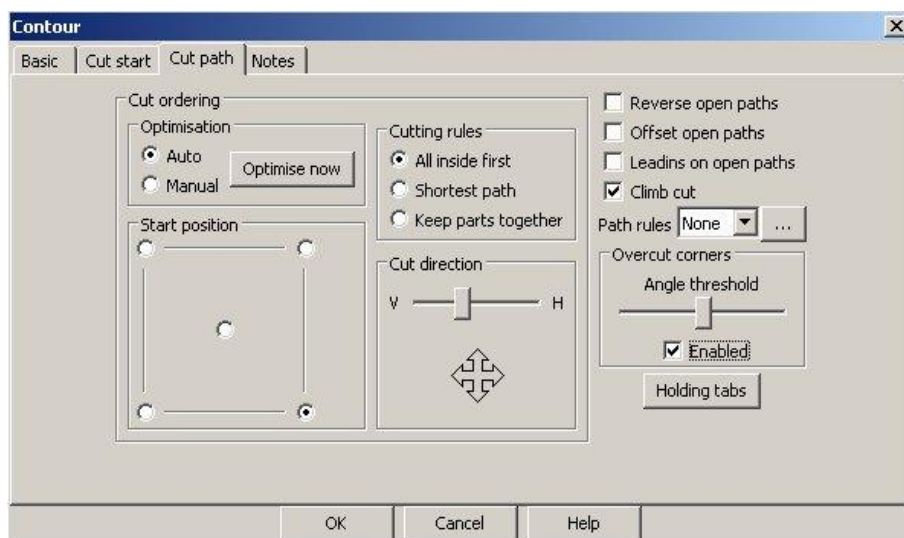
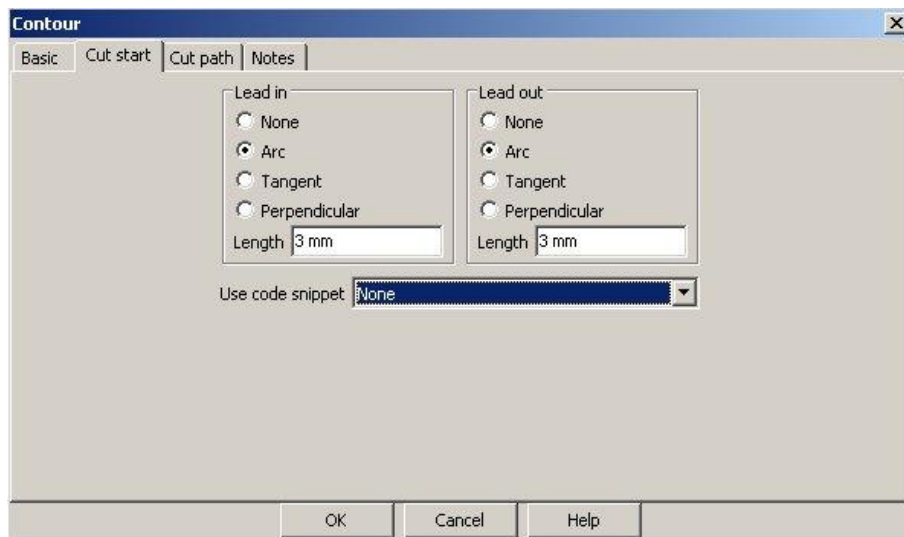
Coolant: ☒ None ☐ Mist ☐ Flood

Material thickness: 27 mm

Edit material

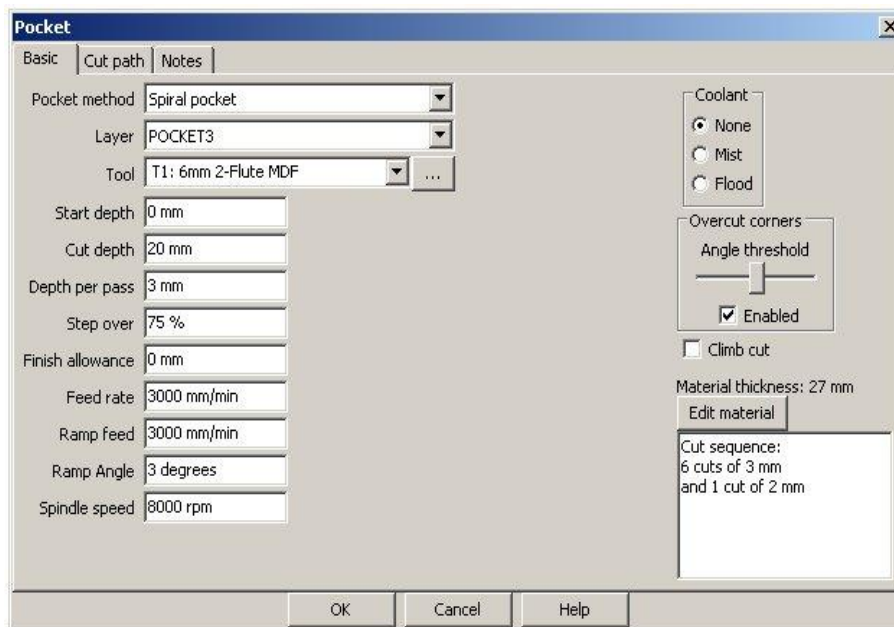
Cut sequence:
4 cuts of 3.3 mm
and 1 cut of 2.8 mm

OK Cancel Help

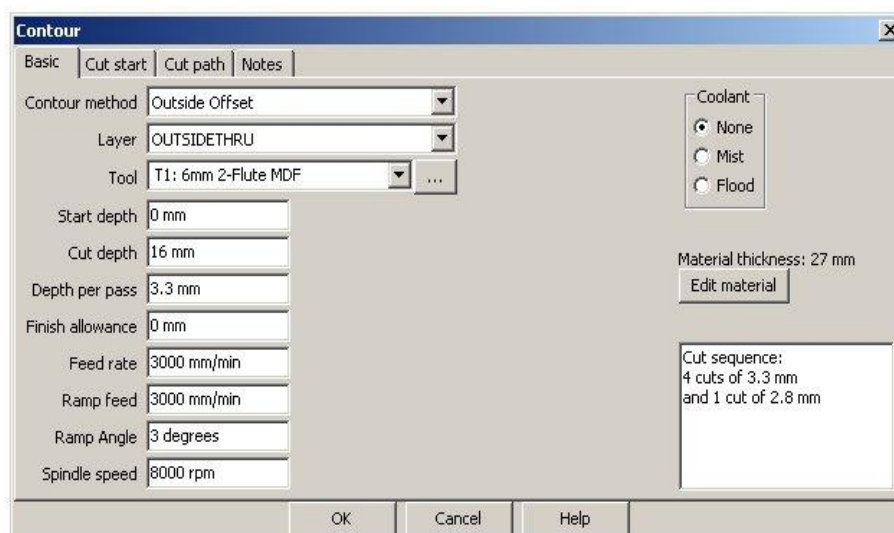


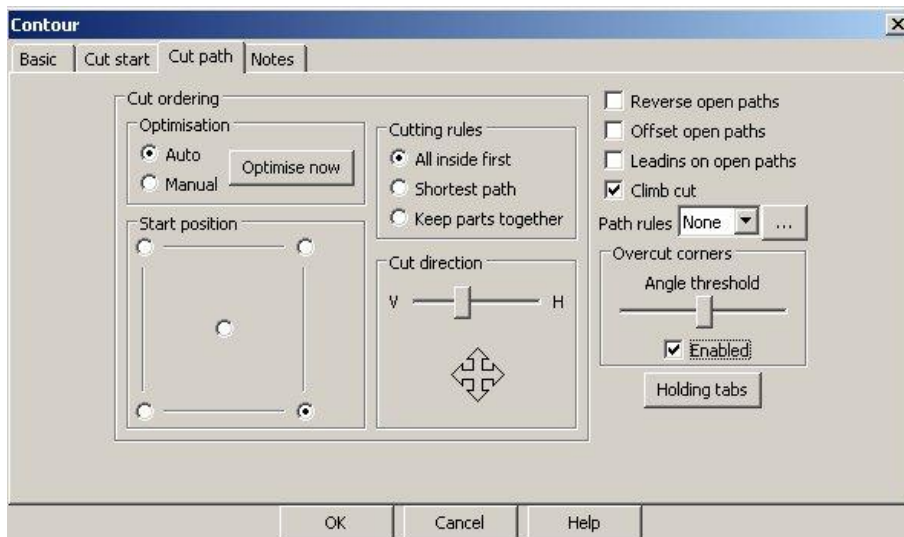
Since the profile to be cut happens to be smaller than twice the diameter of the cutter it's guaranteed there is no core which might break out and so there is no need for tabs on this operation.

The layer Pocket3 is to be pocketed to a depth of 3mm:
Menu->Operation->Pocket...



The layer OutsideThru is to be machined around the outside of the profile:
Menu->Operations->Profile...






Click OK when done. Do not click 'Holding Tabs' yet.

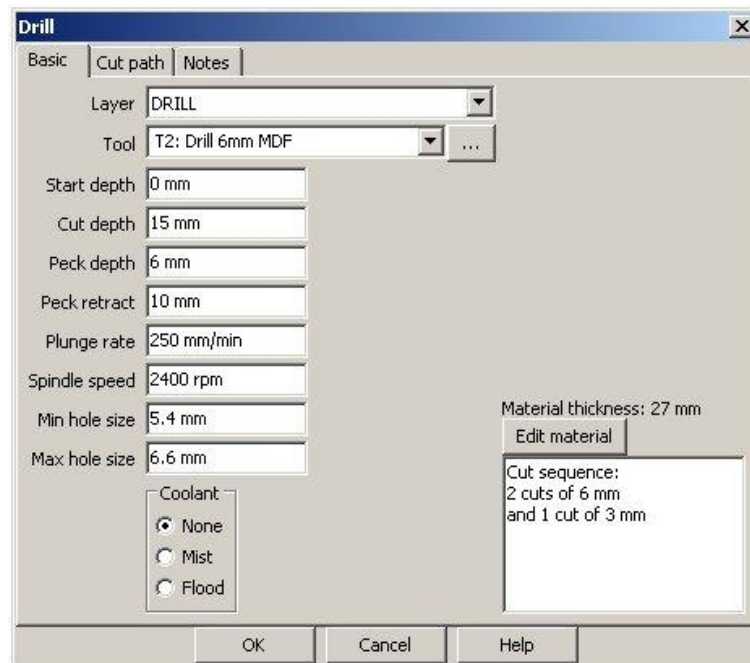
The outside profile requires holding tabs, to hold the part in place as the cut is finishing. Now that the path exists, it is possible to apply tabs to it. So re-open the previous outside offset operation by double-clicking on the 'Outside Offset, OUTSIDETHRU...' entry in the 'Operations' panel (bottom-left):

On the 'Cut Path' tab, click the button 'Holding Tabs'...



After entering the values, click 'Place Tabs', then 'OK'.

Add drillings for the fixings which hold the material in place during machining:
Menu->Operations->Drilling.



Time now to save the output files.

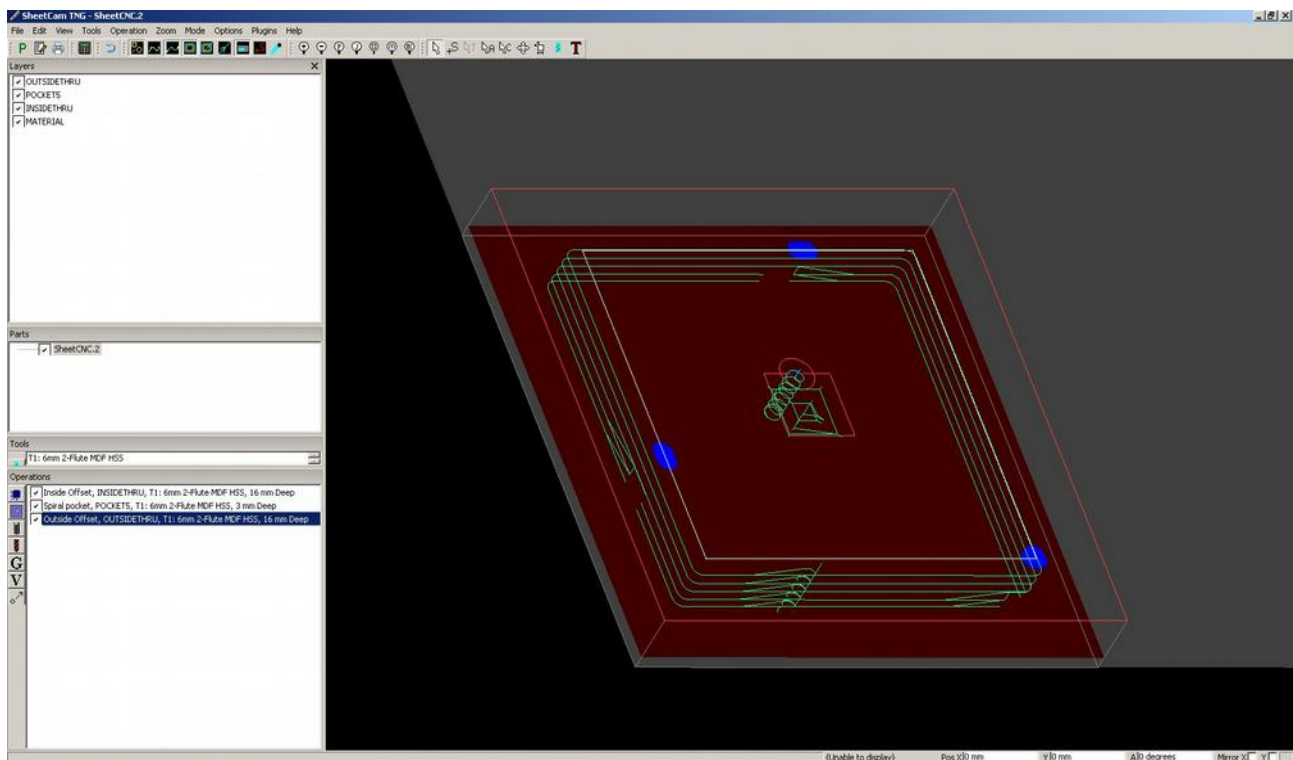
Select only the operation which is required for the fixing of the workpiece:
In the Operations frame (bottom-left) tick only the DRILL operation.

Run the post-processor on that single operation, to create the first program:
Menu->RunPostProcessor...
Save the G-Code as file SheetCNC.01.Drill6mm.FIX.ngc

Select only those operations that are required for the cutting of the workpiece, by inverting the layer selection:
Right-click (in the operations frame)->ReverseSelection
Double-check that the DRILL operation is not selected.

Run the post-processor on these operations, to create the second program:
Menu->RunPostProcessor...
Save the G-Code as file SheetCNC.02.6mmHSS.CUT.ngc

Before quitting, quickly review the operations: Middle-click-and-hold on the centre of the drawing, and move the cursor to adjust the angle of view. Left-click-and-hold to pan as necessary. The results should look something like this:



You may have to change some of the Menu->View settings to see a similar image. The red box is the notional workpiece as specified by drawing layer Material. The green lines are the programmed cutting moves. Direction arrows and rapid moves can optionally be shown too. Blue dots are tab locations - visible also as the lifts in the lowest few cuts of the outline. In the shallow pocket, the diagonal moves to overcut the corners can be seen. The short arcs on the nearest edge of the outside profile keep the cutter away from the finished edge during the start and end of each cut.

Save the job so that it can be reopened later:

File->SaveJob...

Etc.

Explanation

Some of the settings used above may need explaining. The notes below are in the same order as the screenshots above.

Machine Options Window

On the Post Processor tab, the Z-zero is set to be Top of Work. This means that when you touch off the Z axis you should do so by setting the tip of the cutter exactly at the height of the top of the workpiece material. Tip: There are good arguments for using Top of Spoilboard, or even Top of Table, especially if running repeated jobs cutting outlines from material whose thickness tends to vary.

On the Working Envelope tab, the size given is the total possible movement of the axes. I.e. the maximum rectangle which could be swept out by the tip of the cutter. The numbers should match the *Table Travel* values used during the setup of LinuxCNC with Stepconf. See *Spindle, Control and Power Systems*. The origin is at 0,0 which is the location of the X and Y home positions.

On the Table Display tab, the size given is the size of the bed of SheetCNC. The origin is slightly offset from the home location of the X and Y axes because the axes are over-size to allow for the radius of any cutter. The setting is for display only and has no effect on the G-Code produced by SheetCAM.

Drawing Options Pop-Up

The tick-box 'Use points for drilling' means that points inserted in the CAD drawing can be used as inputs to 'Drilling' operations in SheetCAM. Be aware that SheetCAM only recognises AutoCAD style points, which have to be specially selected when drawing in TurboCAD.

Setting the 'Drawing position' to be 'Drawing Origin' has the effect that the cutting operations will be performed on SheetCNC relative to the origin of the part in the CAD drawing. So if the parts were originally drawn at a random position in the CAD drawing then SheetCNC would move the gantry *from* the touch-off point *to* that random location before beginning to cut the job. This is usually undesirable so when working with CAD it pays to place parts in the drawing at the origin of the drawing before saving the DXF. Alternatively, the origin can be reset while importing the part into SheetCAM, using one of the other drawing position options on this pop-up.

Rotary Tool Window

Cutting depth per pass is typically about half the diameter of the cutter for these HSS bits. We've selected 3.3mm just so that the 16mm finished depth can be reached in five cuts instead of the six which would be needed for a 3mm cut. *Tip: for faster cutting, obtain some carbide tooling: hardwood sheets of 27mm or more can be cut in a couple of passes with a 6mm carbide cutter.*

Feed rate of 3000mm/min is typical for MDF.

A spindle speed of 8000rpm gives a chip-load of about 0.19mm with a 2-flute cutter.

These settings have the additional benefits that they are both quiet and produce very little airborne dust, so they are ideal as starter trials.

Job Options Window

Rapid Clearance is the height above the workpiece that SheetCAM will move the tip of the cutter up to, before horizontally moving between cuts at the full speed of the machine. This is a safety clearance to avoid uneven material, clamps and other obstructions. Even if you use flush clamps, consider leaving a reasonable margin to allow for possible large chips which might come away from the material and rest on the surface. A particular issue occurs when machining inside profiles if insufficient tabs are used: the core of the cut-away material may lift above the level of the remaining material and later foul a rapid move.

The Height of Bottom of Material Above Table (4mm) may seem wrong, given that the specified 9mm spoilboard appears to allow 9mm between the bottom of the workpiece and the top of the table. However, this height is used by SheetCAM to warn if the cutter is going to move low enough to foul the table. And 'the table' ought to include any items that are higher than the actual table-top that also might foul the cutter. In particular the screws that hold the spoilboard in place. These screws are countersunk 4mm into the spoilboard which means the usable depth of spoilboard is 4mm.

Plunge Safety Clearance is used when the cutter is descending from the Rapid Clearance height ready to start a cutting operation. The cutter is initially lowered rapidly, but below the Plunge Safety Clearance height the speed is reduced to the same as that used when knowingly cutting material. This allows that the material top surface may be uneven. For MDF, 3mm is extremely conservative.

Contour Window (Inside Offset)

Many of the settings on this form are automatically copied from the selected tool.

Note how the layer names chosen in CAD usefully prompt the user which operation to apply to each layer in CAM. This helps reduce operator errors.

Contours intended to penetrate right through the material are cut to 16mm Cut Depth, even though the material is specified to be only 15mm. This means the bottom edge of the material will be cleanly cut by the side of the cutter, rather than by the tip.

The Climb Cut box is ticked. This means that as the cutter moves through the material, its front edge is rotating towards the finished surface, rather than away from it. With most timber products this produces a better finish with less splintering. In MDF the effects are less pronounced but on grained materials such as plywood the difference is significant.

Finish allowance is 0. Finish allowance is the amount by which the cutter will be offset away from the intended cut. This setting is used when the part is going to be cut in two operations: a 'roughing' operation to approximately cut out the part, and a 'finishing' operation to cut the part to its exact, finished dimension. (Do not confuse the finish allowance with the offset that is automatically applied by SheetCAM to allow for the cutter radius.) In this case, there is no subsequent finishing operation, so no allowance needs to be set.

Pocket Window

A spiral pocket is cut by the cutter repeatedly looping at a given depth, each time moving sideways further out than the previous time. The Step Over setting defines how far the cutter will be moved sideways each time, as a percentage of the diameter of the cutter. Some overlap is desirable, especially if cutting a part-depth pocket.

Overcut corners are specified. This is useful when cutting a square mortise that must accommodate a square tenon: the cutter is moved sideways into the corners of the mortise to allow room for the square corners of the tenon. Otherwise the finished corners of the mortise would be the same radius as the cutter and would foul the corners of the tenon.

Holding Tabs Pop-Up

Holding tabs are used to prevent a cut-out piece - the finished part - from moving after the cutter breaks through the bottom of the material to part it off from the workpiece. Without tabs the cut-out would detach from the material, lift up, and be damaged itself and/or cause damage to the cutter.

The tabs are specified to be 2.5mm thick. But the material is known to be 15mm deep, whereas the depth of cut is 16mm, so only 1.5mm of that tab thickness will be in the material. The remaining 1mm will be lost in the spoilboard. Tip: tabs of just 1mm thickness can be used successfully on production jobs - a 1.5mm tab is very secure but more difficult to clean off the part afterwards.

Saving Multiple G-Code Files

The example creates two G-Code program files from the single imported drawing. The first G-Code file contains only the fixing-points where the material will be screwed to the table. The second G-Code file contains the actual cutting instructions for machining of the part. This allows the fixings to be precisely positioned in locations that will not conflict with the cutter during the machining. (See *Cutting the Part* below.)

It is critical that the cutting program does not include, in the output G-Code, any move that intersects the fixings used to hold the workpiece to the table. *A possible serious error would be to mistakenly include the fixings layer/operation when saving the cutting program. This would result in SheetCNC trying to drive the cutter vertically through the centre of each of the fixings, destroying the cutter.*

Other SheetCAM Actions

The SheetCAM program has more modes and actions to explore. Some of them are listed below.

Tip: When a mode other than normal selection mode is active then SheetCAM permits only a 2D view. In addition the left mouse button is then dedicated to the mode's action while the middle mouse button is repurposed for panning the view.

The start point of each cut can be specified.

Menu->Mode->EditStartPoints

Tabs can be moved, deleted, and inserted individually or in bulk.

Menu->Mode->EditTabs

SheetCAM can reload a modified drawing almost seamlessly.

Return to the CAD package.

Modify some of the objects in the drawing.

Add new objects.

Add a new layer and place some new objects on the layer.

Save the file - either with the same filename or a new filename.

Return to SheetCAM.

Menu->File->ImportDrawing

Etc

Observe that the changes on existing layers are automatically incorporated in the cut paths.

Add a new operation for the new layer.

Observe that, if significant changes are made to the shape or position of an object, the tabs may need recreating.

Cutting the Part

The G-Code created by SheetCAM is ready to be cut by SheetCNC.

Transfer the files to the LinuxCNC machine controller. If you can set up home networking and file sharing on the LinuxCNC machine via Ethernet, then that's ideal. Tip: installing the package `cifs-utils` is the first step. Wireless networking may be somewhat trickier to set up.

For the time being, a memory stick is likely to be the most convenient form of file transfer.

If in any doubt about the following procedure or the validity of the programs to be run then dry-run the programs with no cutter in the machine and with the VFD powered down.

Before powering up the machine, make sure it is safe to do so and don PPE.

Procedure

The following description assumes that your installation does not yet have a touch-off probe available.

Cutter and workpiece must exactly match those previously specified in SheetCAM. The spoilboard must allow at least the clearance below the bottom of the workpiece that was specified in SheetCAM. The touch-off point should exactly match that specified in SheetCAM, most especially in the Z direction.

Place a workpiece on the spoilboard. If it needs temporary clamping, use a screw(s) into the table driven through parts of the material that are *well clear* of any location where the cutter may pass.

Insert the first cutter - a 6mm HSS twist-drill bit - in the chuck, power up the machine, start LinuxCNC, and home the machine.

Use the LinuxCNC interface to manually drive the axes so that the cutter is above the origin point of the program to be cut. The origin point for the programs created above is 'bottom left' meaning that all cutting will take place at positive offsets of X and Y from the touch-off point. When the cutter is at the chosen position, touch off the X and Y axes in LinuxCNC (use the X/Y radio buttons and the 'Touch Off' button.)

Move the cutter so that it is approximately over the centre of the area to be cut. Using *small* movements, move the cutter tip downwards until it just touches the top of the workpiece. Then touch off the Z axis in LinuxCNC. There are several ways to do this. A popular method is to put a small scrap of paper on the workpiece and, with the spindle turning slowly, gradually lower the cutter until the paper is moved by the cutter tip. Alternatively, with the spindle unpowered, the paper can be used as a feeler gauge. An allowance for the thickness of the paper is normally made. Tip: that allowance can simply be entered as a positive number when touching off.

Raise the cutter well clear of the workpiece.

Power up the VFD and the dust extraction, if not done already, and run the program *SheetCNC.01.Drill6mm.FIX.ngc* . This will drill two holes through the workpiece. Then it will stop and park the cutter.

Check the cutter is safely clear and, using the LinuxCNC controls, manually drive the gantry away from the workpiece. Screw two #8x2" woodscrews through the holes, through the spoilboard and firmly into the machine bed to clamp the material. Ensure the screws are flush with the surface.

Power down the VFD. Remove the 6mm drill bit from the spindle and insert the second cutter - a 6mm 2-flute HSS end-mill.

Load the program *SheetCNC.02.6mmHSS.CUT.ngc* . Do not run the new program yet.

Re-touch-off the Z axis for the second cutter, just as for the first cutter. Do not re-home, re-touch-off or otherwise alter the X and Y axis calibration.

Power up the VFD and the dust extraction and run the program. This will cut the finished part.

Turn off the VFD. Check the cutter is clear and, using the LinuxCNC controls, manually drive the gantry away from the workpiece. Vacuum away all chippings. Unscrew and remove the workpiece from the bed. Working from the top of the workpiece, press the part out through the bottom. Tip: Pressing out the other way results in damage to the under-surface of the part where the tabs tear out. If the material must remain screwed on the table for cutting another part, the part's tabs can instead be freed by gently levering the part *sideways* (not upwards) with a large flat-bladed screwdriver or similar blunt object.

Remove the tab remnants with a suitable tool. We find that placing the part flat onto the spoilboard and using the edge of a medium-sized woodcutting chisel is convenient.

Follow Up

Examine the part.

Observe the machining marks that occur especially at the start-point where the cutter first starts each cut. Also the machining marks at the tabs where the cutter moves vertically. Note the linear marks around the part made by each pass of the cutter. Consider how these could be reduced so as to improve the part:

- Limit the number of tabs.
- Move tabs to locations where machine marks are less obvious - corners perhaps.
- Move the start-point to a corner and/or to coincide with a tab.
- Use the lead-in and lead-out features of SheetCAM.
- Adjust the tab thickness so that the tabs are only in the lowest layer. It may be necessary to thin the tabs or to adjust depth per pass so that the lowest layer is thick enough to accommodate the whole tab.
- Turn off the exit ramp cleaning after the tabs. It may be necessary to shorten the tabs and/or increase the ramp angle to compensate.
- Increase the depth per pass.

Experiment with other cutting speeds, depths and other settings. See the chapter *Chips* before making significant changes to these parameters.

Cutter Breakage

If cutter breakage occurs part-way through a program, stop the program by pressing <ESC> on the machine controller. This ends the program but maintains the touched-off position of the machine. Remove any broken parts of the cutter from the workpiece. Fit a new identical cutter and touch it off (*Z only*.) Use SheetCAM to recreate the G-Code program but excluding any layer that has already been completely cut successfully. Load and run the shortened program to complete the cutting-out of the part(s).

Chips

When cutting material on a CNC milling machine with a typical end-mill cutter, the waste material comes away as chips which are sliced off the material by the edges on the side of the cutter. Cutting timber products inevitably also produces dust as the chips break apart, but it is the chips themselves that we're really interested in here.

A quick bite of often-confused terminology: A cutter has *cutting edges* - the sharp parts along the side of the cutter which slice away material. It also has *flutes* - the slots between the cutting edges. The number of cutting edges on a cutter is normally denoted by the number of flutes - 2-flute, 4-flute, etc.

Chip Load

The thickness of chips generated during machining is known as *chip load*.

The chip load can be calculated fairly simply:

Chip Load = FeedRate / (RPM * EdgeCount)

Chip Load = thickness of each chip in mm

FeedRate = Speed of travel of the cutter through the material in mm/minute

RPM = Rotation speed of the cutter in rev/minute

EdgeCount = Number of cutting edges on the cutter

In reality each chip will vary in thickness along its length, but this calculated chip load value remains a useful measure of how hard the cutter is working. There are serious disadvantages to both very large chip loads and, perhaps surprisingly, very small chip loads.

As each chip is cut away from the material, it removes some of the heat that is generated during the cutting process. Bigger chips mean better heat removal. This is a good thing, because excessive heat causes damage to both material and cutter.

If chip load is very small then the cutter is trying to slice incredibly thin layers off the material with each pass of each cutting edge. Most passes will, in fact, fail to cut any material at all and will instead simply cause the cutting edge to slide across the surface of the material doing nothing but creating heat and polishing the surface. This burns the material and damages the cutter. So very slow feedrates combined with normal spindle speeds will burn out the cutter. This is a common novice error.

If chip load is very large then the cutter is trying to slice large bites out of the material with each pass of each cutting edge. This puts heavy loads on the cutter, risking damage to it or even stalling the spindle. At extremes, the chips can be so large that they clog the flutes of the cutter, creating friction between the clogged chippings and the material. This generates heat and, again, burns out the cutter.

A happy medium is needed. For initial setup when machining timber products, a chip load of 0.2mm is likely to be sufficient to avoid overheating while not being so great that the cutter becomes overloaded or clogged.

Trade-Offs

The formula above suggests that - for a given chip load - feed rate can be increased by increasing the number of cutting edges on the cutter and/or by increasing the spindle rpm. Higher feed rates are desirable to reduce total machining time. However, there are other considerations.

Increasing the Number of Cutting Edges

Using a cutter with twice the flute-count and doubling the feed rate while leaving the RPM constant results in the same chip load and the same number of chips per flute. Which seems like a win.

Unfortunately the flutes themselves are less than half the size on a 4-flute cutter compared to those on a 2-flute cutter. This means that chips clear along the flutes less easily resulting in a greater propensity to clog.

In addition the loads on the cutter are roughly doubled too - after all it is now cutting twice as much material in the same time. This means that cutter deflections are increased leading to reduced accuracy.

Small diameter cutters are often made with single flutes to allow room for better chip clearance.

Increasing Spindle RPM

Doubling the spindle speed and doubling the feed rate, without changing the cutter, results in the same chip load and the same number of chips per flute. Another apparent win.

However, heat generation is at least doubled too, and the ability of the extra chips to carry the heat away may be insufficient to compensate.

Furthermore, the feed rate has to be slowed ready to turn sharp corners or reverse direction at the ends of slots: the machine controller does this automatically so as to stay within the acceleration limits of the machine. The reduced feed rate during acceleration and deceleration results in a decrease in chip load. The higher the dialled-in feed rate then the longer the distance the machine needs for deceleration and so the greater the amount of material that is cut at suboptimal chip load. Worse still, the higher the spindle speed the more quickly the heat will increase. So doubling the spindle speed will much more than double the excess heat generated at those points where the machine needs to decelerate.

With careful design, so that sharp corners and other causes of intermittent slow feed rates are eliminated, it is possible to use higher spindle speeds to increase the feed rate and so reduce the job time. The time investment needed to modify the CAD design, and the possible compromising of the product due to lost features, means that these kind of adaptations are worthwhile only for large production runs.

Even if the excess heat is safely dealt with, the loads on the cutter are still doubled and so cutter deflections are again increased. This means still more time invested in selecting suitable finishing clearances to compensate for those cutter deflections.

Higher spindle speed raises one more issue - vibration and noise. Generally the higher the rotational speed the more vibration suffered by the cutter. This creates noise, which is a nuisance. It also means that the cutter is cyclically moving toward and away from the finished surface. These

movements can show up as waves or ripples in the finished surface, reducing the finish quality.

Increasing Depth of Pass

An alternate way to reduce job time is to increase the depth of cut per pass. Since this leaves rpm, feedrate and number of flutes unchanged it has no direct effect on chipload. Increasing depth of cut reduces the number of passes required when cutting deep pockets or profiles and so speeds up the job. Indeed, massive commercial machines will routinely cut boards in one pass. There are even special 'compression' cutters available for this purpose. However, this kind of performance is not within the capabilities of most home-built CNC machines. Even if it were, there are still issues.

Doubling the cut depth results in twice as much material being cut by each cutting edge in the same amount of time. Some of this cut material is cleared away to behind the cutter. However, due to the effective expansion of the volume of cuttings compared to the raw material, much of it has to rise up the flutes and clear at the surface of the material. Doubling the cut depth means the cuttings have, on average, twice as far to travel along the flutes before reaching the surface. So there will be some depth limit, beyond which chipload will have to be reduced in order to prevent clogging from these extra cuttings in the flutes. This means increasing spindle rpm or reducing cutting speed, which leads to the same problems as above. And, worse, the cutter is already generating twice the heat in the first place because it is cutting twice as much material, so overheating will - yet again - occur much more rapidly at those points where the machine has to decelerate.

And, again, even if the excess heat is safely dealt with, doubling the cut depth means that the loads on the cutter are doubled and so cutter deflections are again increased.

Despite the drawbacks, doubling the cut depth not only halves the job time but also halves the wear on the cutter tip - which is the part of the cutter that usually wears out soonest. So it is always worth working with the maximum depth of cut possible.

Roughing and Finishing

One final option is to use separate roughing cuts and finishing cuts in the machining of a job.

The parts can be cut quickly and roughly using aggressive settings but leaving material 'spare' on the cut edge. The machine can then be instructed to make a finishing cut around the edges at lower speed.

Roughing and finishing in two operations has other advantages. The final cut will remove very little material so cutter loads are small resulting in minimal cutter deflections and better accuracy. The final cut can be a very thin 'slice' off the edge but in a single pass over the whole depth of the material, resulting in a better, toolmark-free finish. Running two operations would also allow the use of two different cutters - a fast-cutting roughing cutter, followed by a slow-cutting finishing cutter. In between the operations the workpiece can, if necessary, be cleaned of dust and cuttings to further improve the finish achieved by the finishing cutter.

There are some operations where roughing is not possible - for example when cutting a slot in the part when the slot is the same width as the diameter of the cutter.

Materials and Settings

Most users will wish to run their machines at the best possible speed. Where 'best' could be defined as the maximum cutting speed that gives the required finish within the dimensional tolerances allowed while wearing out cutters at an acceptable rate.

There are, of course, so many variables that a prescription of the best speed is just not possible. Bear in mind, too, that cutter specification covers much more than merely the size and number of flutes. These specifications have a huge bearing on the performance of the machine.

There is far too much to cover in this document. Besides, the finish required and dimensional tolerance of each project that the user undertakes will vary, resulting in a variety of optimum cutters and settings depending on the work being carried out.

Materials for First Projects

We suggest initially experimenting with plain unfinished MDF, starting with 12mm to 18mm sheet. The thicker sheets are more rigid, which means fewer hold-downs are required for machining a given area. No special techniques are needed for machining MDF. The boards are not usually sided - being similarly finished on both sides. We find that MDF is a superb engineering material, being of very consistent dimensions, having a smooth hard surface, and being predictable to cut. Sheets are most economically sourced from timber merchants, with builders merchants a more expensive option, and DIY sheds more expensive still. Water-resistant varieties are available.

We also suggest obtaining a length or two of builder's general purpose rigid uPVC board. This is often referred to as foam soffit board, to distinguish it from the cheaper hollow board. The board is available in white at nominal 10mm thickness in widths up to 500mm and lengths of 5m. We find that this lightweight and versatile material is superb for making gears and gadgets. It is generally very economic to buy from local building plastics specialist suppliers. The sheets are best cut upside-down to minimise spalling from the hard fair face. A reasonable number of hold-down fixings are required because the material is somewhat flexible.

Gain experience on MDF before moving on to cutting plywood, especially hardwood ply such as the ubiquitous 'red-and-white'. Hardwood ply requires carbide cutters, while MDF can - just - be cut using HSS cutters.

Tip: When buying materials from trade merchants, phone ahead and ask for a quote. The price can be up to 50% cheaper on the phone compared to going direct to the trade counter. Delivery will, of course, be an option at all trade merchants. Merchants are always happy to see new customers, discuss the types of sheet material they keep, and show customers the material before purchase... so long as the new customer doesn't turn up five minutes before lunch hour or half an hour before closing time on a Friday. And these days, almost all merchants will serve non-trade customers on a cash or card sale basis - there is no need to be a trade customer or to hold a trade account with the merchant.

First Settings for SheetCNC

We offer some discussion here on choosing settings that suit the starter cutters for new users listed

in *Spindle, Control and Power Systems*.

The settings are generally conservative but are good starting-points for experimentation.

Chip loads in the range 0.1mm to 0.25mm work well enough for general milling work on SheetCNC for most plastic and timber sheet products, when cutting with the HSS cutters we've suggested using for first trials.

A table below summarises our suggestions.

Choice of Cutter Dimensions

The primary consideration is that - for upcut spiral cutters - the flutes on the cutter must be at least as long as the deepest feature to be cut. Even if the feature is cut in multiple passes, the flutes must still reach beyond the top surface of the material so that chip clearance can be effective. This applies even for finishing cuts where the side of the cutter may be open. On the whole, small diameter cutters tend to have shorter flutes which limit the depth of material they can cut. Tip: Vendors of budget cutters frequently confuse cutting edge length and flute length. Flute length tends to be longer than cutting edge length, though on most budget cutters the difference is slight.

Secondly, a rotating cutter moving through material has sideways forces acting on it at right-angles to the direction of movement, due to the front of the cutter rotating against the material it is cutting into. This causes the cutter to bend sideways resulting in reduced accuracy of cut and - in extreme cases - breakage of the cutter. So the cutter must be substantial enough to withstand these forces. Clearly, all other things being equal, a longer cutter will need to be a larger diameter to resist the bending. When cutting soft material the sideways forces are reduced so a smaller diameter cutter can be specified.

Finally, a smaller diameter cutter permits more precise machining of tight inside corners. A smaller cutter forces less overcutting on mortise corners too. Precise components - such as gears - can be manufactured smaller.

In general, start by using the 6mm cutters for MDF work and using the 1/8" cutters for cutting 10mm uPVC boards. The 1/8" cutters can also be used for thinner MDF sheets.

Depth of Cut per Pass

With a 6mm 1-flute carbide end-mill, SheetCNC can cut 27mm MDF or hardwood ply in a couple of passes at 3000mm/minute or faster. However, we strongly recommend that you start with much less aggressive settings.

Cutter specification and the nature of the material being cut influence the depth of material that can be cut in each pass.

See the table below for recommendations.

Plunging and Drilling

End-mills can be used to drill lightweight materials such as plastic, but never plunge cut timber products with an end-mill as this causes rapid heating resulting in blunting of the cutter tip, quickly destroying the cutter.

Summary

Suggested settings for new SheetCNC users. Begin with the settings shaded cyan.

Material	Cutter	Spindle Speed (rpm)	Feed Rate (mm/min)	Depth of Cut per Pass	Ramp Angle (degrees)
Plywood / MDF	1-flute 6mm carbide	23000	2800 - 3000	13.5mm	6
Plywood / MDF	1-flute 6mm carbide	16000-19000	3000	3mm - 6mm	6
MDF	1-flute 1/8" carbide	12000	3000	2mm	6
MDF	1-flute 1/8" carbide	15000	3000	3mm	6
uPVC	1-flute 1/8" carbide	12000	3000	10mm (full material thickness)	90
uPVC	1-flute 1/8" carbide	24000	6000	10mm (full material thickness)	90
Any	Any twist drill used for drilling operations	Choose parameters just as for the same drill in a regular drilling operation such as in a pillar drill. Do not use large diameter drills - the spindle will have insufficient torque at the required low speed. Example: a 6mm twist-drill works well in 25mm MDF at 2400rpm and 400mm/minute, peck depth 4mm, with retract equal to drilling depth minus peck depth.			

All settings are adaptable with experience. As a rule of thumb: if changing the feed-rate then change the spindle speed approximately in proportion so as to keep the chip-load constant.

There is a trade-off between the time it takes to finish a job, and the precision of the result. Faster more aggressive cutting speeds and depths impose much greater forces on the cutter, causing deflections. Therefore, any one job is likely to combine high-speed operations for medium-precision cuts, such as the outlines, and medium-speed operations for high-precision cuts, such as holes and slots.

Probing

Probing means automating one or more aspects of touching-off to the workpiece.

The degree of precision when touching-off in the XY plane is generally not critical to success. It is a simple matter to measure the distance of the cutter from the corner of the workpiece and touch off appropriately. So - for the majority of work - probing is not greatly useful for X and Y touching-off.

The Z-axis, on the other hand, needs to be touched-off very precisely because the depth of pockets depends on the accurate vertical positioning of the tool. Furthermore, any one job may require multiple cutters and each cutter has to be touched-off separately. Automated touching-off for the Z-axis is therefore a great bonus.

Probing The Z-Axis

An effective means of probing for the Z touch-off point is to place a conductive probe plate on the top of the workpiece directly below the tool, then have the machine slowly lower the tool until it detects - electrically - that the tool has touched the probe plate. The Z-axis can then be automatically touched-off at that position, after making due allowance for the thickness of the probe plate.

Typically the conductive probe plate will be a piece of heavy brass sheet, about 75mm square and a few millimetres thick.

To set up probing requires some setting-up of LinuxCNC plus the physical addition of a probe.

Setting Up LinuxCNC

We have already provided all the software necessary to enable the probe with the configuration files we offer for download. See *Spindle, Control and Power Systems* for instruction on installing these files.

You will need to edit the files *120.ngc* and *130.ngc* and therein change the declared value of the probe plate thickness to be equal to that of the plate which you choose to use, unless your plate is the default thickness of 3.25mm*. The files are found in the directory `/home/sheetcnc/linuxcnc/nc_file/`. The values which need to be changed are declared at the top of each file in a line which reads `#<platethickness> = 3.25` and the changes required to those lines are self-evident. Any text editor, such as *mousepad*, will suffice.

The Probe Plate

A physical probe plate must be created and installed. These instructions assume you've used a BL-MACH-V1.1 0302 breakout board, as described in *Spindle, Control and Power Systems*. Adapt the instructions if some other board has been used.

Parts:

- Resistor, 10kOhms with a short length of instrument wire soldered to one lead.

* Measure your probe plate: do not rely on the nominal thickness being correct.

- Capacitor, tantalum 4.7 μ F rated at least 16V.
- Length of red flexible wire sufficient to reach from the breakout board to SheetCNC.
- Metal plate. Something heavy enough to reliably weigh itself down flat onto the workpiece is required. Brass 3mm x 75mm x 75mm is ideal.
- Length of green-and-yellow wire sufficient to reach the spindle body (see text.)

All the parts are cheaply available from eBay.

Shut down the machine controller and turn off all electricals before continuing.

Solder or otherwise connect one end of the length of red wire to the top of the brass plate. Ensure both faces of the plate are flat and smooth, and are free from lacquer or other insulation.

Twist together the other end of the red wire, the end of the wire soldered to the resistor, and the '+' leg of the capacitor. Insert these into the P13 terminal on the breakout board and tighten the terminal.

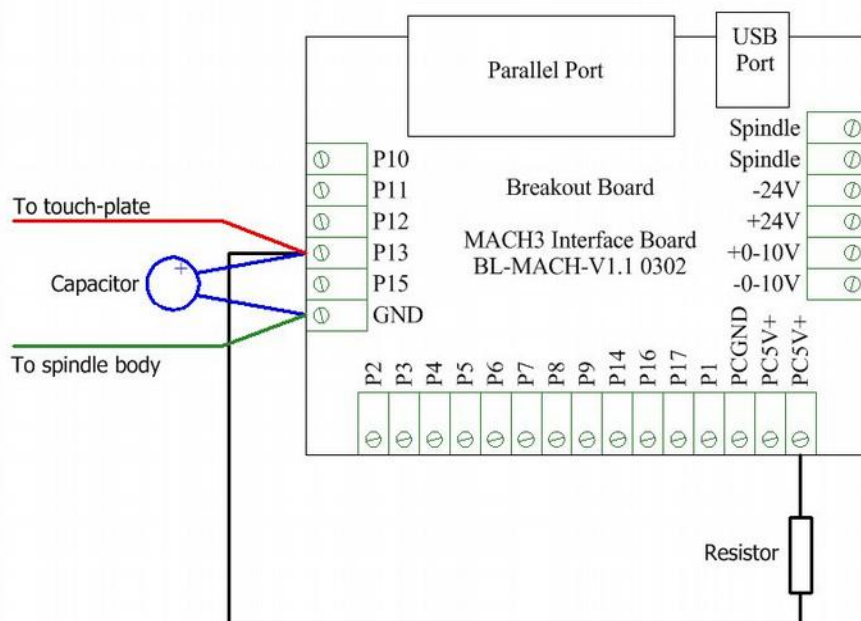
Insert the '-' leg of the capacitor into the nearby GND terminal on the breakout board but do not tighten that terminal yet.

Insert the end of the green-and-yellow wire also into the GND terminal on the breakout board. Now tighten that terminal.

Connect the other end of the resistor to the spare 'PC5V+' terminal on the breakout board.

Tug all wires and components in the breakout board terminals to check good connection.

Connect the green-and-yellow wire to the spindle body. At SheetCNC.co.uk we achieve this by connecting the wire to the shield of the cable which goes from the VFD to the spindle. Take great care to keep the wire clear of the high voltages present on the mains power terminals and the motor supply terminals of the VFD. If in any doubt, engage the services of a qualified electrician.



This completes the construction and programming of the probe.

Principle

The resistor ensures that pin 13 is normally held at a +5V, which equates to the probe being 'off'. When the red and green wires are connected (via the plate, the cutter, the spindle body, and the shield of the power supply cable) then pin 13 is pulled to ground, which equates to the probe being 'on'. The capacitor provides debouncing: when the cutter first touches the probe plate, it is likely to make-and-break contact several times in a few milliseconds and the capacitor smooths these multiple events into a single contact event. Without debouncing, LinuxCNC would detect the multiple contacts and raise an error. Tip: You will find reference online to software debouncing (within LinuxCNC's HAL). This is an entirely different approach to debouncing but requires extensive understanding of LinuxCNC.

Operating The Probe

Turn on all systems, including the VFD, but do not spin-up the spindle itself. Run LinuxCNC and home the machine as normal.

Warning: If for any reason the probe circuit is broken, then the probe will never be triggered. The result will be that the cutter is driven down through the top surface of the workpiece or the probe plate without stopping. Workpiece damage and/or cutter damage are likely. Small cutters, such as those used for engraving, may even shatter. Wear eye protection whenever probing.

Note also that the probe can be triggered only when 24V power is supplied to the breakout board. The 24V power is supplied by the VFD (not by the PSU), so it is essential that the VFD is powered-up when probing. The spindle does *not* need to be turning - indeed if it were turning the probe would likely malfunction. If using the VFD suggested in *Spindle, Control and Power Systems*, then its display should be showing "00000" flashing.

Mount a cutter in the chuck. For first testing, use an old cutter.

Jog the X and Y axes so that the cutter is above the workpiece. Place the probe plate on top of the workpiece, directly under the cutter. Jog the Z axis down to within 25mm of the probe plate.

In LinuxCNC, click the button labelled 'Touchoff To Plate'. The cutter should lower rapidly until it contacts the probe plate. Then it should rise a little and then descend again much more slowly until it contacts the probe plate again. Then it will rise to a safe parking position 1.5mm above the probe plate. The Z position in LinuxCNC should now display exactly the thickness of your probe plate, plus 1.500. The cutter is now touched-off at the exact height of the top of your workpiece.

The other button, 'Touchoff Direct' is used for conductive workpieces - for example the copper face of printed circuit boards. In such case follow almost the same procedure but instead place the probe plate onto the PCB and away to one side of the cutter. The probe plate must be in good electrical contact with the copper, so check the surface of the PCB for obstructions. Click the 'Touchoff Direct' button. After touching-off the cutter should finish 1.5mm above the surface of the PCB. The cutter will be touched-off at the exact height of the top surface of the copper.

In both cases, the maximum distance the cutter will be lowered is limited so that even if the cutter fails to electrically contact the probe plate the probing will eventually stop. If the distance limit is reached without the probe being triggered, the cutter will cease lowering and an error message will show. Clear the error message and investigate the problem before continuing.

Once touching-off is complete, remove the probe plate and continue as normal with the machining process. If the cutter needs changing during machining, re-touch-off as above. See also *Working with Multiple Cutters* in this document.

X and Y Probing

The code we provide already includes a routine for touching-off to the sides of the probe plate. This is useful for touching off an approximate location of the starting corner of the work, and is much safer than probing against the sides of a clamped workpiece because the probe plate will simply drag sideways if the cutter fails to make proper electrical contact.

The probing code assumes that a 3mm dia flat-ended cutter (or probe rod) is in the spindle. To use this code, first position the probe plate square on the bed and at the corner of the workpiece that is nearest to the home position. Touch-off to the top of the plate using the 'Touchoff to Plate' button. Then jog the cutter so that it is about 10mm nearer to the X-home than the nearest side of the probe plate, and about 10mm nearer to the Y-home than the nearest side of the probe plate. Then click the 'XY Touchoff APPROX' button. SheetCNC will automatically seek both sides of the probe plate and will touch-off X and Y appropriately. If successful it will then position itself above the corner of the probe plate.

Probeless XY Touching-Off

The classic way to touch off in X and Y without using a probe is to jog the machine until the cutter is directly above the corner of the workpiece, then to click 'Touch Off' for each of X and Y axes. This works fine, but can be time consuming. There is a much easier way, if no probe is available.

Begin by homing the machine as usual.

Jog the cutter to within a few centimetres of the desired touch-off position.

Lower the cutter to within a few millimetres of the top of the workpiece.

Use a ruler to measure how much further the cutter must travel in the X direction so as to reach the corner of the workpiece. Allow for the radius of the cutter.

In LinuxCNC, select 'X' from the Axis radio buttons and press the 'Touch Off' button just underneath.

In the pop-up box, enter the distance you have just measured, *as a negative value*. Click OK.

The X axis is now touched off correctly.

Repeat for the Y axis.

To confirm the touch-off point is correct, raise the cutter to a safe height then go to the 'MDI' tab in LinuxCNC and enter the command `G0X0Y0` (which means rapid move to x-zero y-zero.) The cutter should move to be above the intended touch-off point.

Machining Tricks


Here are just a few simple design tricks when working with timber products.

T-Nut Lap Joints

T-nuts are inserts that can be pressed into holes in material to provide a thread into which a machine screw can be inserted from the other side. The job is neater if the nuts are recessed cleanly into the surface.

To ensure a T-nut's threaded part will position centrally in its hole, always machine any countersink *larger* than the overall diameter of the T-nut flange. Additionally, machine the hole for the threaded part of the T-nut only just deep enough, then finish through the material with a hole the same diameter as the *actual** diameter of the screw to be used. This ensures that the joint has the minimum play.

In the other half of the joint, machine a through-hole equal to the *actual* diameter of the screw.



Blue: countersink (if any) for T-nut *flange*. Green: recess sufficient for *threaded* part. Orange: through-hole for fastener.

Screwed T-Joint

You'll see these in the construction of SheetCNC itself, where parts meet at right-angles.

Machine a 9.9mm through-hole as an interference fit for a M6 barrel-nut. Locate the hole sufficiently far from the edge of the material so that the tip of the screw, when fully tightened in the joint, will reach *exactly* to the furthest side of the hole.

For strength, the hole should be at least 15mm clear of the edge of the material. For minimum sideways play in the joint the hole should otherwise be as close to the edge as possible - so select the screw length to match the materials being jointed. Tip: Screw sizes are usually in 5mm increments. Barrels come in various lengths. eBay seller MassiveAttack is a good source.

Machine a slot from the edge of the material *almost* to the edge of the hole. This accommodates the


* The fastener is likely to have actual diameter smaller than nominal.

screw. The slot should end 3mm clear of the hole. Preferred slot width is 5.5mm but 6mm will work - the joint will just be a little less precise. Depth of slot should be half the actual thickness of the material plus half the actual diameter of the fastener. Tip: MDF is frequently 0.5mm thicker than spec.

In the other half of the joint, machine a through-hole equal to the actual diameter of the fastener.


To avoid extraneous cuts, which would weaken the joint, ensure all radii to be cut are at least as large as the cutter radius. Turn off overcutting of corners for the layer that contains the slots. Or, if the cutter is the correct width (6mm), machine a 'no offset' contour along the centreline of the slot, rather than cutting an 'inside offset' around the perimeter of the slot.

After machining, use a drill the same diameter as the fastener (6mm) to break through the bottom of the slot into the barrel-nut hole.



Orange: through. Blue: partial-depth. Grey: represents screw-thread in barrel. Black: material.

Combining screwed T-joints with mortises and tenons gives extremely accurate positioning and hugely increases the strength of the joint. The screws can optionally be hidden in the tenons, if the material to be jointed is thick enough.



The tenon outlines must be overcut (or additional material must be removed from the mortise edges) to allow the surfaces to mate.

Quick and Dirty Fastener

uPVC soffit board will readily take threads. Simply cut an undersize hole (say 7mm) then screw a machine screw (M8) into it. The hold will be remarkably strong, although repeatedly removing and replacing the screw quickly wears the thread. As an additional benefit the screw will be placed *precisely* where the drawing dictates, with no lateral play or error.


Cable-Tie T-Joint

Useful when joining thin materials at right-angles where a glued joint would normally be used. Especially useful when designing a new product because it saves waiting for glues to set and allows a subset of parts to be swapped out during the prototyping phase.

In one part machine a row of 4mm holes (or slots) along one side of the joint and a row of matching holes on the other side of the joint. Make mortises in the clear spaces between the holes.

In the other part, drill a matching row of holes along the edge to be jointed. Make tenons along the edge.

Bring the parts together and tie-wrap through the holes to hold the joint.



Orange: through hole. Blue: pocket mortise for tenon, usually through-cut and with overcut corners. Black: Material

Optionally, recess the tie-wraps into the slotted part to semi-conceal them. Copper wire can be substituted for the tie-wraps to create a permanent stitch-and-glue assembly in the production version.


Indexing and Orienting Parts

If two parts must lap-joint in a precise position relative to each other, machine matching 9.9mm diameter pockets into each part to accept a M6 barrel nut, with interference fit. This indexes the parts precisely. Use a pair of such barrel nuts to align both ends of the parts if necessary.

If two parts are to be assembled by a third party, and it is unclear which way round they fit together, machine matching 9.9mm diameter pockets into each part at an offset location. When a barrel nut is inserted in the pocket of one part, the other part will then only fit one way round.

Adding Bearings

Deep-groove ball races can be recessed into the surface and nipped in place with flange-head screws (you'll see these fitted to SheetCNC). Bearings can also be sandwiched between two parts by cutting half-depth recesses in each part. MDF is denser near the surface than in the body of the material, so sandwiching bearings is the stronger option and it usefully indexes the two parts as well. Either way, a further recess is required around the inside to clear the inner moving part of the race. Allow at least 2mm extra depth for this recess.



Blue: recess for race. Green: further recessed to clear inner race. Orange: shaft clearance

Precisely Matching Mating Parts


Every CNC machine suffers cutter deflections, to a greater or lesser extent. These deflections may be slightly different along the X axis than along the Y axis. This can mean that two parts that are otherwise identical, but machined at different angles on the machine, may differ very slightly. So, if two parts must mate with each other, it is good practice to lay out those parts on the machine bed so that the critical matching dimensions on each part are cut along the same axis of the machine.

Gears

Check out web page <http://hessmer.org/gears/InvoluteSpurGearBuilder.html> . Dr Hessmer has kindly provided an online gear drawing program. It will - for free - create a DXF file containing a pair of gears of any size you specify. Racks and internal gears, too.

This webpage <http://lcamtuf.coredump.cx/gcnc/ch6/> accurately describes how involute gears work and the design choices that must be made when specifying gears.

Wacky gears can be created by importing gears of different sizes into CAD then cutting and joining them. The slotted gear shown below will swivel then slide then swivel etc when driven by a sprocket (or by another slotted gear). The gears would be fastened to the backplate shown by a single pivot point each. uPVC soffit board is ideal material for experimenting with gear designs. Tip: It can be tricky to peel off the protective film from a complex shape, so it is best to remove it before machining.



For scale, the sprocket is 40mm diameter.

This pattern was cut using a 1/8" single-flute HSS cutter in a single pass at a feedrate of 6000mm/minute.

Other strange gears available from the web include a pair of square cogwheels which, when correctly meshed, will drive each other without slipping.

See also *Overcutting*, below, for some tips on increasing clearance to maximise the strength of the cogs.

Arcs

An arc is a curve that can be completely described by a finite set of data. Circles, half-circles, ellipses, and bezier splines are examples of arcs.

Genuine arcs are reproduced accurately by SheetCNC on the machined component as pleasingly smooth curves. However, some care is needed to ensure that any curves are faithfully passed all the way through the CAD/CAM process so that they can be correctly machined, as arcs, on SheetCNC. The problem is the poor way that many CAD packages handle arcs.

In TurboCAD, as in most CAD packages, any arc is initially created as a single entity - a smooth curve defined mathematically. For example a circle is defined by its centre position and its radius. However, if the CAD user *modifies* an arc, TurboCAD will sometimes opt to replace the arc entity by a collection of short straight lines that only approximate the curve.

As the designer works on a drawing, true arcs are mixed with these approximated curves, seemingly almost at random. The more editing that is done to an object, the more likely it is that the object's curves will be approximations rather than arcs.

Even in the CAD stage, this can be a significant frustration. For example a corner may be chamfered using the chamfer tool to insert an arc where the two lines of the corner meet. And, for as long as that chamfer is represented by an arc, the chamfer radius can be changed - or removed - simply by reapplying the chamfer tool using a different value. But, once TurboCAD chooses to change the chamfer to a collection of line segments, this becomes much more difficult.


Worse, though, is the effect on the manufacturing process. If the approximated curves, with their many straight-line segments, are passed into a CAM package then they will be faithfully reproduced in the resulting G-Code as many short straight-line cuts. These in turn will be reproduced in the finished product causing a poor finish. Even worse, depending on the cutter diameter and the angle between the segments, LinuxCNC may decelerate into each joint between the segments and

accelerate afterwards, increasing the job time and worsening cutter life.

So, if you find that your curves are being machined with jagged edges, suspect that the CAD package has converted your arcs into straight-line approximations.

The 'Edit Tool' in TurboCAD can be used to determine whether an apparent arc in the drawing is, in fact, a series of lines: an arc will appear with a single node at each end and a green handle in the middle; whereas a series of lines will appear as a large number of nodes close together and with no green handles.

TurboCAD offers a solution to these frustrating approximated curves, albeit something of an afterthought: the ConvertToArcPolyline tool. This tool will analyse any polyline* and convert any approximated curve into a true arc. The tool is not perfect and often it misses a few segments at each end of the curve. Sometimes it adds spurious arcs.



On the left: the original polyline made of arcs and lines. In the centre the effect of boolean addition (with the rectangle) is to convert all the arcs to multiple straight lines. On the right the imperfect effect of attempting to restore the arcs using ConvertToArcPolyline.

On the whole it is better to avoid the problem in the first place. The most common cause of arcs being converted to line segments is the use of 2D boolean addition and subtraction tools, and the Join Polyline tool. So do not use these tools on objects that contain arcs. If a polyline is made entirely of straight lines then apply any chamfering to the corners only after all addition, subtraction and joining have been completed. The best alternative method for combining the two objects above would be to use the chain polyline tool, which simply draws a new polyline around the outline of the combined (or subtracted) objects. This method is not as quick nor as easy as a boolean operation, but the preservation of arcs and the time saved subsequently makes the extra effort worthwhile.

Just One Tab

Tabs are a nuisance. They cause the machine to make otherwise unnecessary Z movements and reversals along the cut path, slowing the job, increasing wear on the cutter, and creating extra machining marks on the part. There is a risk of damage to delicate parts when breaking-out the tabs, and tab remnants have to be cleaned off of the part after machining.

Minimising the number of tabs is, therefore, generally beneficial.

* A polyline is any end-to-end linked set of lines and arcs. It behaves as a single entity - so that the polyline can only be selected and manipulated as a whole. A polyline may be open (i.e. with two ends) or closed (in which case it typically defines the outline of a shape or hole to be machined.) A rectangle is an example of a closed polyline.

The main purpose of tabs is to hold a part in place while its final outline is being cut out. But, until the cutter reaches the last part of the final pass of cutting-out, the part is still being held by the as-yet uncut material. The time when the part is most at risk is when the cutter breaks through the final part of the final pass - only then is the part free to lift and cause havoc.

With these considerations in mind, it is often possible to get away with just a single tab. Simply place the single tab at a strong point on the part, ensure that there are no operations on the part *after* the cutting-out has been completed (which is good practice anyway), and set the cut start/end point of each pass to be at the same place as the tab.

One warning: if using a single tab, there is still a risk that the part will tilt so that an un-tabbed side of the part will protrude above the workpiece and foul the cutter during a later rapid move. So single tabbing is appropriate only when no other parts will be cut out subsequently. It's also important to ensure that the single tab is positioned so that there is a direct line from the tab to the parking position of the cutter without crossing the perimeter of the newly cut part. The cutter should then make the final rapid directly away from the part without risk of fouling any protruding edge. Otherwise it may be necessary to use extra tabs just so that the part does not lift anywhere around its outline. Alternatively- if using SheetCAM - it is possible to turn off the X and Y parking moves. (See SheetCAM's Job Options window.)

Whether using one tab or multiple tabs, it is always worth trying to design the tabs so that their height is entirely within the final pass of the cutting-out operation. That way each tab only has to be machined once. Some juggling of cutting parameters - particularly depth-per-pass and tab height - will ensure that the final pass has enough depth to accommodate the tabs.

Or Even Zero Tabs

Tabs can be completely eliminated. Simply place at least one (preferably two) fixing points within the boundary of the part to be cut out. The fixings can be placed during the initial process of fastening the workpiece to the bed. Alternatively, if the part has through-holes as a feature of its design, the fixings can be placed after those through-holes have been machined but before the outline is cut. Beware that the fixings must be recessed, or the clearance height of all operations which follow the placement of the fixings must be increased to allow for the height of the fasteners. Typically, the fasteners will be woodscrews with penny washers placed under the heads to avoid damage to the workpiece.

If using existing features of the part, and to allow safe placement of the fasteners mid-way through the job, split the machining operations across two separate G-code files. See *Working With Multiple Cutters*, above. There is no need to re-touch off after fitting the fasteners, unless the cutter is also being changed.

Zero Tabs in MDF

When cutting parts from MDF, the use of tabs may be redundant. Depending on the cutter, the waste dust and chippings may pack the cut so tightly that the part is held in place even without using tabs. Experiment *with great care* to determine whether your combination of material and cutter would allow you to do this. Tabs will almost certainly be required on any part which has overall dimensions of less than about 100mm x 100mm. Here at SheetCNC we rarely use tabs on MDF parts at all. Work with caution and always wear PPE appropriate to the risk of having unsecured parts on the machine table.

Perfect Perforations

If you are using a new and unfamiliar material or cutter, or if you have altered the feeds and speeds, then the machine performance will change very slightly. This means that features which were before cut perfectly using your previous settings may come out slightly undersize or slightly oversize.

In any case, changing the cutting parameters will certainly change the quality of the finish.

So after any parameters have been changed it is best to test-cut the critical components of a design on a small piece of scrap before committing to cutting the workpiece itself. For example, if you have holes, slots, tenons or other dimensions which must be perfect, create a test-program which cuts copies of those holes and shapes into a 100mm x 100mm test-piece.

Then measure the actual cuts and adjust SheetCAM's 'finish allowance' setting appropriately.

The *position* of a feature will not be affected by changing materials or settings. But the finish and/or size of a feature may subtly differ from expectation. Hole diameters, in particular, should be tested and checked.

Overcutting

Overcutting is a process which compensates, at internal corners, for the fact that the cutter has a finite radius. Overcutting is usually employed to permit a square tenon to fit inside a square mortise.

In SheetCAM - as in other 2.5D CAM programs - overcutting can be included in a cut-path automatically.

As discussed in *Spindle, Control and Power Systems*, automatic overcutting may cause issues with many homebuilt CNC machines: the machine must bring the cutter to a complete halt as it cuts into the corner, resulting in very sharp deceleration and acceleration. The halt can trigger a resonance which results in the cutter subsequently deviating from the cut-path in a decaying sinusoidal movement. The deviations are small and brief, but may be a cosmetic issue when working with fine detail. When cutting very small cogs the deviations may even be sufficient to spoil the result significantly.

The issue can be overcome in a number of ways. The pros and cons are discussed below.

Reduce Machine Acceleration Limit

It may seem logical to simply configure the machine with a reduced maximum acceleration on the X and Y axes. This would be done in the machine controller's configuration files. On a LinuxCNC controller, for example, the .INI file specifies the maximum acceleration. Note that maximum acceleration cannot, unfortunately, be specified via the CAM process.

Reduced maximum acceleration exacerbates issues with chip-load, heating, cutter wear, and job duration. It is not a perfect solution.

Reduce Cutting Speed

If the cutting speed is reduced, then the duration of each deceleration/acceleration event is also

reduced, so the resonances are also reduced. However, limiting the cutting speed has even more drawbacks than limiting the maximum acceleration. And on small objects with complex outlines, for example cogs, changing the cutting speed may even have no effect at all because the machine may not ever, in fact, reach the specified cutting speed. So reducing cutting speed is very much a last resort.

Avoid Overcutting

In most cases, overcutting can be replaced with alternate strategies.

When cutting cogs, overcutting would sometimes be specified for the roots of the cog-teeth. However, overcutting undesirably weakens the cogs and should be avoided. Instead, the clearance between meshing cogs should be increased by a few tenths of a millimetre, simply by cutting the troughs of the teeth very slightly deeper. This has no effect on the precision of the gears and does not increase the backlash of the finished geartrain. The Hessmer cog-generation program (see *Gears*) allows additional clearance to be specified. This technique is highly recommended whenever cutting gearwheels. An increase in clearance of between one quarter and one half of the cutter diameter is usually sufficient.

In a mortise-and-tenon joint, the mortise can be cut overall a few tenths of a millimetre oversize and/or chamfered. If the material is soft and the radius is small, the interference fit at the corners of the mortise can be incorporated as part of the design - potentially even improving the rigidity and precision of the assembled joint. This technique is occasionally useful, but not generally recommended.

Overcut by a Separate Operation


If a precise, overcut, square corner is absolutely necessary to the design, then consider machining the corner in two operations: a non-overcut hole-cutting operation, followed by a diagonal move into the corner to make the overcut. This method disconnects the overcut resonances from the hole-cutting process. It also permits the user to specify the depth of the overcut, rather than having to accept the overcut provided by the CAM program. This technique is highly recommended.

The overcut can alternatively be made by plunging the cutter at the corner (only *after* the main cut has been made) or by drilling the corners (*before* the main cut is made).

Overcutting by prior drilling has the advantage that a small diameter drill can be used to make the overcut, so reducing the amount of the overcut. For example a corner made using a 6mm cutter can successfully be overcut by drilling 4.5mm dia at the corners beforehand.

Printed Circuit Boards


Right at the beginning of *Build Manual* we showed a magnified image of a tiny PCB created on SheetCNC. You probably wondered if that was a genuine and practical piece of work, or just a gimmick. Here it is again...



The tracks are 0.6mm wide. Pads are octagons on a 0.1" spacing. The isolations were cut in multiple passes using a 0.2mm 10° bit. The 1.9mm text (top and bottom) was also scribed with the same bit. One tool-change was made, for drilling the holes with a 0.8mm drill.

This particular PCB mounts a simple 555 circuit to make a pin-badge with a flashing LED. The PCB was created on a bog-standard SheetCNC machine which ten minutes earlier had been cutting parts from a full 8'x4' sheet of 25mm MDF. A dozen boards were manufactured - with 100% success - from a single 100mmx150mm workpiece. So, this is not a gimmick nor were any special modifications made to the machine for the purpose. It just works.

The PCB was designed on freeware Eagle v7.5.0 Light. The G-Code was created using John Johnson's excellent PCB-Gcode v3.6.2.4 plug-in for Eagle, which also is freeware. Eagle has a peculiar user interface which is slightly less intuitive than quantum mechanics, making it extremely hard to get started with. On the other hand Eagle is by far the most popular PCB freeware meaning that there is a wealth of help on-line.



Here are some quick tips that you'll need to get the most out of PCB-Gcode and SheetCNC when cutting PCBs:

- Install the Eagle and PCB-GCODE configuration files which we provide in our download. These will work for almost all through-hole components, including DIL integrated circuits.
- Check that SheetCNC's bearings are correctly tensioned and are free-running.

- Do not worry about SheetCNC being chain, rather than leadscrew, driven. Historically - and even currently on some low-end machines - chain-drive CNC employs bicycle chains with consequent high chordal errors and backlash, resulting in a poor reputation for chain mechanisms. But SheetCNC uses 04B micro-chain which offers considerably higher precision*.
- Purchase some 0.2 or 0.3mm 10° carbide engraving bits. These are about £5 for 10. The 1/8" (3.175mm) shank varieties are commonplace and will fit your 4mm-3mm collet. With careful choice of drilling depth, these bits can be used for etching, drilling, and final cutting-out of the board without requiring a tool change. They are delicate so buy spares.
- Optionally, purchase some 0.8mm or 1mm drills or end-mills, for hole drilling. Again, these are cheaply available.
- Use only FR2 grade PCB. The FR4 grade contains glass and wears out budget bits almost instantly, resulting in a terrible finish.
- A good, workable DRC setup for Eagle is: clearances all 0.8mm except pad-pad 0.55mm*; track width 0.6mm-1mm; minimum drill 0.2mm; restring all 0.6mm; shapes all octagonal.
- Declare the etching tip width to be 0.4mm in PCB-Gcode - this allows that the tip is wider as it goes deeper. Specify just 0.25mm step size - this is the dimension which really matters. Isolation maximum is a matter of taste and depends on how steady the soldering hand is. We generally use 1mm but, when cutting with a conical tool, don't expect the actual results to match exactly. (Also note there is a known bug in PCB-Gcode which causes it to wrongly offset the cutting paths by half the CED.)
- Set slow feeds with high spindle speeds - around 200-400mm/min at 12,000rpm for etching. All preceding advice about chipload is almost irrelevant for these tiny cutters. Reduce feedrate to just 100mm/min when cutting-out, to lessen tool bending forces. The tips of these tiny tools snap easily, otherwise.
- Use a rapid clearance height ('Z Up') of at least 2.5mm and set Z feed rates of no more than 200mm/minute.

* SheetCNC's chordal variation (the difference between commanded position and actual position due to the chain being made from solid links) is +/-0.13mm. In other words, some items on the PCB will be placed out-of-position by as much as 0.13mm. This is not an issue and is better than a typical chemically etched PCB. Furthermore, the out-of-position is *repeatable*: meaning that every time a feature is positioned at a given coordinate it will always be placed at the same physical location - not with a random error. Much more worrying would be the potential variance in size of individual features. Fortunately this variance is much, much smaller. For example, the worst possible chordal variation across a 0.6mm track is just 0.03mm which is a negligible error. Not only that, but any linear features that are commanded to be all in a perpendicular line are guaranteed to be actually placed in a straight line. So the holes for DIL packages will be drilled properly in-line, with no chordal-induced side-to-side variation, which ensures component fitting won't be compromised.

Backlash in a well set-up chain system is virtually zero. The chain's rollers register with the sprocket very precisely whichever direction the mechanism is driven from. Provided that the grooved bearings are correctly tensioned, backlash will be negligible. For convincing proof, examine the perfectly parallel and evenly spaced 0.25mm step-over cuts at the bottom of the etched isolations in the image of the PCB. Backlash error would show up as uneven spacing or lines crossing over each other.

In summary: chain-induced chordal errors have no significant effect on SheetCNC, even when cutting something as delicate and precise as a PCB. Backlash is virtually non-existent with a miniature chain drive, provided the bearings are smooth-running.

* Pad clearance settings are likely to be ignored by PCB-GCODE for DILs, but the result remains workable. Compensate by setting PCB-Gcode's isolation minimum to 0mm to force at least one tool-pass between adjacent pads.

- Set drill dwell time to 0.5s.
- If using LinuxCNC, add a 'G64 P0.01' command into PCB-Gcode's .pp file as part of the METRIC_MODE string. If you're working in inches, make the appropriate change to the INCH_MODE string instead. The G64 is essential to prevent LinuxCNC rounding-off the tiny corners of the octagonal pads.
- If both etching and drilling the PCB using the same 10° cutter, then PCB-Gcode's 'Spot-Drill Holes' option can be set to a significant depth - around 2.5mm - so that the cutter will sufficiently penetrate the board to create usable holes. (Normally this setting is just used to make shallow centre-marks ready for a follow-on drilling operation with a twist-drill.) Forcing the single cutter to do both operations is very convenient for prototyping and saves having to run the drill operation separately, but it would be poor practice in a production run because the resulting holes are tapered and much more copper is removed than necessary.
- To create a smooth platform for machining, fasten a 300mm square piece of 25mm MDF atop the spoilboard using a screw at each corner. Drive the screws well into the table, not just into the spoilboard. Make sure the MDF sits *absolutely* level and flat - you're working with cut depths of just 0.1mm to 0.3mm. Either tape the PCB workpiece to the top of the MDF using insulation tape around all edges of the PCB, or (better) machine a clamping frame from 3mm ply. Start by experimenting with small pieces of PCB - say 100mm square - which are more likely to lie flat and so give good results. Alternatively, make a simple vacuum table to hold the PCB - this allows quite large PCBs to be cut.
- Touch-off using a contact sensor: an effective manual electronic touch-off sensor can be created by simply connecting a resistance meter between the PCB and the spindle. The spindle must not be spinning and the meter must be set to megohms for this to work reliably. We find that a cheap digital multimeter can repeatably distinguish a 0.005mm movement on the Z axis, resulting in a perfect touch-off. It helps to solder a tag to the board, for reliable connection to the meter.
- If regularly creating PCBs, then an automated electronic touch-off probe will make life much easier. See *Probing* above.
- Engraving depth should be about 0.2mm. If the program fails to cut all the copper due to the board being warped, touch-off 0.05mm or 0.1mm lower and re-run.
- Double-sided boards, if they are symmetrical about one axis, can be cut with reasonable precision as follows. Touch-off and cut the bottom etch, bottom drilling and bottom milling files. Lift up the cut board, flip it over around its Y-axis and put it back into the hole it came from. Centralise it and tape it in place all round. Touch-off the Z-axis again. Read the X-width of the board from PCB-Gcode's milling preview, say Nmm wide. Re-touch-off the X axis at a location Nmm *closer* to the machine's home position. Then cut the top etch and top drilling files. Similarly, the top side of single-sided copper boards can be etched with component names and other annotations.
- The finished PCB can be cleaned-up by application of a kitchen scouring pad on the copper side.

Cutting PCBs on *any* home CNC machine is a big ask, but it is entirely possible so long as care is taken to level the workpiece and to touch-off accurately. But of all the obstacles that stand between a complete beginner and their perfect PCB, it is Eagle - rather than any practical consideration - that is the greatest stumbling block. Eagle is an extremely trying program to get to grips with, despite the many free tutorials and forums out there full of help. In the end the effort is worth it: the program is very powerful and, once its quirks are mastered, quick to use too.

We made our first PCB on SheetCNC merely to demonstrate that chain driven CNC is a practical proposition, even for miniature jobs. The technique was so immediately successful that SheetCNC became - overnight - our main means of prototyping circuits. We hardly use breadboarding any more. It is quicker and much more reliable to bang out a board in Eagle and cut it on SheetCNC than it is to fiddle with endless loose components and messy wires on a plugboard. And as for strip-board... that never even comes out of the drawer these days.

Aluminium and Brass

We'd like to be very clear about this: if cutting metal is one of your reasons for buying a CNC machine, we would suggest that you seek something designed for the task: a small machine with a coolant/lubricant supply to the cutting head and with a coolant collection bath around the cutting bed. Seek, too, a machine that has very, very high levels of rigidity, and the heavier the better. Have realistic expectations for what can be achieved when working metals on any home-built CNC machine.

Having said that, SheetCNC can *occasionally* machine aluminium and brass with *some* success provided that conservative cutting settings are employed to overcome the lack of coolant provision. However, cutting such hard material is very tough on the machine, so we strongly suggest that you just don't do it. In any case, you would find that accuracy is significantly worse than when machining wood.



A 50mm dia smiley dry-cut from 5mm aluminium sheet using 1/8" single-flute cutter.

Production Jobs and Repeat Jobs

If you are using SheetCNC to create products for sale, then you'll almost certainly be producing identical items repeatedly. Here are some tips to maximise productivity when machining multiple repeat jobs.

Touching Off Repeatably

The spoilboard will suffer much less wear if repeat jobs are always touched-off in exactly the same place for every job. The easiest way to do this is to position the workpiece at a known offset from the home position, and touch-off X/Y numerically rather than physically.

So, for example, the corner of the workpiece may be positioned at X=500 and Y=100 from the home position of the machine. To touch-off numerically to this corner, first home all the axes as described in *Spindle, Control and Power Systems*. Then, with all axes still at their home positions, select the 'X-Axis' radio-button and click the 'Touch Off' button just below. Enter the value '-500' and click 'OK'. Note the 'minus' sign. Then select the 'Y-Axis' radio-button and click the 'Touch-Off' button again. Enter the value '-100' and click 'OK'. X and Y are now touched-off to the chosen position.

To check that the touch-off position is correct, first ensure that the tool is raised to a safe clearance height, then select the 'MDI' tab and enter the command 'G0X0Y0'. The tool will then rapid-move to a position centred over the corner of the workpiece.

Touch-off the Z-axis in the normal way. A probe (see *Probing* above) is well worth the investment.

Positioning the Workpiece Repeatably

Production work benefits from having each workpiece clamped onto the machine table in exactly the same position for each run of a job. This saves operator time and minimises material wastage. It also reduces tool wear because the cutter does not cut the spoilboard repeatedly. And it saves having to clean up any burrs from the spoilboard after every run.

To achieve identical positioning of every workpiece in a production run, we like to mark the outline of the workpiece onto the spoilboard. For example, if our workpiece is rectangular we simply create a G-code file which engraves that rectangular outline 1mm deep into the spoilboard, and engraves a text legend which states the job name. Using G-code to do this, rather than simply drawing around the first workpiece by hand, means that we can accurately re-mark the position of the job onto a new spoilboard in the future.

Clamping the Workpiece Repeatedly in the Same Position

Clamping and re-clamping identical workpieces onto the machine using woodscrews that are repeatedly screwed into the same holes is not a practical long-term proposition. Soon, the holes will wear and the screws will fail to grip. Repeat jobs require a more sophisticated solution.

To overcome potential wear we instead insert M6 captive nuts into precision-drilled holes in the machine's table. Then we use M6 machine-screws to secure the workpieces.

To drill the holes for the T-nuts we create G-code to drill right through the machine table at the same locations as the fastener holes for the job itself. Then we mount a 7.5mm drill in the spindle and drill the holes through the spoilboard and through the table. We fit the T-nuts from below the table. There is, of course, one gotcha here to beware of: don't position any fasteners which would intersect with the fasteners that make up the X-table itself. See the tip below for an easy way to avoid doing this.

The T-nut technique works equally well whether the workpiece is being held to the table by screws that pass through the workpiece itself, or is being held by clamps that are fastened to the table outside of the workpiece.

Workflow

In summary, the workflow for a repeat production run of many identical jobs will be:

- Draw the job in CAD. Include the following items in the drawing:
 - The item to be cut.
 - The outline of the workpiece.
 - The locations of all the fasteners used to clamp the workpiece on the table.
- Create multiple G-code files in CAM, with the following separate files:
 - A procedure to engrave the outline of the workpiece onto the spoilboard, together with any identifying text. The text could, usefully, include the numerical touch-off position.
 - A procedure to drill the workpiece (6mm dia and down to a couple of millimetres into the spoilboard) for the fasteners. Not required if external clamps are being used to secure the workpiece.
 - A procedure to drill through the spoilboard and through the table (7mm dia) in the same locations as the fasteners which will hold the workpiece or the clamps.
 - A procedure to drill through the spoilboard (7mm dia) but *not* through the table at the location of the fasteners. This would be used to drill-out a future new spoilboard so that the T-nuts in the table are made accessible again.
 - A procedure to countersink the fastener heads into the workpiece. Optional.
 - And, of course, all the procedures necessary to cut the product from the workpiece.
- All the above procedures should have the same touch-off point which in most cases will be the corner of the workpiece itself.
- Record the numerical touch-off position for future use. For example, include the two numbers in the names of the files, or create a README.txt in that directory.
- Home the machine and touch-off numerically as described above. Check the touch-off position using the MDI command G0X0Y0.
- Engrave the spoilboard with the outline of the workpiece.
- Drill the spoilboard and the table, ready for the captive nuts.
- Fit the captive nuts.
- Place the workpiece on the table, drill it for fasteners, fasten it in place, and machine it. Repeat this step as many times as necessary.

For the next production run, which might be at any indeterminate time in the future, simply home the machine, touch-off numerically, drill the spoilboard (if necessary), then repeat the final step

above.

Tip: When prototyping or first machining a new product, it is best to use simple woodscrews to fasten the workpiece in place the first time. The holes made by the woodscrews can be checked to see if they conflict with any fasteners already in the table, before finalising the location of the production run on the table and committing to drill the table to accept the T-nuts.

Troubleshooting

Here are a couple of troubleshooting tips.

Why are deep holes sometimes tapered so that they're wider at the top than the bottom?

There are several possible faults which can cause this.

Deep holes must be cut in a spiral descent, usually 3mm to 6mm per turn. The top part of the hole is re-cut each time the cutter comes around the spiral. But the bottom part of the hole is cut exactly once. Inevitably the top part of the hole is finished and refinished so many times that it is slightly enlarged. This effect is, however, relatively small.

Other effects are more significant:

The top part of the hole is cut by the highest part of the cutter, nearest the chuck. This part will be sharper than the tip due to having had less wear. The highest part of the cutter will also be deflected less than the tip of the cutter because it is closer to the chuck so bends less. The deflection is increased when using a climb cut, too.

If the cutter has flutes that are too short, the chips will fail to clear and will instead jam between the shank of the cutter and the top part of the hole. This deflects the entire cutter sideways away from the cut, so that the lowest cutting passes at the bottom of the hole are cut around a smaller radius. The same effect occurs if chip load is too high for the depth of cut so that chips are not cleared all the way up the flutes to the surface. Sometimes, when a cutter is deep in the material, chips can be heard 'cracking' as they are forced between the cutter and the workpiece. This must be avoided.

Most of these effects can also be seen on vertical edges, too, resulting in a slight tapering of the edge.

Solutions: The most effective solution is to switch to a better cutter. For example, the carbide cutters we mention above. If that does not correct the issues then also explore the following: Try using conventional (not climb) cutting. Use a new, sharp cutter when cutting thick materials. Never cut deeper than the cutter's flutes can handle. Reduce chipload and depth-of-cut if necessary when cutting deep recesses. Keep the cutter short in the chuck, to minimise cutter flexing, although always leave at least 6mm clearance between the chuck and the workpiece. Use 1-flute or 2-flute cutters for better chip clearance. Use roughing and finishing operations, with the finishing operation having the minimum possible number of passes needed to achieve the necessary depth.

Why are my holes coming out undersize when they used to come out fine when I previously cut the same material with the same settings?

Chances are the cutter is blunted.

Other signs of a blunt cutter are the top (and sometimes the bottom) surfaces will have lots of burrs after cutting. If cutting plywood the whole cut face may be ragged and splintered.

Tapering of deep holes will also be much more significant with a blunt cutter.

Check, too, that the machine is correctly adjusted.

Drilled holes are mis-positioned relative to milled features. The holes are positioned up to 0.3mm too far along the X-axis.

This can occur if too much down-force is applied to the drill bit resulting in torsion of the gantry. To prevent this when drilling the holes, either increase spindle speed or reduce Z feed-rate.

Mis-positioning *will* occur if the drill bit is bent, or is too long and so flexible: the bit tip will 'walk' when it first touches the surface and will then penetrate at a random position. Subsequent drilling will follow the first penetration so that the finished hole is randomly mis-positioned by up to a couple of millimeters. The hole will also be off-vertical. So, never use a bent drill and do use stubby drills whenever possible, especially for sizes under 4mm dia.

Probing has stopped working. Now, the cutter crashes into the probe and just keeps going downward without stopping.

Check that the VFD is powered-up and is providing 24V to the 24V+/24V- inputs on the interface board. This input is required for the limit switch and probe sensors to work.

Check that the brass plate probe is electrically connected to the interface board, and that there is also a good electrical path from the cutter back to the board either via the spindle's earthing or via a separate wire. Check both with a multimeter - a visual check will not detect internally failed wires. The most common failure point is where the ground wire connects to the spindle - particularly if relying on a connection from the plug body to the spindle body. Ensure that the screw-ring which fastens the plug to the body is very firmly tightened.

My tools burn out and go blunt quickly.

If the tools are HSS, then perhaps you're cutting material which is too hard. Don't cut plywood with HSS. Most brands of MDF can be cut with HSS up to about 10-12000RPM but at higher speeds the heat generation becomes excessive. We strongly recommend using only carbide tooling.

Check the chippings which come off the tool. If they are more like dust than chips, then chances are the feed-rate is too low. Either reduce the spindle RPM or increase the feed-rate. Calculate the chip-load and see if it is reasonable.

Never use end-mills for drilling, except in thin plastics. Although most end-mills can theoretically make vertical plunge cuts and so be used to bore holes the same diameter as the cutter, it is in fact really, *really* hard on the tool. Chip clearance when boring is very poor, so heat builds up rapidly and destroys the cutter. It's tempting to use an end-mill to bore holes because it can save a tool change. But it's not worth it. Always use twist-drills for drilling operations, even if there is only a single hole in the job. When spiral-boring, a good rule of thumb is never to bore a hole which is less than 1.5x the cutter diameter.