Feeds and Speeds Tutorial for CNC

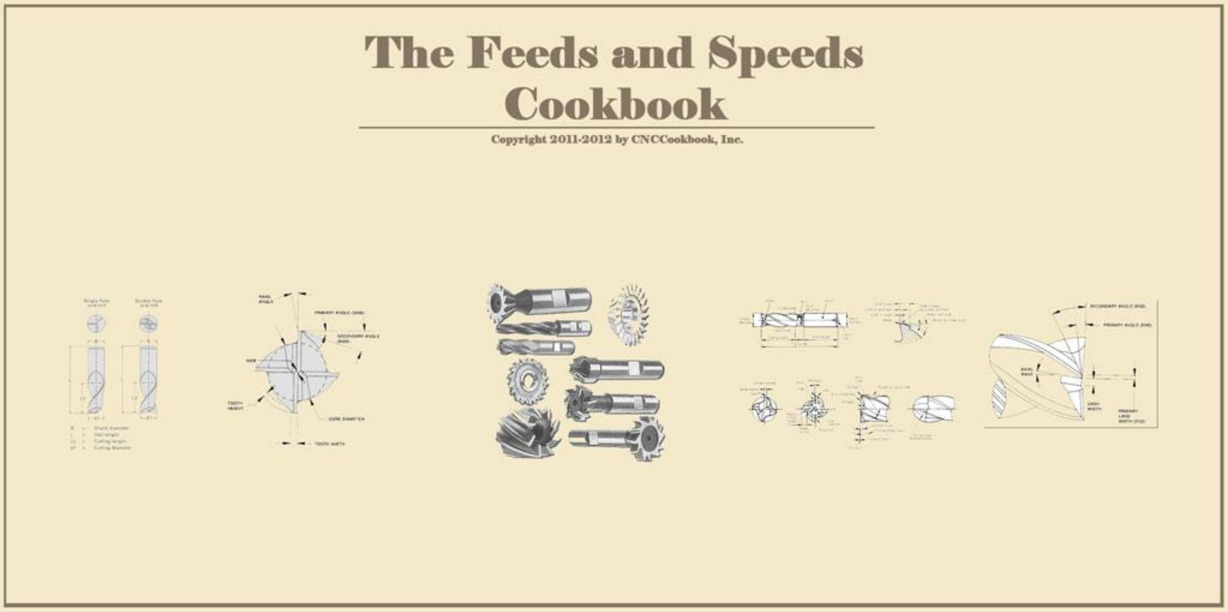

Getting the best possible results from your CNC machine covers a lot of ground. To begin with, there are several different goals we might be trying to optimize for:

- Best Material Removal Rates

- Maximizing Tool Life

- Good surface finish

You'll need to decide what your priorities are among these three, as different techniques emphasize each goal and you can't necessarily get great surface finish, maximum material removal rates, and the best tool life all at once.

In this series of articles, we'll go through some CNC Cookbook recipes that help you optimize for your goals, whatever they may be.

Table of Contents

Here are the articles available for maximizing your CNC machining productivity.

✓ signifies the most popular articles in each section if you want to skip ahead!

Basics Every Machinist Should Know

Basic Concepts for Beginners: You gotta start somewhere!

✓ Chiploads, Surface Speed, and Other Concepts: Key concepts and terminology.

✓ Calculating Feeds and Speeds: How do you calculate your feeds and speeds? (Hint: Use a state of the art Speeds and Feeds Calculator)

✓ Toolroom vs Manufacturing Feeds and Speeds: Do you know the difference between toolroom and manufacturing feeds and speeds?

Coolant and Chip Clearing: Best practices for coolant and chip clearing on the mill.

Intermediate

Turning Down the Heat in a Cut: Reducing heat prolongs tool life.

Dry Machining: Yes, you can machine without flood coolant. Often, it's even better.

✓ Tool Deflection Control: Critical to Your Success: Are you in control of tool deflection, or is it something that just happens?

Climb Milling vs Conventional Milling: There are times when Conventional is better!

Toolpath Considerations: How is your CAM program treating your tooling?

Ballnosed Tools and 3D Profiling

How To Choose a Stepover: Thoughts on how to optimize your stepover for best finish and shorter toolpaths.

Twist Drill Feeds and Speeds: Delving into the considerations for twist drill feeds and speeds.

V-Bits, Dovetails, and Other Exotic Cutters: Finding feeds and speeds for these less often used cutters.

Advanced

High Speed Machining: Tool Engagement Angle, Trochoidal Milling, Peeling, and all that Jazz.

Dealing With Chatter When Milling: Fixing Chatter via Rigidity or Tuning Speeds and Feeds.

Micro-Machining: How to avoid breaking those tiny little cutters.

Advantages and Pitfalls of Rigid Tapping

Long Reach Tools, Thin Walls, and Other Rigidity Challenges

Milling Tough Materials

Using a Cutting Knowledge Base to Beat Manufacturer's Recommended Feeds and Speeds

Quick Guides

10 Tips for CNC Router Aluminum Cutting Success: Take these shortcuts and skip a lot of pain.

Tips for Getting the Best CNC Milling Surface Finish: And the truth about mirror finishes.

Feeds and Speeds for Fly Cutters and Manual Mills: Tips and tricks for using G-Wizard with Manual Mills

Maximizing Material Removal Rates

Tips for Longer Tool Life

Recipes for Increasing Workholding Rigidity

Recipes for Toolholders and Tooling

CNC vs Manual Cutting Speeds: Why is CNC so much harder to figure out?

Videos

A Quick Video Course in Feeds and Speeds

Basic Concepts for Beginners

CNC Feeds and Speeds Cookbook

What are Feeds and Speeds and Why do they Matter?

There are some basic concepts and terms machinists use to discuss feeds and speeds that everyone should be familiar with.

"Speeds" refers to your spindle speed in rpm (revolutions per minute). In a series of experiments performed early on in machining, it was determined that your spindle speed is the biggest determiner of your tool's life. Running too fast generates excess heat (there are others ways to generate heat too), which softens the tool and ultimately allows the edge to dull. We'll talk more in our series about how to maximize tool life, but for now, consider your spindle speed to be largely about maximizing tool life.

"Feeds" refers to the feedrate, in some linear unit per minute (inches per minute or mm per minute depending on whether you're using the Metric or Imperial system). Feedrate is all about the tradeoff between maximizing your material removal rate. Material removal rate is how fast in cubic units your mill is making chips--the faster the better for most machinists, right up until it creates problems. The most common problem is tool breakage or chipping when you feed too quickly.

I'm a Beginner, How About if I Just Run the Machine Super Slow?

It's a common mis-conception that you can "baby" the cut in order to be ultra conservative. Just run the spindle speed way slow and the feedrate way slow and you won't break anything, right? Well, not exactly. Metal is a very unforgiving material. Plastics, wood, and other softer materials can also have problems from improper feeds and speeds, but metal is the most sensitive.

Here's some examples of what can happen if you run too slowly:

- If you reduce your spindle speed too much relative to the feedrate, you're forcing the flutes of your cutter to take of too much material. The endmill is being pushed too fast into the cut and the chips get too big. You can easily break a cutter this way.

- If you reduce your feedrate too much relative to spindle speed, you will soon cause your cutter flutes to start "rubbing" or "burnishing" the workpiece instead of shearing or cutting chips. Many machinists will tell you the fastest way to dull a cutter is just to run it with the spindle reversed and make a pass, but having too slow a feedrate creates a similar effect. We'll talk more about how this happens in the Feeds and Speeds article, but suffice it to say that running too slow is just as hard on your cutters as running them too fast, if not harder.

Okay, I Get It--There's a Sweet Spot for Feeds and Speeds

Yes! That's exactly right, there is a Sweet Spot for every cutting operation. It's not a point that has to be hit exactly, but at the same time, it is not very large either, and there are penalties if you miss it completely. The more difficult the material you're cutting, the smaller the sweet spot and the greater the penalties. Once you know where the Sweet Spot is, you can tweak your cutting parameters within that envelope to maximize Material Removal Rates, Surface Finish, or Tool Life. In fact, you can often maximize any two of the three, just not all three at once.

Let's take a look at the sweet spots for different things, as well as the danger zones:

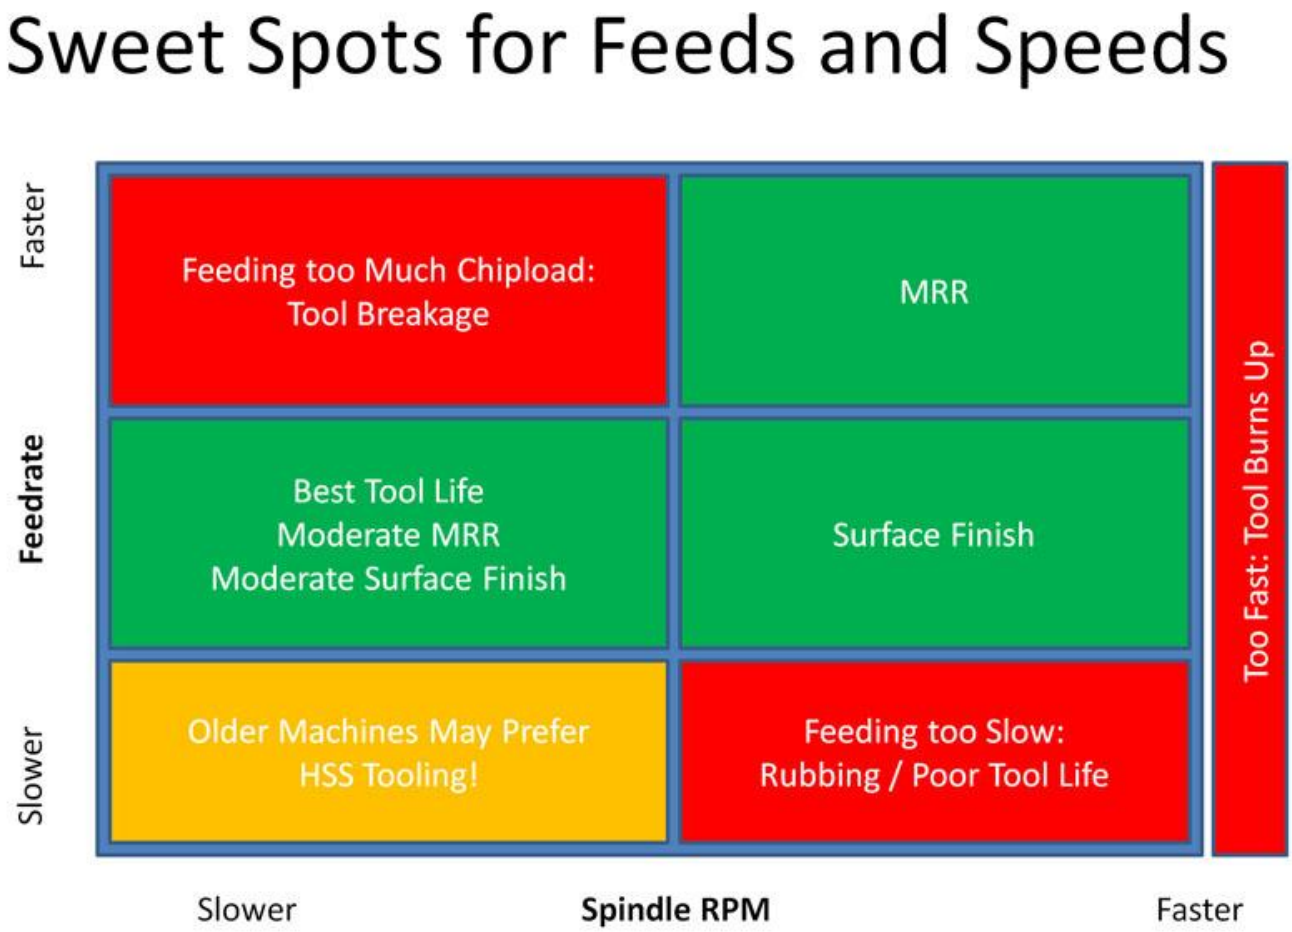

This chart is relative, meaning you can't assume anything about the proportions or scale. Just look at the positions of the regions relative to one another, and relative to the idea of faster and slower spindle speeds and feedrates.

Let's consider the different labeled zones, left to right, top to bottom:

Feeding too Much Chipload: As we've discussed, when you feed too fast for a given spindle rpm, you're likely to break the tool. The more you exceed the appropriate speed, the more likely. At some point, you'll always break the tool. Consider the absurd case where spindle rpm is zero and you rapid the tool into the work. Pop! Just broke another tool.

MRR: Running the spindle as fast as we can without burning the tool, and feeding as fast as we can without breaking the tool is the sweet spot for maximum material removal rates. If you're manufacturing, this is where you make money by getting further up and to the right than the competition.

Too Fast: Too much spindle speed will generate excess heat which softens the tool and dulls it faster. There are exceptions and mitigating circumstances we'll talk about in more advanced installments.

Best Tool Life: Slowing down the spindle a bit and feeding at slightly less than appropriate for maximum MRR gives the best tool life. We'll talk more below about Taylor's equations for tool life, but suffice it to say that reducing the spindle rpm is more important than reducing the feedrate, but both will help.

Surface Finish: Reducing your feedrates while keeping the spindle speed up lightens the chipload and leads to a nicer surface finish. There are limits, the biggest of which is that you'll eventually lighten the feedrate too much, your tools will start to rub, and tool life will go way down due to the excess heat generated by the rubbing.

Older Machines: So your spindle speed has come way down, and in addition, so has your feedrate. You're probably on an older machine where you can't run the kind of speeds you need to take advantage of carbide tooling. You may need to switch to HSS. It comes as a surprise to many that there are areas of the feeds and speeds envelope where HSS can outperform carbide, but it's true, depending on your machine's capabilities and the material you're cutting. Check the article "Is Carbide Always Faster than HSS" for more information.

Feeding Too Slow: As discussed, feeding too slow leads to rubbing instead of cutting, which can radically shorten your tool life and is to be avoided.

Now that you know how the sweet spots break down, you'll have a better idea how to steer your feeds and speeds to the desired results.

Chiploads, Surface Speed and Other Concepts

CNC Milling Feeds and Speeds Cookbook

Before we get into calculating the best feeds and speeds for your goals, there are a few more concepts we need to understand.

Chipload: Chip Thickness per Tooth

While feedrates are specified in length units per minute, the more important measurement is something called "Chipload". Thing of a chip as looking something like a comma in cross section, or perhaps an apostrophe. One starts big and gets smaller at the end. The other starts small and gets bigger at the end. We'll ignore that difference for a moment, though it is important as we shall see later.

Chipload is a measurement that is independent of spindle rpm, feedrate, or number of flutes that tells how hard the tool is working. That's a very useful thing, as you could imagine. Hence, manufacturers and machinists typically like to talk about chipload for a particular tool.

You can see that a tool with more flutes (cutting edges) has to be fed faster to maintain a particular chipload. Since each tooth is going to take a cut every rotation, a tooth has only a fraction of a rotation in which to cut a chip that reaches the chipload thickness. During the time it takes to rotate the next tooth to start cutting, the tool has to have moved far enough to shave off a chip that is thick enough. Hence, tools with more flutes can be fed faster. A 4 flute endmill can be feed twice as fast as a 2 flute, all other things being equal.

Why do Tools Break from Too Much Chipload?

You can imagine that forces simply become too great if a tool tries to take too much "bite" by having too much workload. This can chip or break the cutter.

But there is a second issue that comes from too much chipload--the chips get bigger and eventually can't get out of the cutter's way. Beginning machinists probably break more tools because they don't get the chips out of the way fast enough than because the force of the feed is breaking the tool. If the cutter is down in a deep slot, the chips have a particularly hard time getting out of the way. We use air blasts, mists, and flood coolant to try to clear the chips out of the way, but if they're way down a hole or slot, it makes it that much harder, and we have to reduce speed.

Making matters worse, chips always take more room once they're chips than the equivalent weight of material takes as a solid. The only place they have to go is gaps between the flutes of the cutter. Of course, the more flutes we have, the less space there is in the gaps. Can you see the point of diminishing returns coming? In fact, some materials have a tendency to throw off such big chips that we prefer 2 or 3 flute cutters instead of 4 flutes. Aluminum is a good example.

Performance Recipe: Cheating on the 2 Flute Rule for Aluminum and Going to More Flutes Elsewhere

Many beginners are taught to use a 2 flute in aluminum for chip clearance, but must we always use 2 (or perhaps 3 flutes) for Aluminum and never 4? Now that we know why fewer flutes must be used (chip clearance), we can think effectively about when we might not be restricted to fewer flutes. In general, when you have plenty of room for the chips to escape, you can use a 4 flute cutter, and you'll get a better surface finish.

How much is "plenty" of chip clearance?

Forget slots and plunging. Those are the worst cases. Try to avoid tight inside corners or interpolated holes whose diameter is at all close to the tool's diameter, those are nearly as tough. But what if we are profiling around the outside of a part and there's no concave curves, only convex? Tons of chip clearance there, so have at it. If you have a sufficiently roomy pocket, you may also get away with a 4 flute, especially if you can open up a big hole in the middle of the pocket to get started in.

The best case for more flutes is when you have a finishing pass, particularly if you're already committed to changing tools to get the best possible surface finish from a newer sharper tool that hasn't been roughing. The finishing pass will be very shallow and the rougher will have opened up plenty of room for chip clearance. Consider using 2 or 3 flutes for roughing followed by 4 flutes for finishing in materials like aluminum. In harder materials that don't need so much chip clearance, tools with as many as 10 flutes may be used.

This doesn't just apply to aluminum either. More exotic tools are available with 5, 6, 10 or more flutes. Experienced hands will tell you that if you're profiling (where there's lots of chip clearance) steel and aren't using 5 or 6 flutes, you're leaving money on the table. Let's run the numbers in G-Wizard. Suppose we're profiling some mild steel--1020 or some such. We're going to profile the outside of a part, so there's plenty of clearance. Cut depth will be 1/2", cut width 0.100", and we'll use a 1/2" TiAlN Endmill. Here are the numbers:

- 4 Flute: 3158 rpm, 29.8 IPM. MRR is 1.492 cubic inches/minute. A little over 1 HP.

- 5 Flute: Same rpms, now 37.3 IPM. MRR = 1.865. 1.3 HP. That's 30% faster cutting.

- 6 Flute: Now 44.8 IPM. MRR = 2.238. 1.6 HP. 60% faster than the 4 flute case.

How much more profitable are your jobs if you could run them 60% faster? The cost to do so is a more expensive endmill and a tool change for profiling. Harder materials can benefit particularly well because they're already up against surface speed limits. More flutes is the only way to get faster feeds.

Sometimes we have to go the other way too. If you've got some really sticky stainless steel that's giving you fits in tight chip clearances, try a 3 flute instead of a 4.

Rubbing: When You're Feeding too Slowly (aka I'm a Beginner, How About if I Just Run the Machine Super Slow?)

I've mentioned several times now that feeding too slowly causes "rubbing". How does that work?

Consider a magnified view of your cutting edge versus the material:

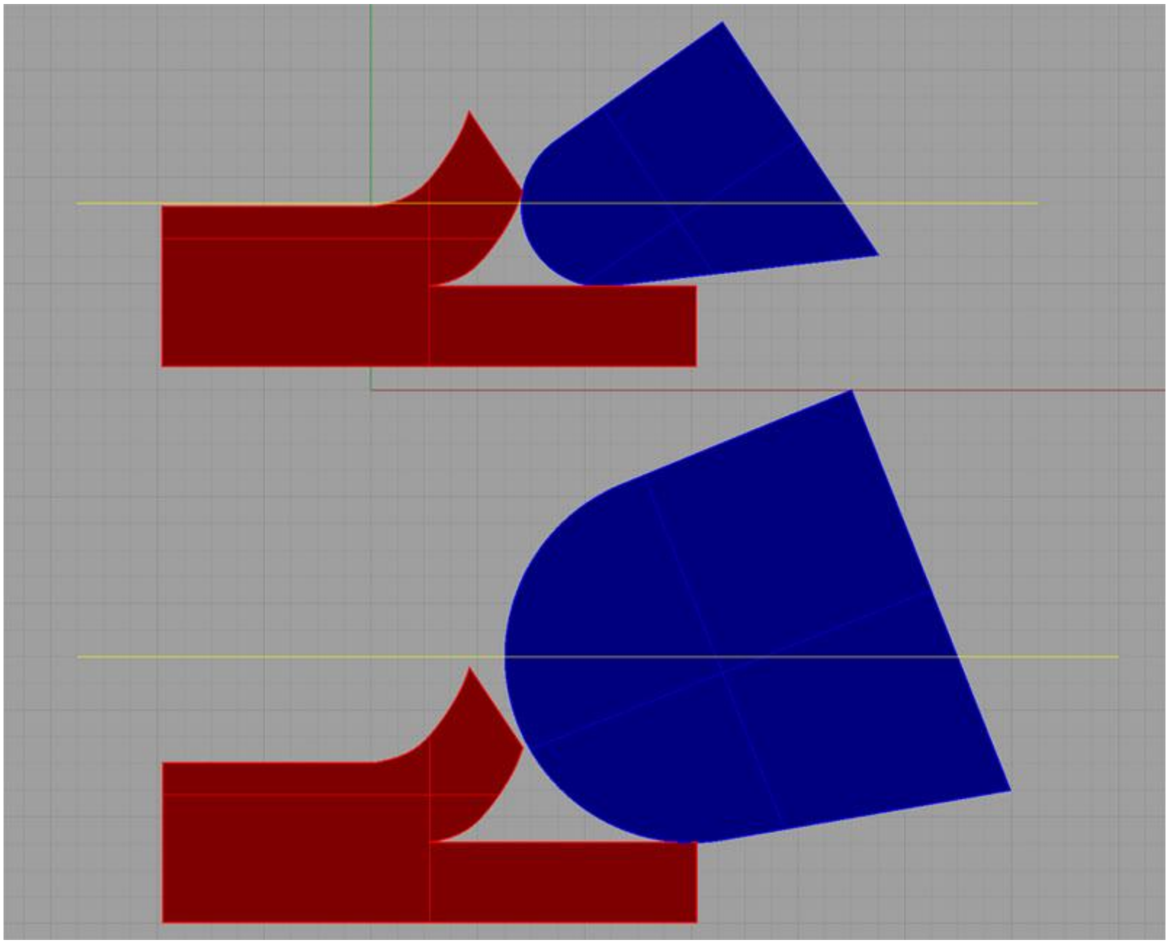

Two chiploads: Top one has chip thickness > tool edge radius. Bottom one has chip thickness < tool edge radius and will rub...

In the diagram, the cutter edge radius centerline travels along the yellow lines. If the radius is too large relative to the depth of cut (bottom), all the force goes to pushing the chip under the edge. This is the "rubbing" effect you'll hear talked about when feedrate and hence chipload are too low.

Tool manufacturers will tell you that too little feed is just as bad for tool life as too much feed (or too much spindle rpm). But how little is too little? That part is seemingly hard to find out. I went fishing around with Google to try to find what speeds and feeds result in a "burnishing" effect with tools. Here is what I found:

- Article on hard milling: 0.0008" per tooth is definitely burnishing because it is "less than the edge hone typically applied to the insert."

- De-Classified 1961 Batelle Institute report on aerospace machining of super-alloys says an IPR less than 0.0035 will result in burnishing and likely work hardening of these alloys. Interesting how well this number agrees with the one above for a 4 flute cutter. 8 tenths times four would be 32 tenths.

- Kennametal says the "highest possible feed per tooth will usually provide longer tool life. However, excessive feeds may overload the tool and cause the cutting edges to chip or break." So feed as fast as you can until you start to chip or break edges. They reiterate this under work hardening. One wonders whether rubbing leading to work hardening isn't the principal risk of cutting with too-low chiploads with respect to tool life in susceptible materials.

- Another reference, like the first, to keeping chiploads higher than tool edge radius. In this case, IPT should be greater than 0.001". This is once again an article on hard machining where work hardening may be a factor.

- Minimum chip thickness is 5-20% of the cutting edge radius. Below that level, chips will not form and the cutter will "plow" across the workpiece causing plastic deformation and considerable heat.

- Ingersoll says that as a general rule carbide chiploads should not be less than 0.004" or you run the risk of rubbing which reduces tool life and causes chatter.

I take away two things:

1. If you cut too little, you run the risk of work hardening if your material is susceptible to it. That will wreck your tool life if you are over-stimulating work hardening. Imagine tossing a handful of hardened chips into the path of our cutter--that can't be good!

2. Aside from work hardening, if you're cutting much less than the cutting edge radius, you're rubbing and not making clean chips. That will heat the tool and material and drastically reduce tool life.

Figuring out the work hardening part is easy. If your material is susceptible, keep the chipload up at manufacturer's recommendations and don't fool around. Figuring out the whole cutting radius issue is harder. Most of the time we don't know what the cutting radius is. I'm not talking about tip radius on a lathe tool, for example. I'm talking about the actual radius of the sharp edge. In other words, the smaller the radius, the sharper the tool. A lot of carbide inserts are pretty blunt. A chipload of less than 0.001" may very well be too little. Modern tools for aluminum are often much sharper, and can take less chipload. In general, indexable tools are usually less sharp than endmills, so they need higher chiploads.

It's ironic that just when you think you are taking it easy on a cutter with a very light cut, you may be doing the most damage of all due to rubbing.

Why chance it though? Use a calculator like G-Wizard to figure out how to deliver the manufacturer's recommended chipload by increasing the feedrates. Not only will the job go faster but your tooling will last longer.

Radial Chip Thinning (aka I'm an Expert and I ran the Machine too slowly without even knowing it)

Would you believe that especially for light cuts, the basic math combined with SFM and chipload tables often gives results that are wrong and radically increase the wear on your tools?

The reason is that there is more going on here than meets the eye. For example, if I poke around various endmill manufacturer's literature in search of speeds and feeds for steel, I can get to a page like this one from Niagara cutter. First thing to note is that the recommended chiploads and SFM vary depending on the exact operation being performed, and in particular, the depths of cuts. If you're just using the basic shop math around SFM and chipload, no such compensation is available. I have built compensation like this into my G-Wizard Machinist's Calculator, but trust me, it isn't so easy just to do it by hand. You'll be constantly referring to pages and tables, or to Excel spreadsheets.

Let's try an example based on doing some peripheral (edge) milling to profile a part made out of mild steel using a 1/2" uncoated HSS 4 flute endmill. We plan to take fairly shallow finishing quality cuts of 5% of the cutter diameter. Further, let's do a pretty deep cut axially, a full cutter diameter of 1/2". So if I am profiling a 1" high part, I can make a full pass by going around twice and cutting 1/2" each time.

What feeds and speeds should we use?

The Niagara page says for cuts less than 1/16 of a diameter (5% is 1/20), we want 210 SFM and a chipload of 0.0035. If I plug all that into G-Wizard, but ignore the chip thinning, I get the following results:

Radial Depth Ratio of 5% = 0.015" depth of cut

210 SFM and 0.0035 chipload gives us 22.46 IPM feedrate and a 1600 rpm spindle speed.

Is that the right speeds and feeds?

Yes and no. It's certainly what the majority of folks would use. In fact, they might even be less aggressive than that if they're trying to be conservative.

Let's see what G-Wizard would suggest by default and why:

For the same depth of cut and cutter, G-Wizard wants a little slower SFM of 160, and a little less aggressive chipload at 0.003.

The spindle speed works out to be 1200 rpm due to the lower SFM, but the feedrate is now 84.69 IPM--nearly 4x the original values.

How can we go so fast?!??

Unfortunately, avoiding the rubbing problem gets harder, even for experts, because of a phenomenon known as "Radial Chip Thinning." With chip thinning, you can be making a cut and following all the recommended chiploads, and still be rubbing. The cut above, 1600 rpm at 22.46 IPM will almost certainly wind up rubbing. The reason is that due to the geometry, when your radial engagement, the cut width looking down the axis of the tool, is less than half the diameter, the chips that come off are actually thinner than the basic formulas everyone learns in machinist's school predict.

A picture is worth a thousand words when understanding why:

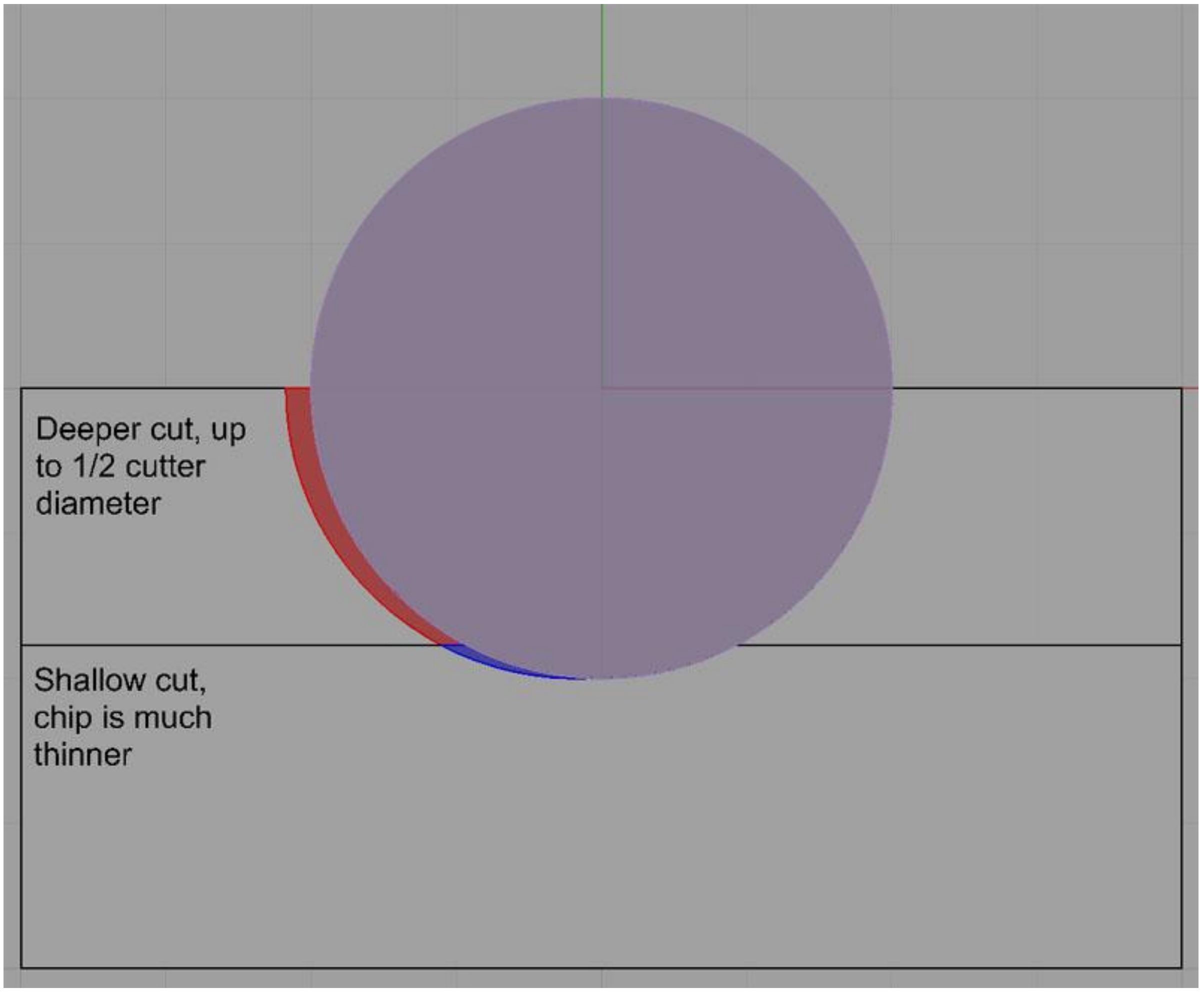

The blue chip is a shallower cut. Note how thin it is at its widest compared to the red chip from a deeper cut...

The blue chip represents a very shallow cut, and the red chip a deeper cut. Note how thin the blue chip is at its widest compared to the red chip from a deeper cut. You can see that the chip gets thicker all the way up to the point where we've buried the cutter to 1/2 its diameter. That's the thickest point.

Chip thinning simply answers the question, "How much faster do we have to go so the maximum width of the blue chip is the recommended chipload?"Because we're just trying to get back to the recommended chipload, chip thinning isn't about going faster in the sense of higher MRR's. You will see your machine going faster, sometimes a LOT faster, but chip thinning is about avoiding rubbing that will dramatically shorten your tool life.

Many manufacturers publish tables that suggest how much faster to feed based on the % of cutter diameter you are cutting. A good machinist's calculator, like G-Wizard, will factor in chip thinning automatically. Unless you never cut less than 1/2 the diameter of your tool, you need to make sure you're adjusting your cuts for chip thinning or you're probably wearing out tools prematurely as well as not taking full advantage of the material removal rates the tool is capable of.

Surface Speed: How Fast the Tool Slides Against the Workpiece While Cutting

When specifying the operation of a tool, surface speed goes hand in hand with chipload. Just as chipload is a better way to talk about feedrate because it is independent of so many factors, surface speed is a better way to talk about spindle rpm. Imagine that instead of a rotating cylinder with cutting edges, your tool was a flat piece of metal slid against the workpiece. The recommended speed to slide when cutting is the surface speed.

Surface speed is measured in linear units per minute: feet per minute (SFM) for Imperial, and meters per minute in Metric.

You can't really cheat on Surface Speed. It is what it is and exceeding the manufacturer's recommendations is sure to reduce tool life fairly drastically except for some very special HSM cases you should only worry about when you've mastered the beginner and intermediate speeds and feeds concepts.

Calculating Feeds and Speeds

CNC Milling Feeds and Speeds Cookbook

How Do Machinists Determine Feeds and Speeds?

By now you're getting the idea that feeds and speeds are important, and that they involve a lot of different concepts. We haven't covered nearly all of them yet, either. But before we go further, it's worth asking, "How do machinists determine Feeds and Speeds?"

There are a number of approaches:

- They can go by seat of pants and experience or by asking other machinists.

- They can use rules of thumb.

- They can punch it up on a calculator.

- They can rely on manufacturer's recommendations.

- They can use a machinist's calculator.

- They can rely on their CAM program.

- They can use trial and error and repeat what works.

Phew! That's a lot of choices. Which one is right? Do we have to follow them all?

Let's look at a few of them in more detail.

Seat of pants, experience, rules of thumb, and asking other machinists

Some machinists feel like they can judge feeds and speeds by sound and experience, with maybe a few rules of thumb thrown in. In the manual machining days, this was probably true, particularly when turning on a lathe. With your hand on the machine's handwheels, a machinist gets quite a bit of feedback about what's going on with the cut. If you have any doubt, try ramping into the cut by turning either X or Y at the same time as Z. The endmill cuts much easier than trying to plunge with your machine's quill. You really can feel the difference, even if it is hard to coordinate the motions as precisely as a CNC can.

Things like cutter engagement angle are largely irrelevant--you'd have to turn two handwheels at once in exactly the right way to mill around a corner. Feedrates were generally a lot slower, so radial chip thinning would seldom be a problem, at least not for longer than it takes to complete a light finish pass. Roughing would be contacted with the biggest hogging cuts that could be made--no dainty swirly HSM toolpaths were available. If you ran a Bridgeport manual mill by hand, never able to exceed 6000 rpm or so, and probably using HSS tooling, it's pretty easy to pick up the feeds and speeds, especially if an old-time master is looking over your shoulder and letting you know when you're screwing up.

Those days are gone, unfortunately. Today we have CNC, much higher speeds, the ability to do things manual machinists can only dream of, and a far more competitive business environment for manufacturers. You need to be able to maximize your machine's performance, or at least take advantage of as much as you can without damaging tools. CNC machines are dumb--they have no ability to sense much about what's going on in the cut. They're going to do just exactly whatever you program them to do, and the better ones will do it so fast it'll be all you can do to press the E-Stop before anything too terrible can happen. Tools can snap or rub and where out very suddenly.

Forget about seat of the pants or rules of thumb. You can ask other machinists who have experience, and many times they'll be helpful. But try to ask your question generally to get the maximum benefit of their experience. For example, they can tell you pretty quickly what range of surface speed and chipload they like to use and perhaps their preferred tooling. That's helpful, but those are just parameters and you'll still need to perform some calculations. If you ask for the answer, e.g. what spindle speed and feedrate for this exact cut, you're doing them a disservice (that's some trouble for them to sit down and figure out) and you're only getting value for one cut.

Experience counts, but experience knows it's important to get smart about what you're doing and not reinvent the wheel. Experience needs to move higher up the food chain than just feeds and speeds where purely mechanical calculations can produce accurate results quickly. Why waste your time on something like that when you can't do any better job?

Can I do those "Mechanical Calculations?"

Absolutely!

All the information is available. But, and this is important, there is a lot more going on than the simple formulas used to derive feedrate and spindle rpm can account for.

In the spirit of full disclosure, you can find the simple formulas in a lot of places, but I'll link to Wikipedia. These formulas accept as inputs surface speed and tool diameter to calculate spindle rpm, and they accept number of flutes, spindle rpm, and chipload to calculate feedrate.

Seems pretty simple, so where is the problem?

We've already seen one fly in the ointment in the form of radial chip thinning. Those formulas on Wikipedia don't account for chip thinning, so anytime you're cutting less than half the diameter of the cutter as your stepover or cut width, they're wrong. The thinner the cut, the more they're wrong, and ultimately they will be very wrong.

So, you'll need to go research the formulas for chip thinning so you can add them too. You'll also want to find a large table of materials, with chiploads and surface speeds. Ideally your table is large enough to be a materials database that considers not just broad classes of materials, but individual alloys as well as the *condition* of the alloy, and adjusts the figures accordingly. You will want to scale back your figures if you are slotting. In fact, you want to adjust based on how wide the cut is as well as how deep. There are manufacturer's tables out there to help you do that, it's just one more step to add to your process.

Speaking of steps, this stuff all adds up, and eventually, you have an awful lot of steps to be punching numbers into a calculator while rabidly flipping back and forth to look at various charts. I recommend using an Excel spreadsheet. In fact, that's how my G-Wizard feeds and speeds software started out, but I'll warn you, you will outgrow Excel if you keep adding bells and whistles like I did. More on that below.

What About Feeds and Speeds Calculators?

LOL, I thought you'd never ask (and I bet you figured I'd get here sooner or later because I sell software that calculates feeds and speeds).

Here's the thing, you can figure out everything you need to know to do what the software does and you can do it yourself. The data is all out there if you want to take the time to research it. To write G-Wizard, I've probably gone through several hundred learned papers by PhD's and countless thousands of pages elsewhere on the Internet. I have standing Google searches that give me alerts every morning if someone publishes a new article about speeds and feeds that might be of interest.

There are really only two reasons why you'd want to look into a feeds and speeds calculator like G-Wizard:

1. They work and produce better results than simpler methods. The software can consider a whole lot more variables than you can punch into your desk calculator.

2. Because you don't have time to do all the research and build the software that brings it all together.

Using a calculator is fast and easy. Take a look at my doc page on G-Wizard's feeds and speeds which includes a 2-part video course on feeds and speeds to get a quick overview of just how easy it is. I won't belabor the point further other than to say I can't understand why every machinist wouldn't want to use a calculator of some kind (whether or not you choose G-Wizard) unless they just weren't worried about being competitive or were so conservative about feeds and speeds (remember the manual machinist case) that it just didn't matter. After all, who wouldn't want the best possible material removal rates, surface finish, or tool life?

What's the Role of Manufacturer's Recommendations?

A number of machinists will pop up at this point and ask about Manufacturer's Recommendations. After all, doesn't the manufacturer know best how their tooling should be used? The short answer is, "Yes, but it's more complex than that."

Some machinists have the perspective that their manufacturer is making claims that are aggressive for marketing reasons. They're suggesting outlandishly high feedrates and surface speeds that the tooling can't actually back up or that won't work right when the machinist tries them. This is true in some cases, but most manufacturers can't afford to do this very much. After all, if the cutters don't perform, are you going to reorder?

What they can afford to do is shade things towards the aggressive. After all, who is to say whether the numbers are a tad aggressive and the tool wears out a little quicker than it has to? There are remedies for this. G-Wizard, for example, considers a lot of manufacturer's recommendations in an apples to apples match up (i.e. same coatings and geometries). It then does some very sophisticated number crunching to try to separate out the fact from the fiction. In other words, it tries to determine whether a manufacturer is overly aggressive (great MRR, lower tool life) or overly conservative (great tool life, lower MRR) to get to some "balanced" numbers. It does this by analyzing a minimum of 3 manufacturers for uncommonly used tools and 12-15 for commonly used tools (e.g. endmills or twist drills). It then provides a slider that lets you configure whether you're more interested in being conservative or aggressive:

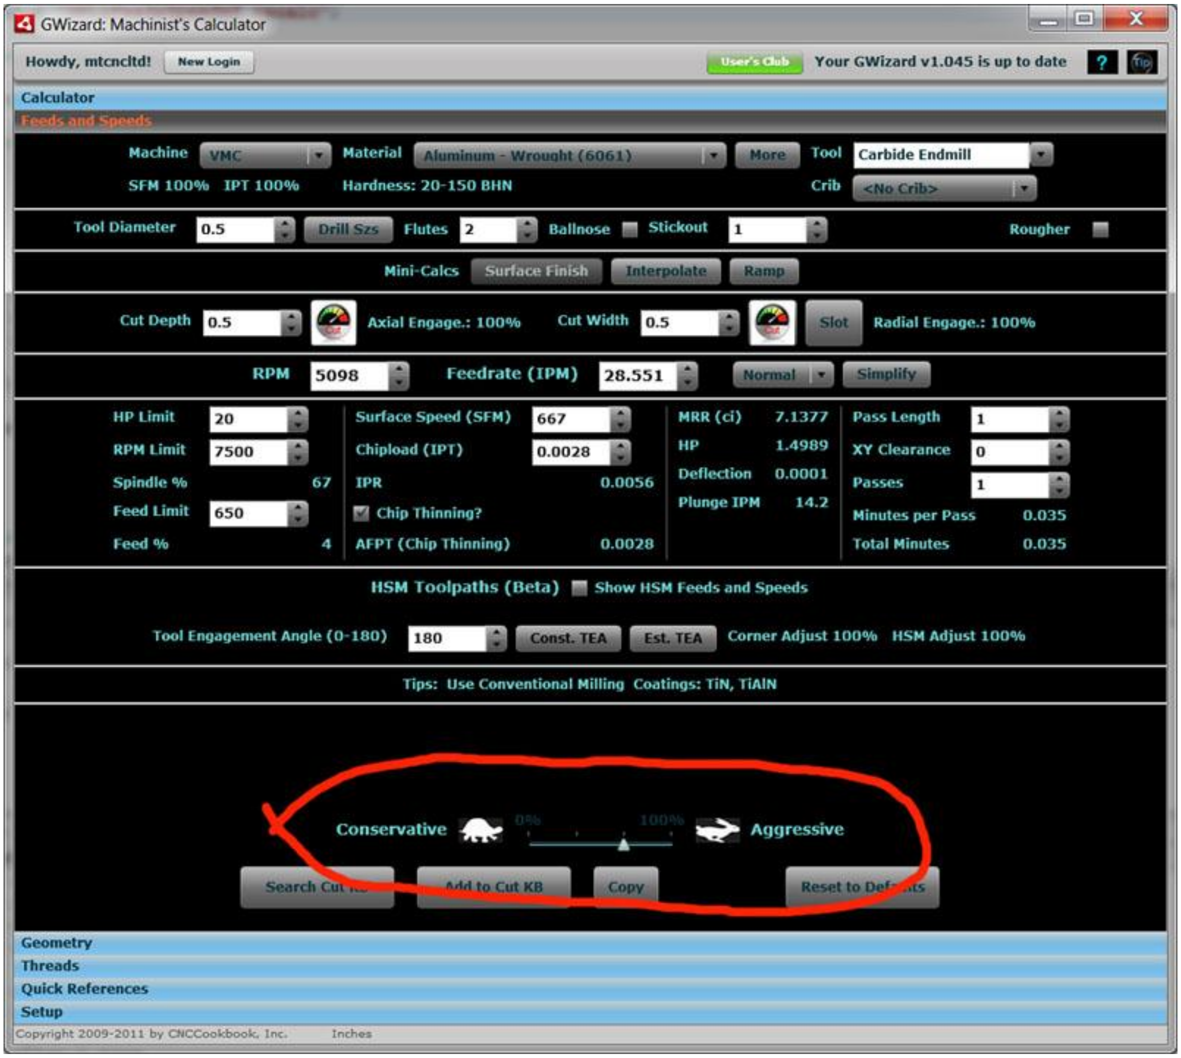

The G-Wizard Gas Pedal

We call this feature the "Gas Pedal", and it is depicted by a tortoise and a hare, much like the old Bridgeport manual mills had for speed control. I'll talk more about how to use the Gas Pedal and how to think about how aggressive you want to be in the article "Toolroom vs Manufacturing Feeds and Speeds", which is the next article after this one.

Being able to make your own choices about whether to be conservative or aggressive is useful, but here is the real way to think about calculators and other machinist's software:

It's all about how many variables you can master.

Sophisticated feeds and speeds software lets you master a lot more variables than you could manage by hand. You can see a number of them in the G-Wizard screen shot above, but there are even more built into the internals of the program. These variables all interact in various ways. Most are quantitative numbers and hard math, but G-Wizard even includes qualitative rules and variables. Note the line right above the Gas Pedal labeled "Tips". It says, "Use Conventional Milling," and, "Coatings: TiN, TiAlN". That's useful information to have handy. If you read many articles on machining, you'll know there are tons of these out there. I used to try to memorize them, but then I thought, "Why bother if I can have the software tell me the right rules at the right time?"

Every time you learn to master some additional variables, you can produce better results. G-Wizard is all about helping to master as many as possible.

To give an idea of how crazy it gets, G-Wizard considers 49 different variables as it is making a feed and speed calculation. Compare that to the half dozen considered by the Wikipedia formulas and you can start to understand the complexity behind modern feeds and speeds calculators. In addition to its 49 variables, it consults a total of 14 distinct databases. The total size of all that data makes G-Wizard the Calculator larger than G-Wizard the G-Code editor as I write this, even though the G-Code Editor is a far more complex piece of software. It's the sheer volume of the databases that makes the Calculator larger. And, it's being able to consider all that data together with all those variables and do the math in the blink of an eye that produces the results.

Let's go back to the Manufacturer's data one more time. Are we saying you should ignore it? No, absolutely not. On the other hand, G-Wizard and other calculators obviously can't incorporate every manufacturers data. Most of the time they don't even tell you which data was used to develop their database. If you use a particular line of tooling as most shops do, you'll want your calculator to be able to import and use the manufacturer's data. Ideally it will import and use it along with all the other rules and formulas built in. That last point is important: you need to apply all that math even if you have the manufacturer's data.

Why?

Because manufacturer's data has to be simplified in the interests of presentation. If the manufacturer gave you a fancy calculator like G-Wizard, they could afford to consider a ton of proprietary variables. But they don't. Instead, they give you tables. Tables limit the number of variables that can be considered. A two dimensional table considers just 2 variables, perhaps material and tool diameter, for example, to look up surface speed and chipload. If you're lucky, they give you a couple of extra tables and maybe some rules of thumb:

- "These numbers are good to 1/2 diameter cut depth."

- "Reduce SFM 50% for full slotting or when cutting more than 2 x diameter deep."

You've surely seen such rules. Once again, a calculator can consider far more complex models. It can interpolate smoothly from 0 to the 2x diameter depth, adjusting all along the way. It can consider any cut width when figuring radial chip thinning instead of just the few in the manufacturer's tables. This is valuable and leads to more performance no matter what you're trying to optimize for. The Manufacturer's data augments the 49 variables and 14 databases inside G-Wizard, it doesn't replace them.

So, enter your manufacturer's data into your calculator so it can add value to that data. G-Wizard lets you import the data as spreadsheet (CSV) files, to make it easy. It also includes a large catalog of downloadable manufacturer's data so you may not have to do any data entry at all. Lastly, if your calculator has tool table (tool crib) support and the ability to import manufacturer's data, they make ideal tools for comparing the performance of different tooling.

What About my CAM Program, Won't it Figure Feeds and Speeds?

Most CAM programs have some sort of simplified speeds and feeds calculator built right in. Unfortunately, most of them are painfully over simplified to the point where they don't do much more than your 4 function calculator would let you punch in with the basic Wikipedia formulas. As I write this, I have no less than 5 different CAM programs installed on my computer. They were all sent to me to evaluate and write about. Every one of them has cool features of various kinds that I love. But every one of them also has a very primitive notion of feeds and speeds. That's probably a good thing because it gives my G-Wizard business an opportunity to grow, but I wonder how many machinists just assume their CAM program is doing a good job for them on feeds and speeds?

You can tell how sophisticated a speeds and feeds calculator is by the information it takes in and the information it gives out. Take a look, for example, at the G-Wizard's Feeds and Speeds documentation page--there's a lot going on there. Now compare that information to what your CAM program is doing. Many of them have a lot of limitations. Here are just a few examples:

- A fixed chipload by tool without regard to material. This may be modified by some "chipload factor" by material, but that isn't how the manufacturer presents the data, so why should you stand on your head to think about it the way the crazy CAM program wants?

- No chip thinning calculations.

- Not much tooling-specific calculations.

- No qualitative rules, like when to use conventional vs climb milling (there are important distinctions there!).

- etc.

The short answer is using your CAM program is better than nothing, but not so great. For that reason, we're working on integrating G-Wizard with various CAM programs to make it easier for you. Meanwhile, it's easy enough to use G-Wizard and enter the values it produces into the CAM program. You'll be happy you did so as our users report it does a better job than even the market leading CAM programs.

Toolroom vs Manufacturing Feeds and Speeds

CNC Milling Feeds and Speeds Cookbook

I recently received an email from a G-Wizard user that prompted me to write this article. Here is what he said:

I ran a toolpath yesterday with Const. TEA and the endmill broke, I had followed manufacturers recommended SFM and chipload. Then I let G-Wizard do its magic. I noticed that the recomended SFM and IPT are lower than those recommended by the manufacturer, I guess those numbers are pretty universal, right?

People frequently ask where G-Wizard's numbers are pegged--are they aggressive or conservative? Closely related are questions about the Manufacturers numbers. This user found G-Wizard was more conservative than what the manufacturer quoted and it was a good thing. As I mentioned inthe article about calculating feeds and speeds, sometimes manufacturers are overly aggressive in the interests of marketing a product that seems a lot better than competitors. In fact, it may only be more aggressively rated than the competition while being no better a cutter.

For this article, I want to talk about something different than how aggressive your manufacturer's numbers may be, and that's this question:

How aggressive should *you* be with feeds and speeds?

Let's talk for a minute about how the G-Wizard Feeds and Speeds Calculator is *calibrated* for aggressiveness. As mentioned in the calculation article, G-Wizard considers a lot of manufacturer's recommendations in an apples to apples match up (i.e. same coatings and geometries). It then does some very sophisticated number crunching to try to separate out the fact from the fiction. In other words, it tries to determine whether a manufacturer is overly aggressive (great MRR, lower tool life) or overly conservative (great tool life, lower MRR) to get to some "balanced" numbers. It does this by analyzing a minimum of 3 manufacturers for uncommonly used tools and 12-15 for commonly used tools (e.g. endmills or twist drills).

What's the Difference Between Toolroom and Manufacturing Feeds and Speeds?

Using some empirical testing and research, we took all of that data and determined the point we felt was appropriate for *Toolroom Feeds and Speeds*. By that we meant the most aggressive feeds and speeds that could be used for the first time on a new part without much risk of breakage or unduly short tool life. We have a lot of time into this and some proprietary statistical algorithms and it works well. This is a very useful baseline for a lot of machining operations. High-Volume Manufacturing Machining can be looked at like racing--if you're not breaking anything, you're not going fast enough. But, for a lot of jobs, it's faster to avoid breaking anything versus gettiing the absolute edge of the envelope material removal rates. The latter, edge of the envelope MRR's, are *Manufacturing Feeds and Speeds*.

Determining Manufacturing Feeds and Speeds is also a lot like tuning a race car--you go faster and faster until you break something, then you back off slightly. If you're making enough parts, this approach works well. You can tune up your manufacturing process to the best performance for profit, break a few things along the way, and back off slightly.

Using the G-Wizard Gas Pedal to Manage Conservative vs Aggressive Cuts

G-Wizard provides a slider we call the Gas Pedal that lets you configure whether you're more interested in being conservative or aggressive:

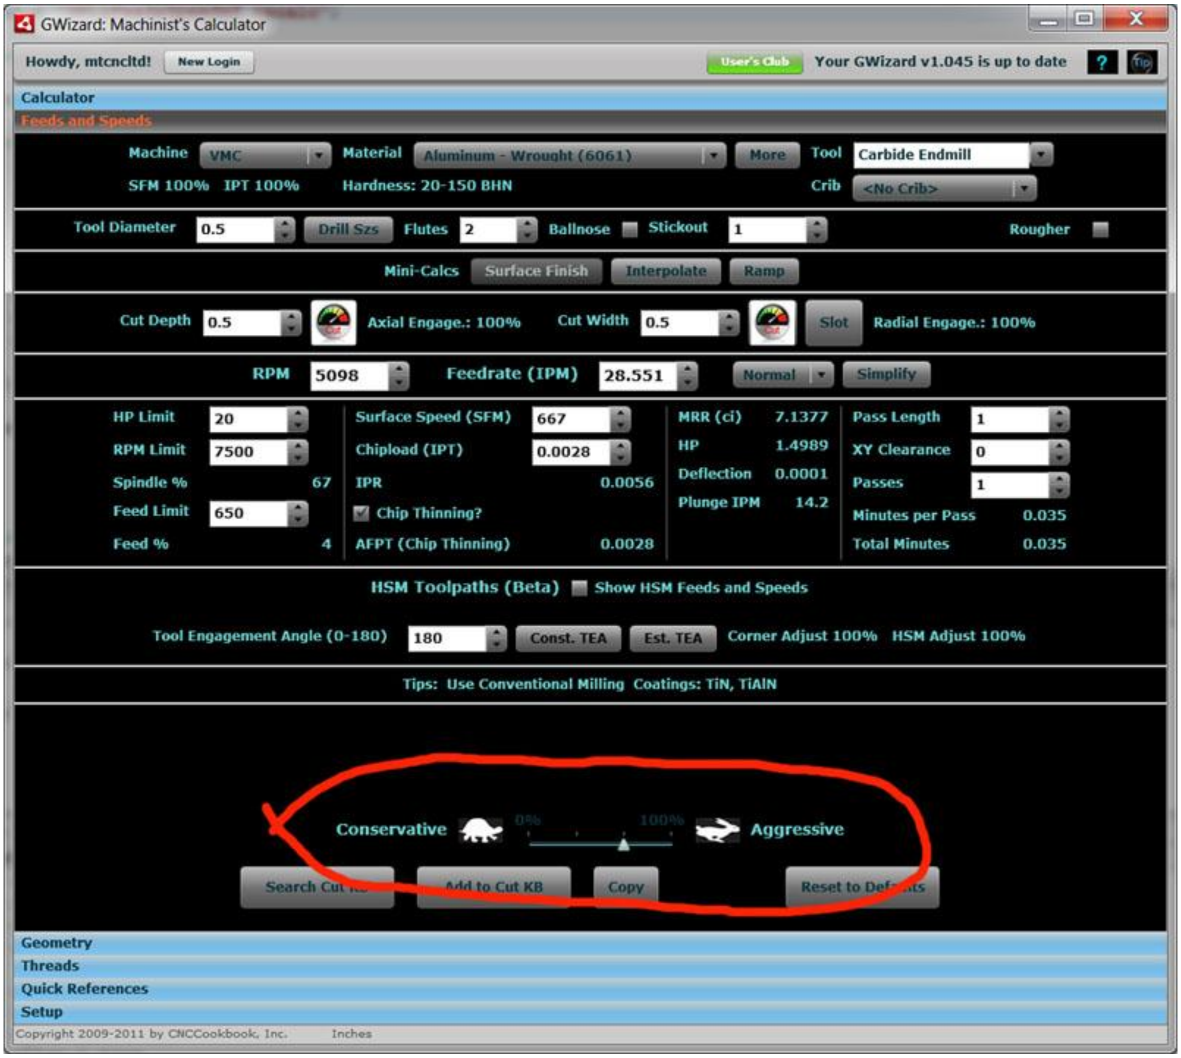

The G-Wizard Gas Pedal

The "Gas Pedal", is depicted by a tortoise and a hare, much like the old Bridgeport manual mills had for speed control. It represents not just conservative versus aggressive but also *surface finish and tool life* versus *material removal rates*. The default best MRR for toolroom situations where you're making one part and you don't want any breakage is one notch to the left of aggressive. This is the default position for the Gas Pedal. The remaining step to the right is only the first rung on the ladder of increasing feeds and spees to get to the best Manufacturing Feeds and Speeds. You're unlikely to break a tool there, but it can happen every now and again with difficult materials and toolpaths.

Where Do We Go After the Most Aggressive Step on the Gas Pedal?

Since the last step on the Gas Pedal isn't really as aggressive as you can get for most jobs, the next step is to add your Manufacturer's data to G-Wizard for your tooling. Unless your manufacturer is cheating pretty good on the numbers by quoting values that are too aggressive, the manufacturer's recommended surface speed and chipload are an excellent next step. G-Wizard will take that data and further massage it. It may well get more performance out of it, or it may soften the numbers somewhat depending on the exact nature of the cut you are contemplating.

Be sure to look at the comparison of the two--G-Wizard's defaults and the Manufacturer's numbers. Check out the results for several typical machining operations you might contemplate. G-Wizard's Tool Crib function under Setup is a good way to setup this kind of comparison. In fact, you can create an entry for several manufacturers including G-Wizard's defaults, tee up a particular set of cut parameters (Endmill diameter, material, flutes, cut width, cut depth, etc.), and then flip through the tools to compare the MRR's.

What Comes After the Manufacturer's Numbers?

After you run the Manufacturer's numbers on a job there are two possible outcomes--they work or they don't. Believe it or not, you're trying to find cases that don't work. They're more valuable than cases that do, because you want to go until something breaks (or surface finish, tolerance, or some other metric suffers) and then back off slightly.

For our friend whose story I quoted at the top, the one that broke an endmill with an HSM toolpath running the manufacturer's numbers, he now has two good data points. He knows G-Wizard's default numbers worked, and he knows he broke an endmill with the manufacturer's numbers. He can now find the spot between the two that gives the best performance without breaking the tool. Depending on how much time he has to experiment, and how much it is costing to break tools, he can get pretty precise.

His next step is to either back off the Manufacturer's Numbers slightly or to pick the midpoint between the two and try that. I prefer the latter approach, but if you feel lucky, just back off a little.

G-Wizard provides a number of tools for "backing off". You could use the Gas Pedal with the Manufacturer's numbers. A better way is to use the surface speed and chipload scaling in conjunction with the Tool Crib. The scaling lets you apply a percentage multiplier to the two values. For example, you could try "0.95" or "0.9" to reduce the values 5-10%.

If you didn't manage to break a tool with the Manufacturer's Numbers, crank 'em up further using the same tools. Try 1.1 or 1.05 to add 5 or 10%. Given the choice, focus on tweaking chipload first. It takes much more time to find out if you have a little too much surface speed because it increases the wear which may not immediately break the tool. You could be already past it after the 5% increase but keep increasing to 20-25% before the tool finally gives out.

The Role of A Cutting Knowledge Base

If this all sounds like a systematic process that will benefit from record keeping and careful analysis, congratulations--you understand!

The right tool to make that systematic process easier and more reliable is a Cutting Knowledge Base. Achieving the best possible cutting speeds are a function of two things:

1. Identifying and tracking the key variables that govern cutting speeds. Remember, G-Wizard manages 40-50 variables for each feeds and speeds calculation, and there are still more I intend to cover over time.

2. Learning by trial and error where the line is based on your shop's way of doing things--you shop's Best Practices, in other words.

This stuff is repeatable. If you put it into a database, and index that database with the right information, it will be a Treasure Trove that can improve every situation--from Toolroom to Manufacturing. Going back to the case of G-Wizard's numbers not breaking a tool where the Manufacturer's numbers did, a Cut KB would track a lot more variables to make it easier to identify when the more aggressive numbers are trouble and when they might be fine. Knowing exactly where the line is and which variables to track is a critical advantage for better feeds and speeds.

I'll talk about exactly how to go about this in an article focused entirely on Cutting KB's. For now, it's important not to lose any valuable data. Whether a cut succeeds or fails, keep a record of it and track as many variables as you can. G-Wizard's Cut KB function is an ideal way to do this and the variables it has chosen to track are backed up by a large body of independent research. Whether you like G-Wizard or not, make sure you write down at least those variables in your notebook. I'll be explaining not only how to go about it, but why those particular variables matter in the Cut KB article.

The sooner you start building your Cut KB, the sooner you'll have enough data for it to start becoming useful.

Coolant and Chip Clearing

CNC Milling Feeds and Speeds Cookbook

The Role of Coolant in Machining

Let's get off Feeds and Speeds for a minute and talk about Coolant. In some ways it's unfortunate that coolant is called "coolant", because it causes machinists to ignore the other roles of using coolant. We actually use coolant for three distinct reasons:

1. Chip Clearing: Spraying a liquid at the cut helps move the chips out of the way of the cutter so it doesn't have to recut the chips or use valuable chip clearance on chips that are not part of the current cut. Recutting chips destroys surface finish and dulls tools much more quickly. In the worst case, a cutter down in a slot or hole can get clogged with chips and get much hotter or even break.

2. Lubrication: Some materials, like aluminum or some steels, are sticky. They have an affinity for the material cutters are made from and will try to weld themselves to the cutter unless we can arrange for some lubrication to make that makes things slippery so the chips are less likely to adhere.

3. Cooling: Liquids, especially water-soluble coolants, are capable of carrying heat away from the cut much more efficiently than air. Plain water conducts heat 25 times more efficiently than air, for example.

I've chosen to put cooling dead last for a reason--while not unimportant, cooling is probably the least important role of coolant!

Let's take a closer look at each of these three critical coolant functions.

Chip Clearing

Chip clearing is by far the most important function. I cringe every time I see a cut being made while the chips pile up. Sure, it's easier to photograph the action, but those piling chips are very hard on your cutter's life and can even lead to breakage. You're much more likely to experience built-up edge (BUE) where chips weld onto the cutter if the cutter must recut the same chips over and over and can't get rid of them. If there is insufficient chip clearing in your machining operation, you may use up all the chip clearance your cutter has available, which can lead quickly to cutter breakage, if the cutter is buried under a mound of chips.

If your machine has no flood coolant, rig up an air blast or mist to get the chips out of there. Get paranoid about having too many chips around. Think about it this way: most tooling manufacturers recommend you turn off the flood coolant when surface speed goes *above* a certain point and you will increase the life of your tooling. If it was all about heat, that shouldn't be the case as more surface speed means more heat.

Lubrication

Lubrication helps the tool to cut more easily and therefore it cuts while generating less heat. As the face of the tool slides across the workpiece, it rubs while cutting. As the chip curls up, it also rubs on the tool, generating more heat. All that rubbing will produce less heat with a little lube, just like any sliding fit would. That's an important role for lubrication, but it's not the most important part of the cooling issue (generate less heat by reducing friction and we don't have as much to cool). A much bigger role for lubrication is reducing the likelihood of Built Up Edge (BUE). This is a big deal as anyone who has seen a big wad of aluminum get welded to their cutter will attest to. Things stop working pretty quickly when that happens!

Fortunately, BUE is material-specific, and mostly applies to Aluminum and Steel that lacks much carbon or other alloying substances. Titanium is another material that has a reputation for being sticky. Use of really sharp cutters with very high rake angles (positive rake is your friend!) can help reduce adhesion tremendously, but it's not enough. Many coatings used on tools can also provide lubrication, although they're fragile and as they wear off and shouldn't be counted on as the primary answer to BUE. In the end, a little bit of mist can deal with this problem as well as flood coolant, so it isn't the end of the world.

Just don't forget to do *something* about lubrication before that wad of aluminum welds itself in all the wrong places. I've heard a number of professional machinsts flat out declare you can't machine aluminum without some form of lubrication (mist or tool coating, and preferably the mist). Even a little spritz of WD-40 makes a big difference with aluminum.

Cooling (and its Evil Twin Shock Cooling)

Which brings us to our next issue, Cooling. The temperature of the tool is probably the biggest factor affecting tool life. A little heat is good, as it softens the work material, making it easier to cut. A lot of heat is bad, as it softens the tool, which means it wears rapidly, becomes dull, cutting forces skyrocket, it gets hotter, and trouble is not far behind. Note that the heat to be tolerated is hugely dependant on the tool material and coating. Carbide takes a lot higher temperatures than HSS. Some coatings, such as TiAlN really need the higher temperatures to do their job properly, and are often used without coolant. The benefits of TiAlN aren't even present until there is enough heat to "activate" it.

There are lots of stories out there where turning the coolant off increased tool life under the right conditions. Carbide is susceptible to micro-cracking under the thermal shocks of uneven heating and cooling. This effect is called "Shock Cooling", and matters a lot to tool life in higher performance applications. Sandvik, in their cutting tool study course, recommend either no coolant or copious amounts of coolant to avoid this problem. It should also be noted that too much heat is not helpful to accuracy, as it makes your workpiece change size.

Let's also talk about coolant type. There are water soluble coolants, and there are oil-based coolants. From a cooling standpoint, the water soluble coolants win. How much? Consider this data:

Coolant	Specific Heat of Coolant	Steel A (tempered) Temp Decrease %	Steel B (annealed) Temp Decrease %
Air	0.25		
Compound oil, high viscosity	0.489	3.9	4.7
Compound oil, low viscosity	0.556	6	6
Aqueous solution of wetting agent	0.872	14.8	8.4
Aqueous "soda product" solution, 4%	0.923	-	13
Water	1.00	19	15

First thing to notice about the table, is that the efficiency of the various coolants at removing heat directly corresponds to the specific heat of the coolant. Second thing to notice is that air is pretty lousy, about 1/4 as good as water. That's actually not as bad as predicted, since water carries heat 25x more effectively than air. The reason for the difference is its hard for the coolant to make efficient contact everywhere it needs to and carry away enough heat. Also, if you're using the right cutting parameters (e.g. feeds and speeds), most of the heat should be carried away in the chips, not the coolant.

It's interesting to note that the oil-based coolants are about half as effective as water-based in terms of their ability to cool the tool and workpiece. Between that and the health considerations, its no wonder a lot of shops have gone to water-soluble coolants--they just cool better. On the flip side, the oil lubricates better (natch), and there are still some applications where machinists may prefer oil (usually turning) to the water soluble coolants.

One last think about flood coolants--above a certain critical surface speed, they all start to work about the same, and the faster you go the less cooling effect they have. One reason for this is that when things are going really fast, there isn't time for a big gout of coolant to make it's way into all the nooks and crannies. Cooling becomes less and less consistent, and this also contributes to the shock cooling effects that make coolant hard on carbide life above certain speeds.

Material Considerations

There are two material-specific considerations for coolant. The first is the tendency to BUE, where the material sticks to the cutter and lubrication is important. The second is the ability of the material to absorb and transfer heat. Some materials do not transfer heat very well--titanium is a good example. Those materials are often more dependent on the coolant for cooling in order to offset the inability of the material to carry heat. That inability makes it harder for the chips to carry away heat and harder for the workpiece to stay cool without changing size due to excess heat. Titanium further compounds the problem by producing relatively small chips.

If the material you're cutting transfers heat poorly relative to aluminum (which is an excellent conductor of heat), steel (a decent conductor), or other common materials, make sure you've got a good flood coolant setup and are using it well.

Performance Recipe: Aiming the Coolant for 5% Higher MRR

Where you aim the coolant matters, whether for chip clearance, cooling, or lubrication. But how many machinists take time to aim the coolant after each tool change? Different tools have different lengths. Different machining operations may also change the optimum aim of the coolant.

You can offset the labor-intensive business of constantly reaiming the coolant by using multiple nozzles preset to a range of heights. With three nozzles, you can cover a pretty decent range. The problem with this approach is you now only have 1/3 of your coolant available at the optimum spot as the other two nozzles aren't doing the right thing. I keep wondering why I haven't come across a machine with solenoid valves that select the right nozzle under program control, but I haven't seen that.

Another solution is just to crank the volume and pressure so that even with only one of three nozzles in the right place, there is a wall of coolant. The high pressure and volume option is available for most machines and is worthwhile.

Lastly, you can retrofit your machine with a Spider Cool nozzle that allows you to pinpoint your coolant by twisting a know on the control panel, and that can track your tool changes and change the aim automatically.

What's good aim worth? Spider Cool says its good for 5% better MRR, and they have a free trial to prove it. Might be worth checking into your coolant aim.

Performance Recipe: Through Spindle Coolant

The most difficult chip removal job is down in a hole. For this case, through spindle coolant (TSC) is a huge benefit and sometimes the only way to drill deep holes. I have also seen data that suggests you can increase the MRR on indexable tooling by as much as 15% with TSC. That's a pretty big bump.

Another advantage for TSC machines: you can often run more flutes, particularly when profiling.

Horizontal Mills and Lathes Have Gravity Helping Clear Chips

Don't overlook the benefits of gravity for machining. On vertical mills, gravity makes it harder to pull chips out of deep holes. On lathes and horizontal mills, gravity makes it easier. Kinda makes you wonder why nobody has a machine that cuts from underneath. You'd need a pretty crazy palette loader so you could drop the workpiece onto the table and then flip it for cutting. Too far out, I thought, but then I discovered such machines actually exist. They're called "Inverted Spindle Lathes" and are a potent alternative to bar fed lathes.

Here is a link to an MMSOnline article about them.

Exotic Recipe: Alcohol as Coolant

Datron uses an ethanol alcohol mist as coolant for its HSM machines. They make a good case for it:

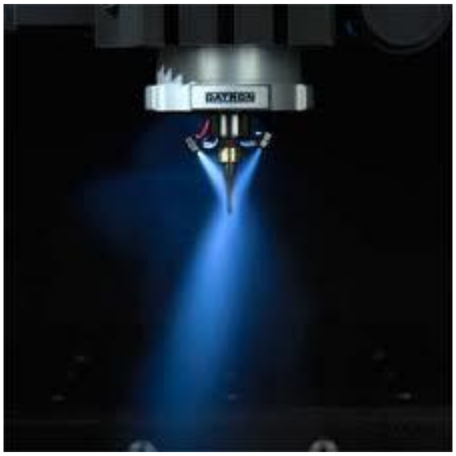

It is ideal for high-speed, micro-tooling of non-ferrous metals and some plastics due to a thinner-than-water viscosity that allows the ethanol to quickly cover and cool more of the surface area on fast-moving parts. The low evaporation point of ethanol makes it an efficient cooling and lubricating solution. Since the ethanol simply evaporates, disposal, recycling and their associated costs are a thing of the past. Plus, ethanol coolants leave no residue on machined parts, which makes costly secondary operations, like de-greasing, obsolete — maximizing throughput, increasing efficiency and ultimately improving a manufacturer's bottom-line.

Conclusion

For many high performance applications, machinists can focus on chip clearance and lubrication and ignore the cooling issues. Above a certain surface speed, many tooling manufacturers recommend turning off the flood coolant and using an air blast (perhaps with mist for lubrication) to clear chips. Materials that don't transfer heat well like titanium will require flood coolant regardless.

Additional Information

Selection of Cutting Fluids in Machining Processes: Good overview of some of the factors involved in selecting coolant type.

Turning Down the Heat in a Cut

CNC Milling Feeds and Speeds Cookbook

Heat is the enemy of your cutting tools. The history of cutting tool improvements has largely been about making tools that can stay sharp when it gets hot, hot, hot in the kitchen. Coatings like TiAlN even thrive on a certain amount of heat and have to be run to achieve their threshold temperatures before the full value of the coating can be activated.

The problem with heat is that it softens your tool's cutting edges. Not long after such softening, that tool will be done and will need to be replaced. What techniques and options are available to machinists to turn down the heat on their cuts?

Reduce Spindle Speed

We like to think intuitively that we should slow down, and indeed, reducing the surface speed by cutting the spindle rpm will reduce the heat generated. In fact, a slight reduction in rpm can radically improve the situation. The amount of tool life increase is given by a model known as Taylor's Tool Life Expectancy Equation. Most manufacturer's feeds and speeds recommendations are based on a tool life of 15 to 30 minutes of cutting. Check your manufacturer's tables for the exact figure. If you want your tool to last longer, the easiest way to improve the situation is to reduce the surface speed (spindle rpm) slightly.

The G-Wizard Machinist's Calculator includes a Taylor's Equation calculator on the Hardware Profile tab, under Setup. That way, you can create a special "long tool life" profile if you want by having G-Wizard adjust surface speed to a percentage (less than 100%) of its normal recommendations. The results may surprise you: a 10% reduction in surface speed adds 20% to the tool life and a 20% reduction in speed adds 60%. A little bit of reduction can make your tools go a long ways.

For most professionals, it typically makes more sense to keep the Material Removal Rates as high as possible, but there are times where tool life is more important. For example, you wouldn't want to lose a tool in the middle of a long 3D profiling operation on an expensive part.

Increase Feedrate?

I list increasing feedrate with a question mark just because a lot of people will see it that way. Increasing the feedrate has got to be harder on the tool, right?

There are two cases where it is actually easier on the tool, and both occur when you're too far below the sweet spot threshold on the tool's chipload. The absolute worst case for feedrate is when you're feeding so slowly that it causes the tool to rub or burnish instead of shearing off a nice chip. This can actually produce a pretty decent surface finish due to the burnishing effect (like a "wiper" insert), but it is extremely hard on the tool. All that rubbing heats it up in a hurry. This effect occurs when the chip thickness (chipload per tooth) falls below the radius of the cutting edge on your tool.

I give more detail on this on the Chip Thinning page, but a picture is worth 1000 words:

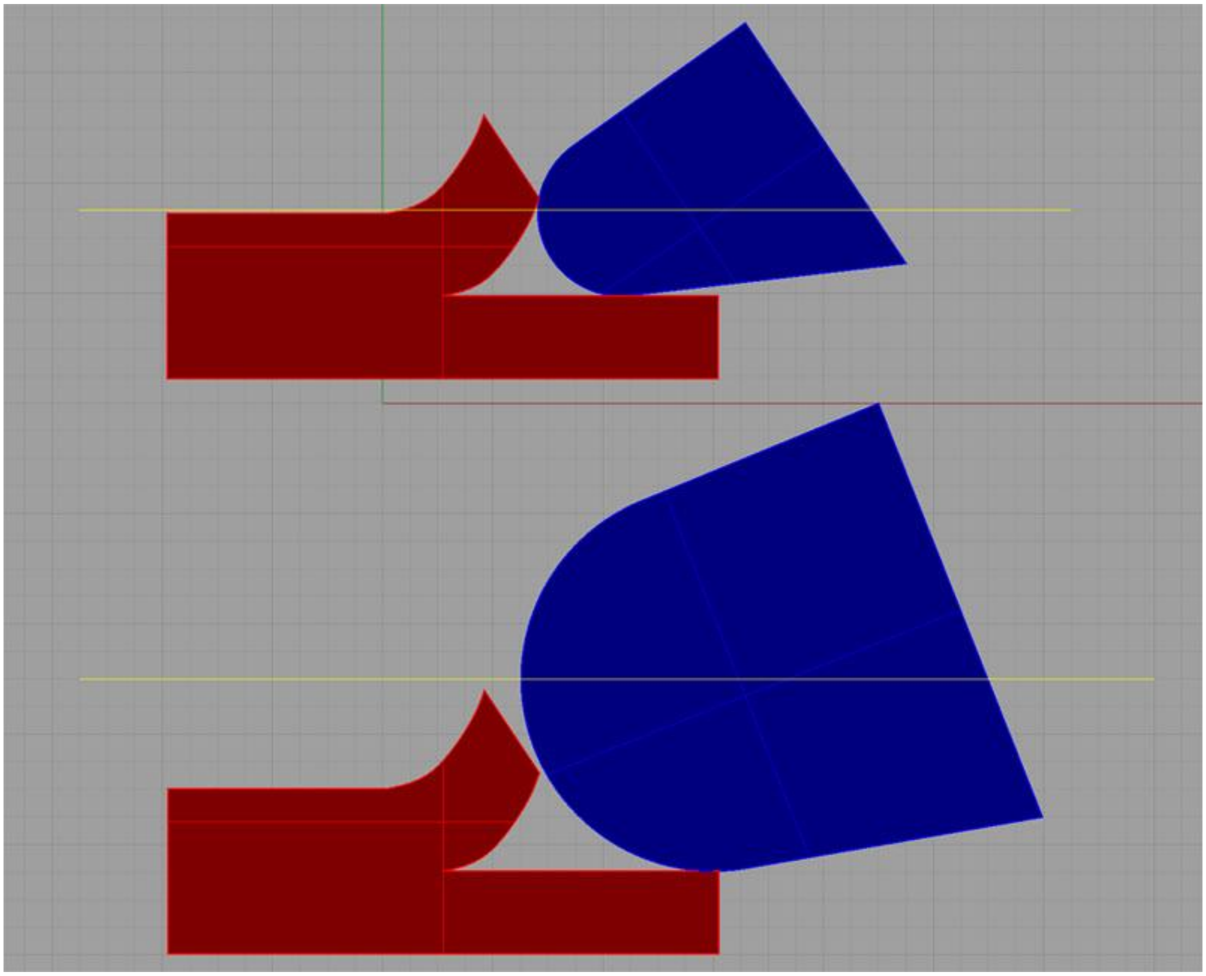

Cutter edge radius centerline travels along the yellow lines. If the radius is too large relative to the depth of cut (bottom), all the force goes to pushing the chip under the edge. This is the "rubbing" effect you'll hear talked about when feedrate and hence chipload are too low. Use a calculator like G-Wizardthat has radial chip thinning to avoid this problem.

The second case where faster feeds help is that larger chips carry away more heat from the cut. Heat is generated on both sides of the shear line, but the chip is also deforming and ultimately tearing off, so more of the heat goes to the chip side. The trick is to bury that heat in as large a chip as possible and fling it away before it can transfer the heat back into the workpiece.

To achieve optimal heat reducing effects through feedrate, use the recommended feedrates. Going to higher feedrates can result in higher MRR, which is an excellent goal, but it won't further reduce the heat buildup.

Crank Up the Coolant

This is another one, like feedrate that deserves a question mark. Coolant can be something of a two-edged sword, and in fact, the name itself is a bit of a misnomer. By far the biggest function of "coolant" is clearing chips so you're not recutting them. For a worst case picture of what happens when you don't clear the chips, imaging cutting a work hardening material like stainless. Leaving those chips in the cut is like dropping a handful of super hard nuggets in the path of your cutter just when its trying to work its hardest. Not a good idea!

The second most critical function of Coolant is lubrication for materials that want to stick to tooling edges like aluminum and stainless.

It's only when we get to third place that we're talking about heat. Perhaps "Coolant" should have been called "Chip-Clearant" or even "Lubricant" (some coolants are oils for a reason!). The cooling effects of "Coolants" are most important for materials that don't transfer heat very well. Aluminum and copper alloys transfer heat extremely well. It turns out that Titanium does so very poorly. So Titanium benefits from the cooling effects far more than aluminum, which benefits more from the lubrication and chip clearing of coolant than the cooling.

That's not to say there is no cooling value when cutting materials like aluminum, but it is far less, and we may do just as well dry machining provided we can deal with the chip clearance and lubrication issues.

The real issue with coolant is a phenomenon called "shock cooling". Picture your cut under microscopic conditions. Imagine something less than the "Wall O' Coolant" we hope we're running. Perhaps at these scales it is more of a "Trickle O' Coolant". A drop hits the super heated tool and flashes instantly to steam. That's a shock to the tool, and carbide is especially susceptible to a reduction in life due to shock cooling, which can make it crack surprisingly quickly. Cooling can also interfere with the action of coatings like TiAlN that have to reach a certain temperature before they're fully activated to their maximum potential.

The bottom line is you'll need to check carefully with your tooling manufacturer about your coolant practices. At a certain speed, a surpising number of manufacturers are suggesting you switch to air to clear the chips and leave the flood coolant off. Also, be aware that different coolants have different levels of cooling efficacy, with water soluble being the best. See the Coolant article for more.

Reduce Radial Engagement

This is the last heat reducing tool we'll talk about, but it may be the most important one from a productivity standpoint. Radial engagement is how much of the radial (diameter) dimension of your tool is doing the cutting. For a slot, it's 50%, because the back side is open. Think of it this way: in that slot, each edge of the tool spends 1/2 of its time cutting and making heat, and 1/2 of its time exposed to air and coolant where it can release the heat. It's pretty easy to see that the worst case cut is plunging into solid material where 100% is involved in cutting and there is no cool-off time. No wonder we plunge slowly or prefer a helix or ramp to start a pocket!

By now you will have seen the tremendous increases in performance that are possible if we reduce the radial engagement and run a high speed machining tool path that avoids pushing the tool into a corner. Such tool paths do two things. One, they up the feeds and speeds tremendously so the MRR's are still very high even though we may only be running 15-30% radial engagement. They can do this because the forces on the tool are constant and don't skyrocket in the corners, and because of the second thing--such light radial engagement gives each flute a lot of cooling time in air rather than in the material.

RobbJack has a great article on how to think about Radial Stepover and heat when machining hard materials. Check it out.

Now that You've Reduced Radial Engagement, Crank up the Spindle RPMs Further to Reduce Heat and Join the HSM Revolution

It hardly seems fair. We just got done telling you that slowing the spindle is a great way to reduce heat and gave you the Taylor Tool Life equations as backup. How the heck can more spindle rpms reduce heat?

The answer is from a peculiar phenomenon that led to the whole HSM movement. Consider this amazing chart from Dr Herbert Schulz's, "History of High Speed Machining":

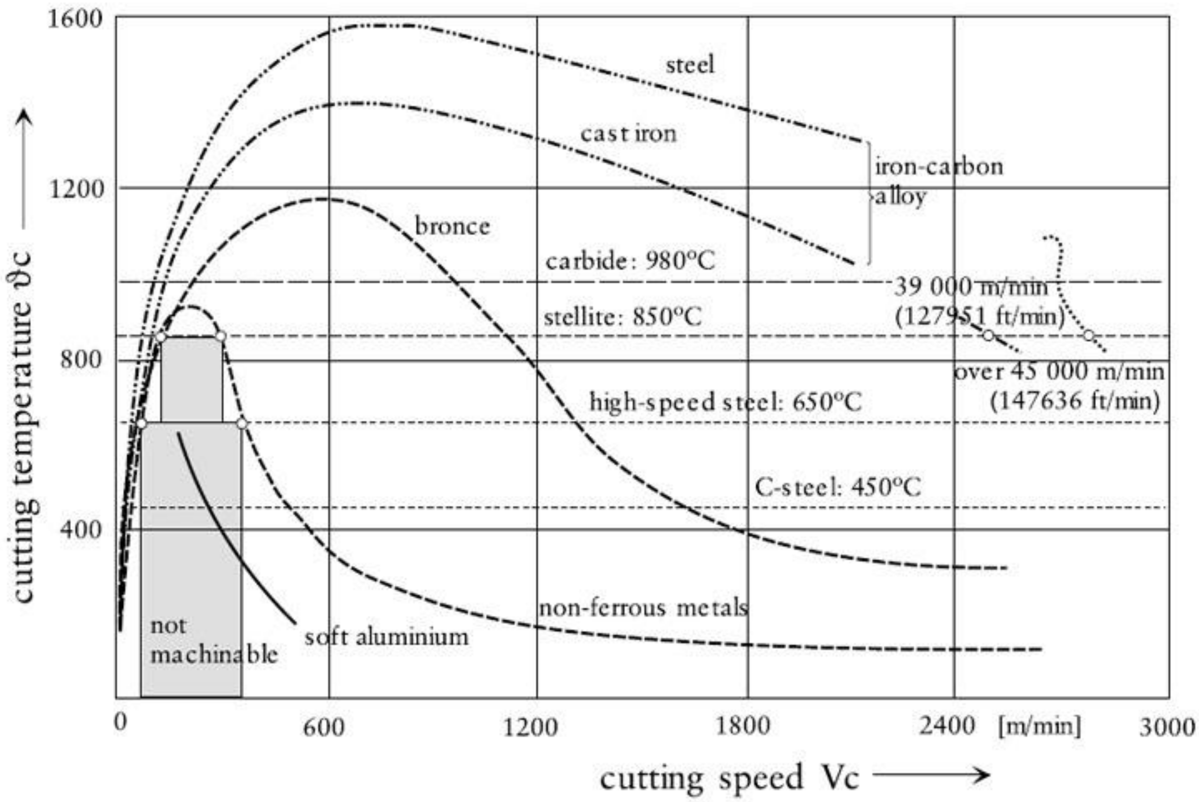

Fig. 1. Machining temperature in milling at high cutting speeds.

The dotted lines represent temperatures at various surface speeds. Note that all of the materials go steadily up and then eventually start dipping back down again as surface speed increases. Somehow, temperatures decrease beyond a certain spindle rpm!

This chart is in meters/minute, so multiple the values by about 3 to get to SFM. For aluminum, we have a pretty good dip by the time we're hitting 1000 SFM, for example. In fact, it's temperature is more equivalent to less than 300 SFM on the other side of the aluminum curve--that's nothing for aluminum. Heck, if we have a fast enough spindle, there's even room to run HSS faster and get lower temperatures (you'll note various cutter materials critical temperatures are also marked off--stay below the line for your cutter!).

Steel and cast iron taper down more gently than aluminum, but the effect is still alive and well. Yes Virginia, there surely is some strange behavior when you start in with that HSM stuff!

The same research showed that cutting forces also come down, and that's at least one reason why the temperatures drop, and why for HSM machining in the right rpm ranges, you can achive high MRR's with lower cutting forces.

Tool Deflection Control: Critical to Your Success

CNC Milling Feeds and Speeds Cookbook

Bad Things Happen When Tools Deflect

The machine's spindle is powered by a motor of a certain amount of horsepower. Think of that motor as pumping energy into the workpiece. The only thing connecting the motor to the workpiece is your cutter, so it potentially has to conduct a lot of force. Hopefully, most of the force is converted into cutting and results in chips. Inevitably, some of the force will be converted into less desirable by products such as heat, tool deflection, distortions in the workpiece, or vibration (ultimately leading to chatter if it's bad enough).

Given the dimensions of a typical cutter, especially smaller ones, and their rigidity, machinists need to be concerned with tool deflection. This should come as no surprise as machinists learn early on that everything bends, it's only a matter of how much force it takes and what scale we're talking about. For microscopic forces, the amount of bending may be measured in ten thousandths or less. For large forces, bending can start to be a problem during machining.

This article is all about understanding what sorts of problems tool deflection contributes to, using the G-Wizard Calculator to figure and optimize cutting parameters for tool deflection, and tips for how to avoid those deflection problems.

Tool Deflection and Accuracy

The first thing you're probably thinking is that if the tool deflects it's going to lead to some accuracy problems, and it sure will. Consider a 3 flute 1/8" carbide endmill. Let's say we're cutting 1/2" deep and we've left 5/8" of cutter sticking out of our collet chuck. Further, we're using it to cut pockets in 6061 aluminum and we're running the spindle at 7500 rpm. How wide can we make our cut each pass and keep tool deflection below 0.001"? The answer may seem surprising, but G-Wizard calculates a maximum cut width (stepover) of 0.0766" will yield 0.001" of tool deflection. That's about 60% of the cutter diameter.

That's what's needed to keep cutter deflection within 0.001". Most would consider that suitable for roughing, but way off the mark for finish passes where accuracy matters. If you're trying to hold half a thou, you'd probably like your tool deflection to be half of that. Let's call it 2 tenths (0.0002") to keep it simple. Now you can only afford a finish pass with a stepover of 0.0016"--mighty thin indeed.

Effects of Climb vs Conventional Milling on Tool Deflection

While we're on the subject of tool deflection and accuracy, let's consider the effects of climb vs conventional milling. The following illustration contains small arrows (often called vectors) showing the direction of tool deflection as the cutter moves along the toolpath:

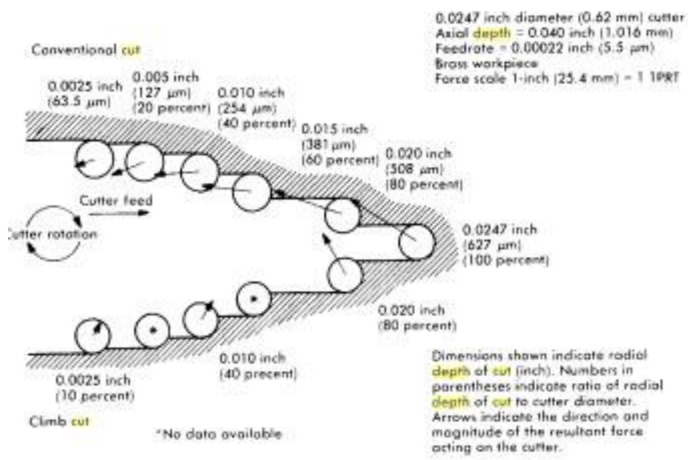

The arrows show where the cutting force is attempting to deflect the cutter. Conventional cut at top, climb cut at bottom.

Note how the deflection force vector is more nearly parallel to the cut with conventional milling (albeit the arrows are longer, showing there are higher cutting forces). With climb milling, the arrow is nearly perpendicular to the cut. If your cutter deflects 0.001", wouldn't you prefer it to be nearly in the direction of travel? The alternative is for the cutter to plow deeper into the wall or pull away from the wall. Either case will introduce more error in the part being machined.

Try climb for roughing, because you can rough faster and the tool deflection effects on accuracy don't matter--the finish pass will deliver the accuracy. You can rough faster because cutting forces are lighter and the thick-to-thin chip profile carries the heat away on the chip. That thick-to-thin + carrying the heat away is particularly crucial for touch work-hardening materials like stainless. It also results in a nicer surface finish if you can afford to climb for the finish pass.

But, you should switch to conventional milling for the finish pass if you're at all deflection challenged (use G-Wizard to see if your tool diameter and stickout result in small enough deflection for your finish pass). At the very least, one should avoid too much depth of cut when climb milling lest it invite deflection. The same article suggests that when deflection is to be minimized, use no more than 30% of the diameter of the cutter for conventional milling and 5% for climb milling. Of course here again, if you have G-Wizard, you'll know what kind of deflection to expect and whether it's a worry.

Climbing to rough and conventional to finish is inline with the consensus over at Practical Machinist as well.

Properly factoring deflection can help you avoid the need for an extra spring cut, which saves time and money.

Tool Deflection and Tool Life

Okay, let's say we have to start that cut we've been talking about by ramping down in a full slot. How much depth of cut can we afford? G-Wizard says 0.1296"--quite a bit less than the 1/2" deep we're aiming for. It'll take a zig zag ramp or helical interpolation to avoid exceeding our allowance. Do we care? After all, this is just a ramp or entry move, the actual pocket or profile is yet to be machined.

In the end of the day we do care, because of tool life. Carbide is a very brittle material. Too much flexing and you'll break the cutter. Think of your tool deflection as being akin to runout. Suppose I told you your machine spindle had 0.001" of runout? You'd be wanting to have that spindle rebuilt and would consider that unacceptible. Yet you can get the same thing geometrically by pushing the tooling too hard through cut deflection. We just saw how little it takes to deflect our 1/8" carbide endmill in aluminum by 0.001".

There are no end of writeups about the evils of spindle runout for tool life, particularly in demanding applications or when using smaller diameter tools where runout starts to be a significant portion of the tool's diameter. Runout causes uneven work for the flutes of the cutter instead of spreading the load evening around all the flutes. Imagine how much worse it must be if the poor tool is flexing in the same direction (remember those deflection vectors from the diagram above) by almost 0.001" while spinning at 7500 rpm!

For best tool life, try to keep tool deflection below 0.001", even when roughing, and much less for small tools. Again, think of it as runout.

Tool Deflection and Chatter

Chatter is a resonant phenomenon, like ringing a bell. The combination of your machine, tooling, and feeds/speeds yields certain frequencies that the combination is particularly suspecptible to.

Now imagine that deflected cutter as it spins. The deflection makes it act almost like a hammer swinging against the bell. The flexing combined with the flutes and the spindle rpm create a rhythmic pulsing, which is exactly what excites chatter. If you do a little research, the major tooling companies will tell you that a tool deflection of more than 0.001" starts to be enough that chatter can set in.

Calculating Tool Deflection

You've probably been wondering for most of this article how you're actually supposed to go about dealing with tool deflection. Is it something you can calculate? Do you measure it somehow? Or do you simply have to be paranoid about it and maximize tool rigidity every way possible. The answers are yes, yes, and yes. Okay, if we're going to be paranoid anyway, why bother trying to determine deflection? Because of the flip side of the deflection coin, which is cutting force. More cutting force equals more deflection, but it also equals more material removal. The reason to be paranoid is so that we can use as much force as possible and thereby gain as much productivity as possible. Knowing how much force can be used is why we want to calculate or measure tool deflection.

Let's dispense with the measuring piece up front. In theory, you could measure tool deflection with the right instrumentation, but outside of a research lab "that ain't happenin'." You could also measure it through its effect on the workpiece. If a vertical was has a draft angle, tool deflection is a likely culprit. However, as we have seen above, the direction of the deflection is very hard to predict without complex analytical models, and it changes throughout the cut. Other than noticing the draft and concluding you may have a problem, this is not a very productive line.

That brings us to calculating tool deflection. This is something that is actually possible to do in the shop, and relatively easily. One of the unique capabilities of the G-Wizard Machinist's Calculator ist that it computes tool deflection as one of the many pieces of information it offers for feeds and speeds. Tool Deflection is shown right below the calculation of spindle power used in the cut:

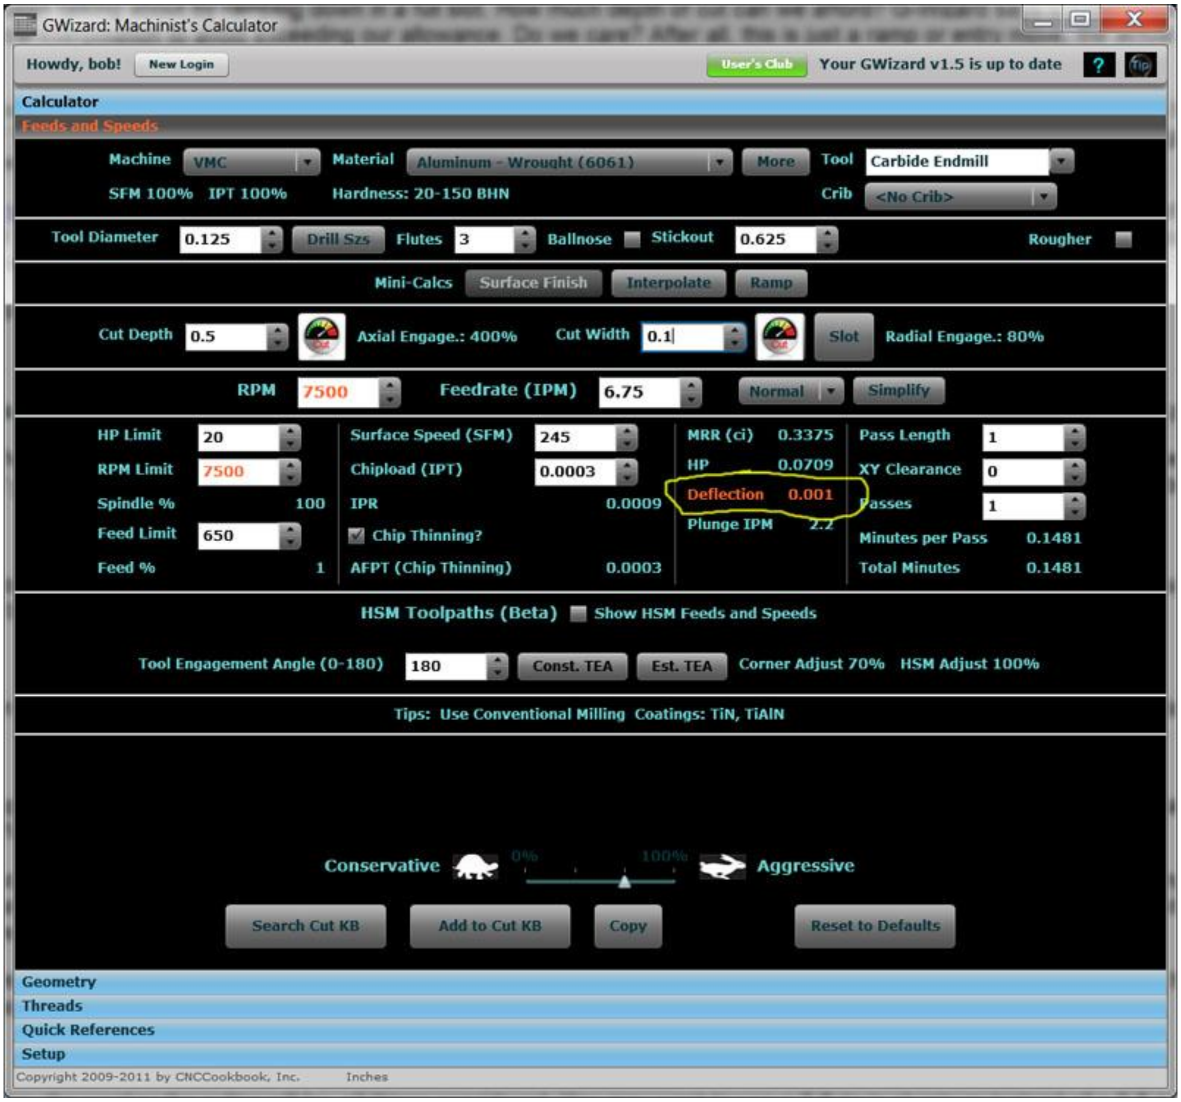

Calculating tool deflection...

This is our old friend the 1/8" carbide EM in 6061 aluminum. Note that the deflection is shown in orange, which G-Wizard uses in a number of places to indicate a limit has been reached or exceeded and there may be a danger zone. To move out of the danger zone, you can use G-Wizard's Cut Optimizer (see below) to calculate the exact depth or width of cut (or if necessary reduction in feedrate) that keeps the deflection within the limits you choose to set.

Are There Rules of Thumb for Tool Deflection?

What if you don't have a calculator like G-Wizard? Are there any rules of thumb for estimating or avoiding tool deflection?

Unfortunately, the math involved is really complex and not very intuitive. It also involves a lot of variables. Tool deflection is affected by number of flutes. It changes as the third power of length and the fourth power of diameter. Et cetera, et cetera. Simple rules like, "Don't use a depth of cut greater than 2x the cutter's diameter", just can't capture the complexities of the math.

If you don't have a deflection calculator, you're left with making the most rigid possible tool holding and keeping an eye on the cut. If accuracy, surface finish, chatter, or tool life become a problem, tool deflection is one thing to consider looking into.

Cutter Rigidity: Resisting Deflection

There are only two ways to reduce tool deflection--either reduce the forces acting on the cutter or increase the cutter's rigidity. We can always reduce the forces, but that will compromise our productivity, so let's focus on increasing rigidity by considering the following factors:

Use Carbide

Carbide tooling is 3x more rigid than HSS. When rigidity is an issue, use carbide. This is particularly important for smaller diameter tools and cases where a lot of stickout is needed for longer reach.

Minimize Stickout

On our 1/8" endmill example, consider what happens if we didn't carefully minimize the stickout of the cutter. Let's say we left it at 3/4" stickout instead of 5/8". In that case, we will find we can now only cut 0.0429" stepover. That's only about 30% of diameter. We sure lost a lot of stepover for only having increased the cutter stickout from 5/8" to 3/4"!

That's because rigidity decreases as the third power of length. Twice as much length is 8x less rigid. Therefore, very small changes in length cause third power changes in rigidity, so always use the minimum amount of stickout that still leaves clearance for the job. Some CAM programs and simulators will report how much stickout is needed to clear the work, fixtures and other obstacles. If your thinking of maximizing your tooling flexibility by leaving a bunch of stickout, you'll be maximizing the bad kind of flexibility which ultimately leads to fewer options.

In the article on chatter, we'll learn that there are strong arguments for standardizing stickout in order to make chatter more predictable. If your shop chooses to do that, be sure to have multiple standards so that there is tooling set up with less stickout as well as tooling with a lot of stickout.

Number of Flutes and Flute Length

Flutes weaken the cutter's rigidity--the more flutes, the weaker the cutter will be, all things considered. You may want to use a 2 flute in aluminum instead of a 3 flute just to pick up some rigidity. In addition, the length of the cutter that is fluted also matters. Stub length flute cutters are more rigid.

Increase Tool Diameter

Diameter is an even stronger determinant of rigidity than length. Length decreases rigidity as the third power but diameter increases rigidity as the fourth power. A cutter with twice the diameter is therefore 16x more rigid. It may sound glib to suggest increasing the tool diameter when you've been handed a print and told to make that part, but consider:

- Will the designer increase the limiting radii even a little bit?

- Can you get a metric dimension tool in that's a little bigger than the largest Imperial that fits or vice versa?

- Can you add a toolchange and do more than one pass, perhaps starting with a much larger roughing tool?

When rigidity increases as the fourth power, even a little more diameter can make a big difference.

Optimizing Cuts for Tool Deflection

Did you ever wonder if there was some combination of cut depth and cut width that might be optimal for your machining operation? It turns out there is, provided your satisfied that the term "optimal" relates to maximizing the depth or width of cut while keeping tool deflection within bounds that you've set. As we've discussed, a reasonable bound while roughing might be 0.001". For finishing passes, much less, perhaps 0.0002" is good.

G-Wizard contains a facility called the "Cut Optimizer" that is expressly designed to help solve these kinds of problems:

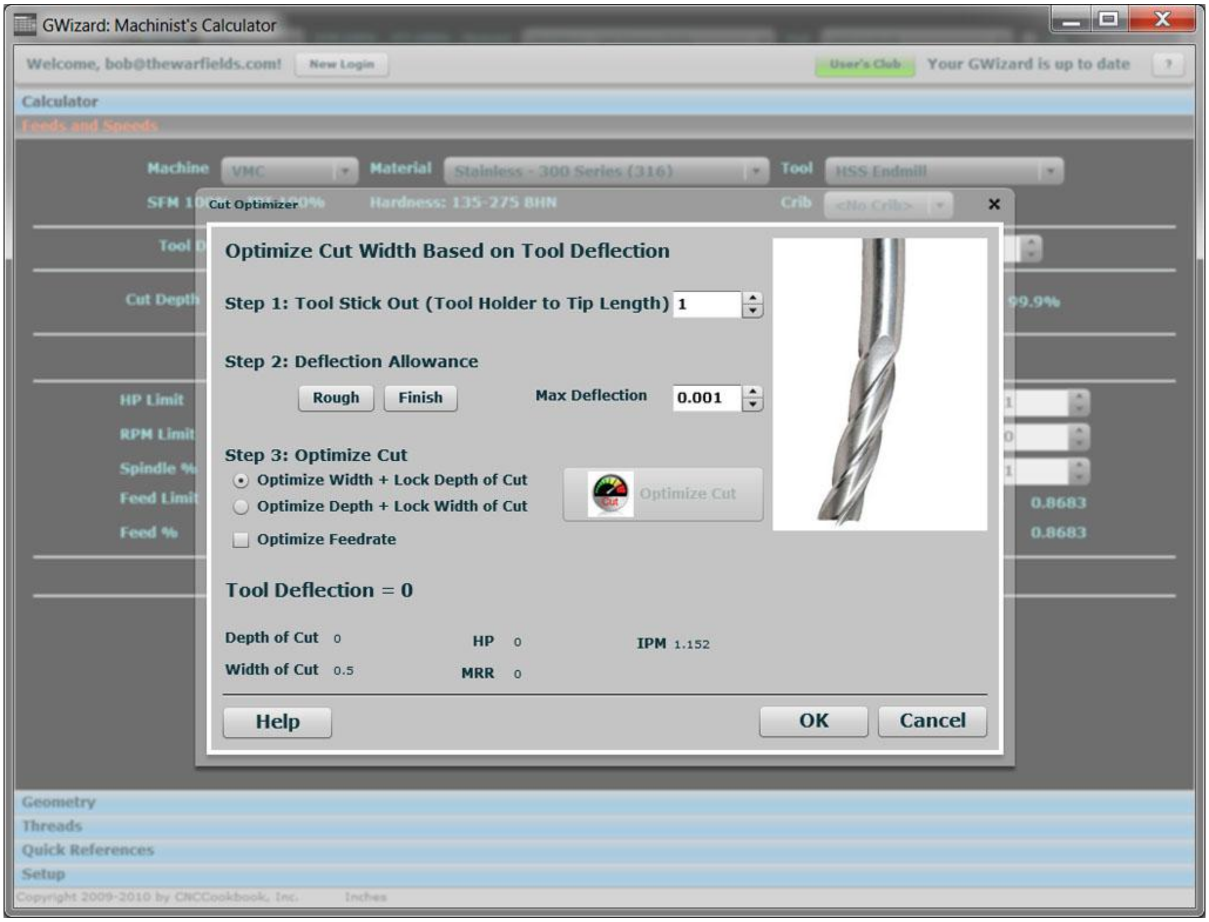

G-Wizard's Cut Optimizer maximizes MRR while keeping tool deflection within limits you've set...

Using Cut Optimizer, you can master the interplay between Cut Depth, Cut Width, Tool Deflection, Feedrate, and Tool Stickout.

Deciding on Best Depth and Width of Cut Using the Cut Optimizer

I got a note recently from a G-Wizard user who wanted to know how to decide on best depth and width of cut when milling. It's a great question. Most machinists, use rules of thumb and habit more than anything else unless the situation dictates something in particular based on the dimensions of the feature being machined. They're used to using some fraction of the cutter's diameter or some figure that they got to some other way through habit (40 thousandths or some such is what they've always used). Perhaps their CAM program has a hardwired default that is a percentage of the cutter's diameter.

But these values, while they have worked over time, are not necessarily optimal figures with respect to material removal rates, tool deflection allowances, or a host of other variables we might choose to consider. What's a more systematic way to approach the problem?

First thing is we have two variables (width and depth of cut), so it's hard to make progress unless we can nail one of the two variables down and focus on the relationship of the other. It's usually pretty easy to nail down one of the variables based on the situation. Let's divide our work into two categories:

- Slotting: I'll generalize this to be any situation where the material to be removed is very close to the cutter's diameter. It may be a slot, or it may involve interpolating a hole or pocket that's only a little bit larger than the endmill's diameter.

- Pocketing: Here again, I will generalize this to be any situation where the cutter's diameter is quite a bit smaller than the dimensions of the material to be removed. That doesn't mean there isn't some inside radius or other feature that isn't more like the slotting example, but for the most part, we have some room to work in. Note that profiling will be considered to be the same as pocketing for this discussion.

Okay, so now we have to take the task before us and decide whether it is closer to slotting or pocketing. The reason I've defined these two the way I have is that it informs our choice of which variable to work on first. If we are slotting, the cutting width is the first variable. If we are pocketing, the cutting depth is the first variable. Why?

When slotting, the feature is very close to the cutter's diameter in size. We can't take a 1/2" endmill and use it to make a 1/4" slot. In general, we want to use the largest diameter endmill that fits the feature, and then we pretty much have to make at least one cut that is full width. Once we're cleared that cut, anything remaining is handled the way we would under pocketing. So, when slotting, we focus first on cut width and make that the cutter's width to get started.

When pocketing, our limitation will be the smallest inside radius we have to deal with as well as the depth of the pocket. Remember, it may be advantageous to make two passes. The first with a cutter that has a diameter too large for the smallest inside radii we have to deal with. That's a roughing pass that uses a larger cutter just to get done faster. The second pass is a finishing pass, and must use a cutter whose diameter is less than or equal to that required to reach into the smallest internal radius the pocket holds. Note that we can go around an outside radius (a boss) with any diameter cutter, it is the inside radius that limits us.

So, we pick a cutter that is either as big as the smallest radius, or we choose to go two passes and go with a larger cutter. Let's leave the two pass issue aside for the moment, because figuring out when that is optimal can take a bit of trial and error. Its similar to think of one pass. Given that the cutter is chosen, we can choose just about any width of cut we want. So how do we nail down a variable when pocketing? On the slotting case, I like to nail down cut width. On the pocketing case, I prefer to nail down cut depth.

In general, we get a nicer finish if we cut the pocket in as few layers as possible. CAM programs are good at layering down into the pocket, so we can pick arbitrary depths of cut. If I can, I like to do it in one layer for a pocket that isn't two deep. If not, I prefer the depths of the layers to be equal. In other words, I wouldn't go down 1/4", 1/4", and then 0.19" on the third layer. So pick a layer depth that satisfies that criterion.

Now, in both cases we have locked one of the two variables--slotting locks width, pocketing locks depth. We need to determine the best value for the variable we left floating based on the value of the one we locked. This is where the G-Wizard Cut Optimizer makes it easy. Enter the values you know for the cut and let the Optimizer figure the value for the floating variable.

For example, let's suppose we need to cut a pocket that is 3/4" deep in 6061 aluminum. The smallest internal radius is 1/4", so we've decided to use a 1/4" 3 flute carbide endmill. Here is the problem set up in G-Wizard:

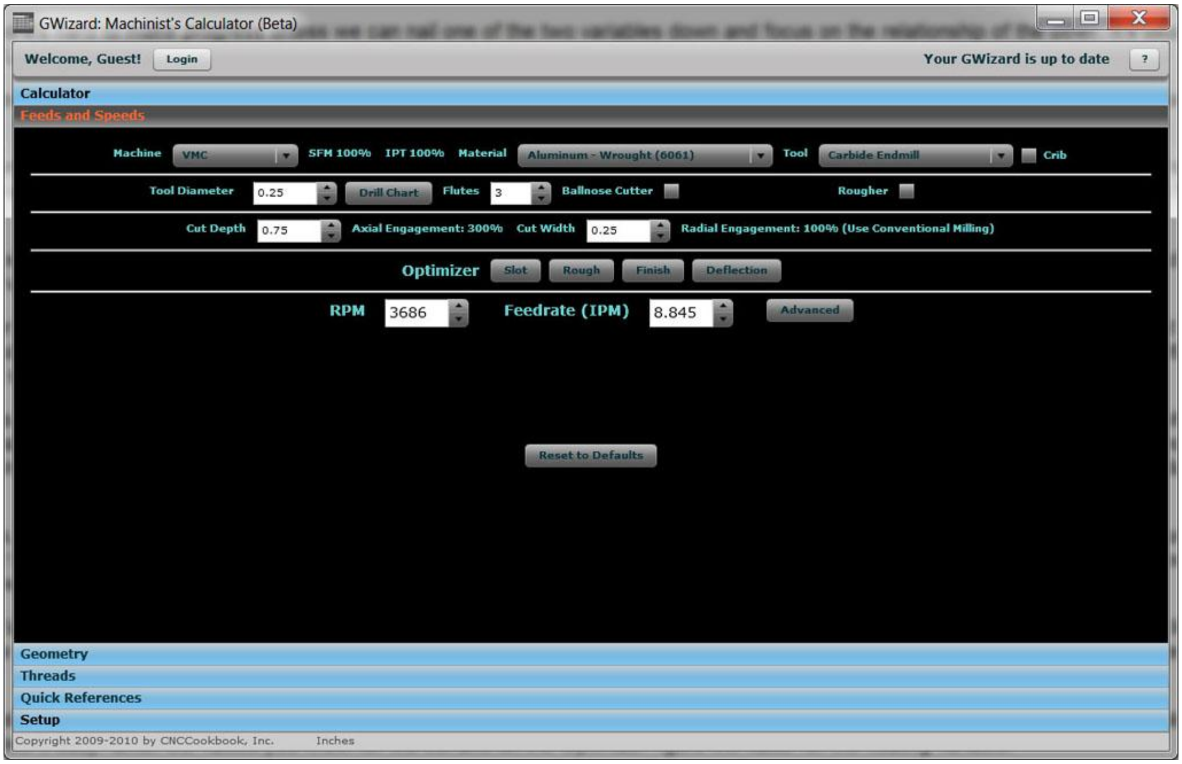

Material, Tool, Tool Diameter, Flutes, and 3/4" Cut Depth Entered...

Now we can invoke the Cut Optimizer just by pressing the "Rough" button:

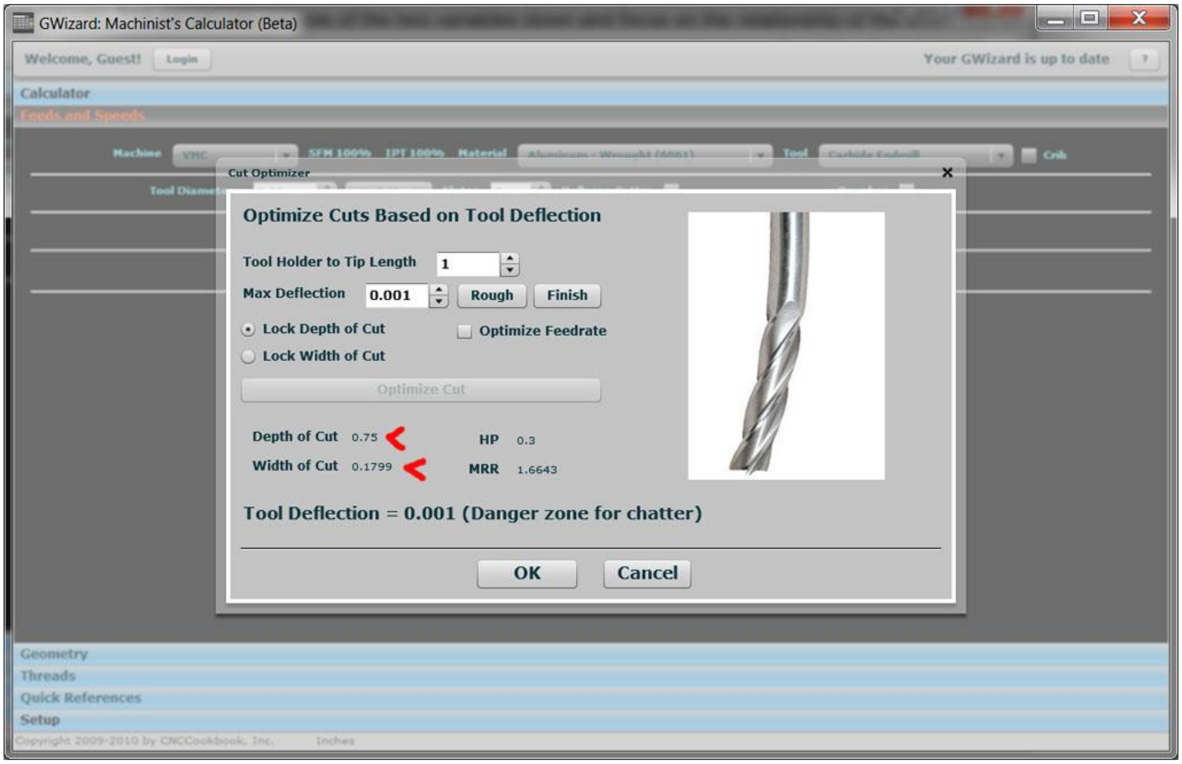

As you can see from the red arrows I added, for a 3/4" depth of cut, this endmill can handle no more than 0.1799" width of cut when roughing. Let's round that down and go with 0.170"

Press the finish button to see what sort of finish allowance we should have the CAM leave for our finish path and we get 0.0052". That's a pretty light pass, but 3/4" is deep for this 1/4" endmill. Here's an interesting thought: if we reduce tool holder to tip length to 0.9" instead of 1", we can increase that cut to 0.0095". That gives you an idea of how important it is to keep the tool stick out as little as possible. I'd be inclined to go with choking up on the tool and a finish width of cut of 0.009" were this my job. The other thing to consider is two levels of finish pass. If we don't mind taking two levels and still choke up on the tool, we can get 0.015" width of cut for the finish. That's about as much as I like to take on a finish pass.

The problem with this cut is its a little bit deep for our 1/4" endmill. That's a 3:1 ratio of diameter to depth. We can tell its straining because the max recommended widths of cut are so light. If I had a CAM program that made it easy to make the roughing pass with a bigger cutter, I would be tempted to jump in with a 1/2" endmill (or maybe even larger) for the roughing pass and then go to the 1/4" for finishing, but you get the idea.

The slotting case is pretty similar, except for that case, instead of trying to compute the width of cut, we want to use the optimizer to figure out the depth. For example, if we continue with our 1/4" 3 flute, let's say we need to cut a slot 0.300" wide to a depth of 3/4". Our plan is to cut a full slot 0.250" wide down the middle, and then finish it up by cutting the remainder on each side. How deep can we make our full slot passes? Once again, dial up the initial parameters, and this time, hit the "Slot" button. For roughing, the Cut Optimizer tells us we can cut to a depth of 0.3466" before we get too much deflection. Two passes at that depth will get us to 0.6932" deep. That leaves 0.0568" on the bottom for us to finish and 0.0259" on either side for the finish pass. Remember, we're not cutting a full slot for the finish pass, so we treat it just like we did our pocket to figure out the width and depth of cut.

That's all there is to it. To summarize:

1. Decide whether you are slotting or pocketing.

2. When slotting, pick a value for width, and use Cut Optimizer to decide depth.

3. When pocketing, pick a value for depth, and use Cut Optimizer to decide the width.

If you approach the problem this way, you'll maximize your MRR's while minimizing your tool deflection as appropriate for either roughing or finishing. That's a much more optimal approach than the old wet finger in the wind!

Breaking Cutters With Tool Deflection: An Anecdote

Not long ago I got a call to go visit a shop and check out their new Volumill HSM module for GibbsCAM. Being a fan of HSM techniques, I couldn't resist. Volumill is indeed very slick, though we noticed it was leaving some pretty severe nicks in the walls and rough spots in the floor of a pocket. My friend commented that the dealer had suggested Volumill was focused on roughing, and so not having a smooth finish was really not an issue. It certainly did radically increase the MRR's on his job, which I think he said went from 26 minutes to 8 minutes or something similar.

I asked to see the job run, and while he was teeing up the tool in the changer, I remarked that it seemed like he had a lot of stickout for the 3/16" EM he was using. Perhaps tool deflection was the reason for the rough finish? So, not only did we change the stickout, but we went to a more rigid collet chuck while we were at it and fired up the job. We couldn't see much through all the coolant, but we distinctly heard the tool break near the end of the job. Darn!

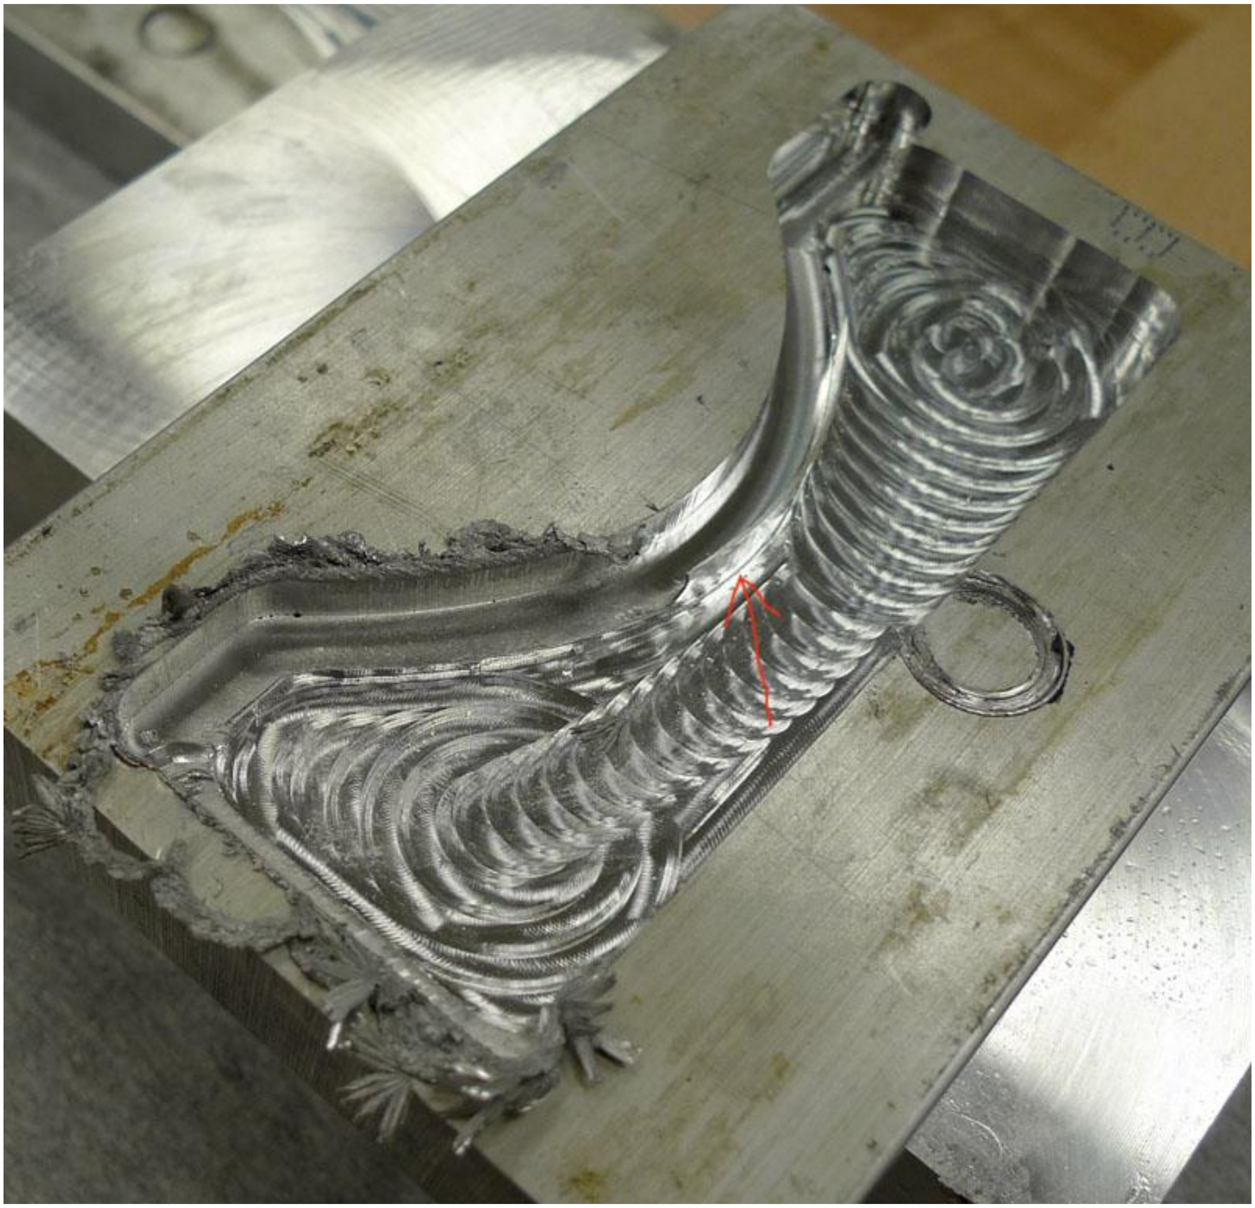

Note the little wall caused by tool deflection (red arrow)...

After pulling the part out of the machine it was pretty easy to see what had happened. The program for the pocket contained two passes. First was a pass using VoluMill to rough out the interior of the pocket. The machinist had then programmed a second pass that was a standard constant stepover that would work from the outside inward to clean up the rough finish of Volumill.

The immediate cause of the failure is the thin wall I've pointed to in the photo with the red arrow. That wall turned a partial stepover cut into a full slot which reduced chip clearance and likely shrouded the cutter from getting enough coolant. The cutter didn't get very far along the wall before built up edge caused the aluminum to start welding onto the cutter. You can see the typical debris from that all along the left side--it looks like mud, but it's aluminum. Most machinists have had this happen before and they know it doesn't go on very long before a tool breaks.

The key question: Why was that little wall there? At first I blamed Volumill because the thickness of the wall seemed on par with some of the other little dips and divots of the toolpath. However, that blame turned out to be misplaced. After I got home with the g-code, I plugged the parameters into G-Wizard and determined that the toolpath had been severely deflecting. The cutter was climb milling, so as we mention above, the deflection would tend to be perpendicular to the direction of cut. In this case, the cutter deflected deeper into the cut and left the wall. If it had deflected away, we never would have seen a problem.

In retrospect, a lot less stepover probably would've improved the performance and finish of the HSM path as well as eliminating the deflection that eventually took out the cutter.

It pays to be aware of and control your tool deflection!

Toolpath Considerations

CNC Milling Feeds and Speeds Cookbook

Introduction to CAM and Feeds and Speeds

The first thing to know is most every CAM package does a pretty poor job calculating feeds and speeds. Most CAM programs have some sort of simplified speeds and feeds calculator built right in, but most of them are painfully over simplified to the point where they don't do much more than your 4 function calculator would let you punch in with the basic Wikipedia formulas.

At CNCCookbook, we have no less than 5 different CAM programs installedr. They were all sent to us to evaluate and write about. Every one of them has cool features of various kinds that we love. But every one of them also has a very primitive notion of feeds and speeds. That's probably a good thing because it gives our G-Wizard business an opportunity to grow, but I wonder how many machinists just assume their CAM program is doing a good job for them on feeds and speeds.

You can tell how sophisticated a speeds and feeds calculator is by the information it takes in and the information it gives out. Take a look, for example, at the G-Wizard's Feeds and Speeds documentation page--there's a lot going on there. Heck, just check out a G-Wizard screen shot:

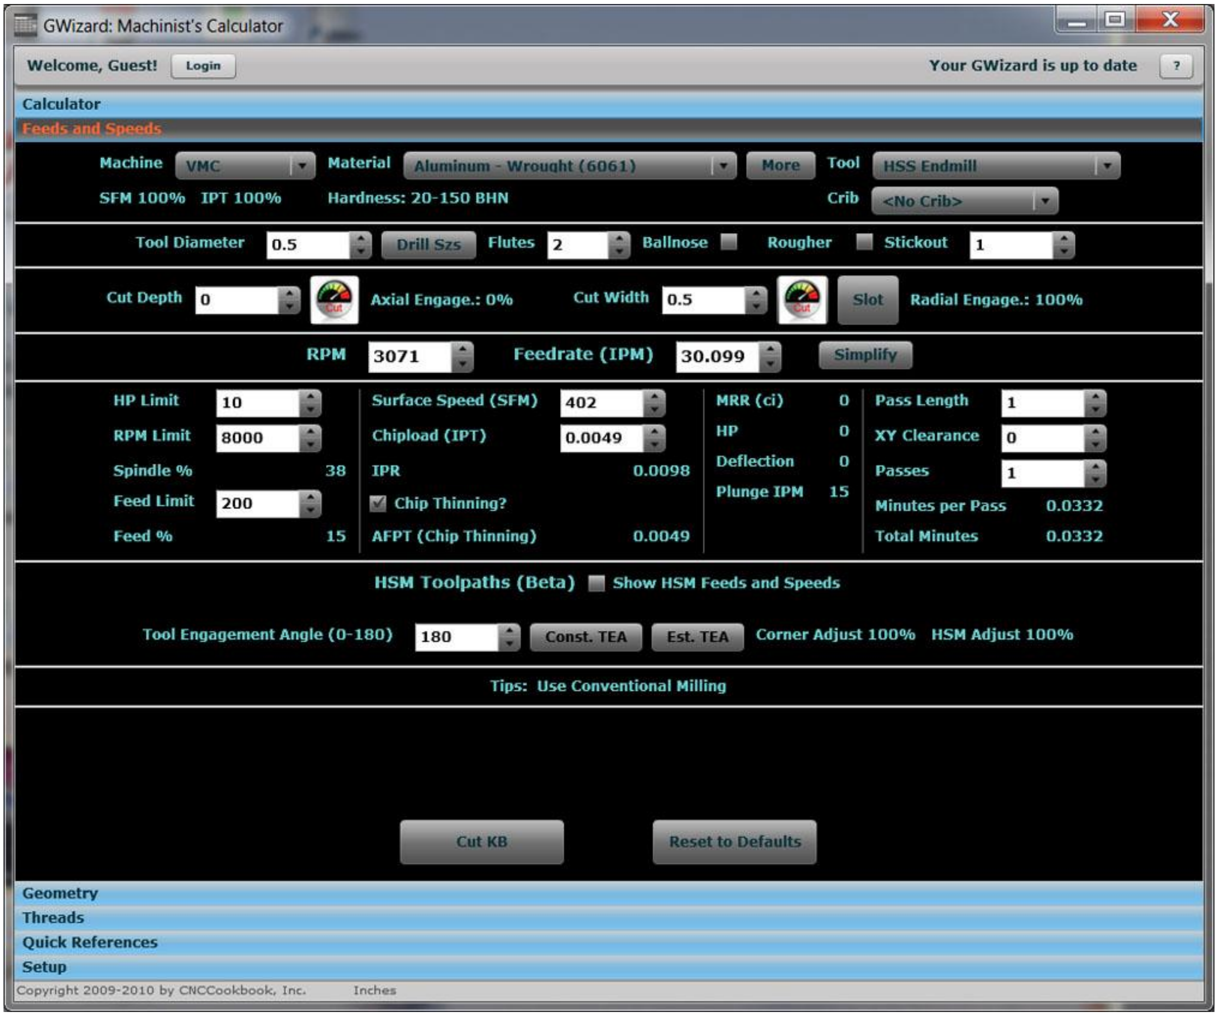

The G-Wizard Feeds and Speeds Calculator...

Now compare that information to what your CAM program is doing. Many of them have a lot of limitations. Here are just a few examples:

- A fixed chipload by tool without regard to material. This may be modified by some "chipload factor" by material, but that isn't how the manufacturer presents the data, so why should you stand on your head to think about it the way the crazy CAM program wants?

- No chip thinning calculations.

- Not much tooling-specific calculations.

- No qualitative rules, like when to use conventional vs climb milling (there are important distinctions there!).

- No analysis of cutter deflection to help optimize your cut widths (stepover) or depth of cut.

- etc.

The short answer is using your CAM program is better than nothing, but not so great. For that reason, we're working on integrating G-Wizard with various CAM programs to make it easier for you. Meanwhile, it's easy enough to use G-Wizard and enter the values it produces into the CAM program. If you've never played with G-Wizard, sign up for our free 30-day trial now to check it out. You'll be happy you did so as our users report it does a better job than even the market leading CAM programs.

How is Your CAM Toolpath Treating Your Tooling?

After you've upgraded your feeds and speeds calculations beyond what is included with the CAM package, the next thing to look at is the effect of various toolpaths. There are an infinite number of toolpaths that will machine the shape you desire, and they are not all equal. For this article, we'll delve into some of the minutiae of toolpaths rather than looking too hard at big picture issues like which toolpath to choose. We'll be talking about things like how the tool enters and exits the workpiece and how it negotiates its way around curves and corners.

Why does it matter? Don't CAM programs all produce pretty good g-code?

Let's digress for a moment to talk about race cars. A fast driver can keep his car right at the edge of its performance envelope around the entire circuit. He knows that if the car is right at that edge, any control input might push it over the edge, resulting in the car moving off the optimal line he wants it on, or in the worst case, resulting in a loss of control. If he goes into a corner a little too fast, he knows that the act of breaking is applying an input that may make matters worse. If he twitches the wheel slightly through a turn, it makes the car hunt for the line and prevents it settling in gracefully. He must apply just the right amount of each input, and he must apply the input smoothly, without any jerking.

Your cutter needs exactly the same treatment for optimal performance. Every change of direction is an input that increases the forces on the cutter. The sharper the change, the greater the forces. It's very important to consider the difference between an impulse and steady force too. Imagine you're applying a force to the end of the tool. You can either apply it by tapping the end of the tool with a hammer using a tap that applies the exact amount of force, or you can steadily press against the tool with that force, perhaps having ramped it up smoothly. Which method do you think will allow you to apply the most force? It's not the hammer!

Let's go through some of the most critical toolpath goings on while milling and see what measures can be taking to keep it smooth and easy. Using these measures will reduce tool wear, improve surface finish, and ultimately allow you to run the cutter faster, further, or both.

Cut Entry and Exit

Entry into a cut is the worst case of interrupted cut, and causes the most wear and tear on the cutter. Ideally, you want to roll the cutter smoothly into the material with an arc motion rather than a straight line. If you must move in a straight line, avoid a head-on collision--try to enter almost tangentially. If you can't roll gently into a cut, try reducing the feedrate to half until the cutter is fully engaged in the cut.

Here's a great video on how to "roll into a cut", courtesy of Sandvik, with a hat tip to Don "Milacron" at the PM Boards:

Another great video from Sandvik's "My Yellow Coat" site...

I'm fascinated by these geometric effects. Chip thinning is another. Isn't it interesting how Mother Nature tends to like circles better than straight lines? Chip thinning, rolling into a cut, and the trochoidal paths of high speed maching are all about the behavior of circles as we try to use them (in the form of rotating cutters) to cut the straight lines that we humans are more comfortable with. Circles are more gentle and natural in these applications. Here is another look at the geometry:

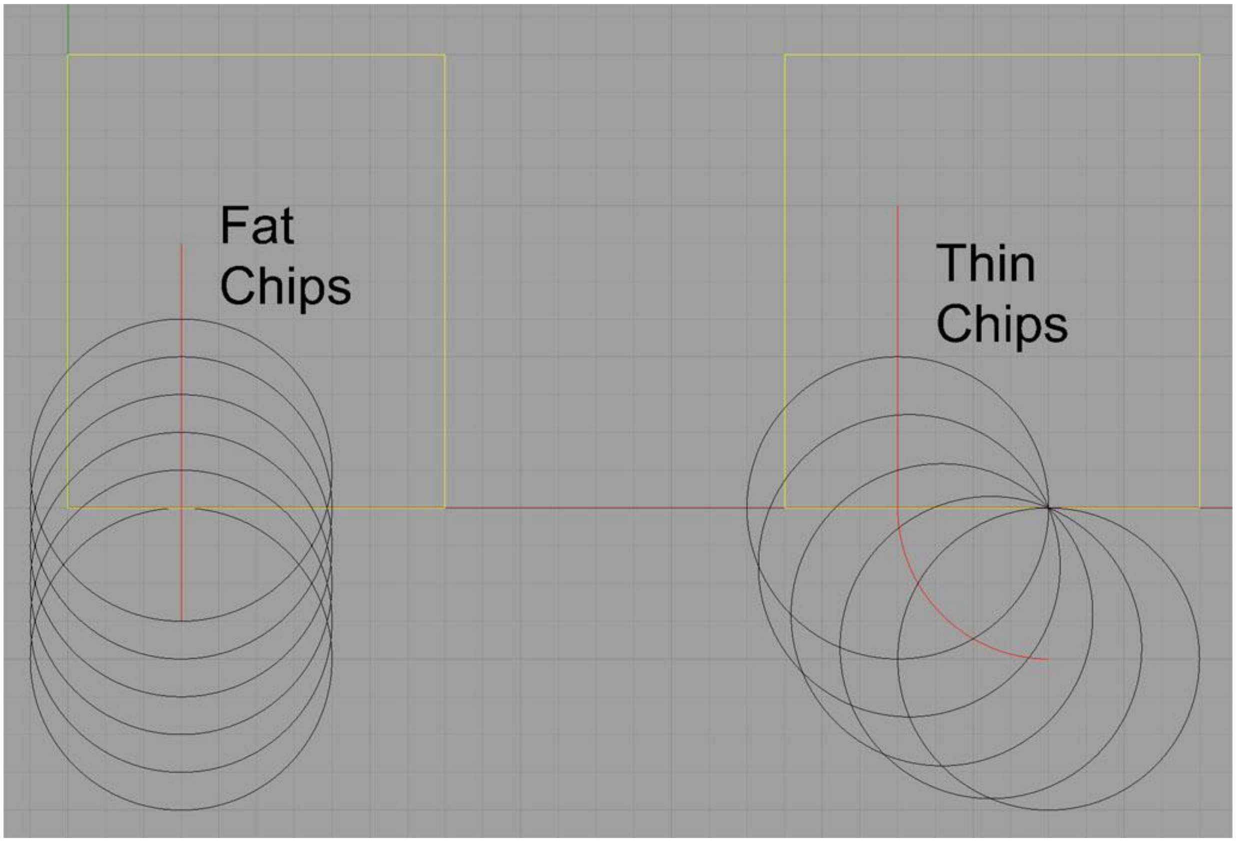

Not the chip shape to the right of the feed line (the red line is the path the cutter follows): thin chips on exit are better!

To execute an entry like this means starting the cutter out one radius to the right of the original starting location and then rolling it in along a path that is an arc with the same radius as the cutter. The folks on PM report that this works as well for endmills as it does for the face mills Sandvik shows in their video. In fact, they say it really helps improve cutter life on materials like Stainless Steel.

Stay in the Workpiece and Avoid Straddling Slots

Try to keep the cutter engaged as much as possible. Every time it leaves the workpiece it has to do another entry, which is an interrupted cut of the worst kind. Your cutter and your workpiece will be happier if you can keep them engaged.

You can really make a cutter's life miserable by travelling over slots, holes, and other openings because you're forcing the cutter to enter the workpiece over and over again each revolution. Try to arrange your toolpath so that instead of straddling a slot, it cuts a pass on either side of the slot.

Anytime you must machine over interruptions, you can improve the situation by reducing the feedrate up to 50% depending on the toughness of the material.

Plunge, Ramp, or Helix (Circular Ramping)?

When preparing to pocket, you have a choice of plunging, ramping, or helixing to get the cutter into the material. Of course you could also use an insertable drill to quickly open a cavity to start from, but you'll have to way the cost of the tool change to see whether it's worth it. Not counting the tool change, a drill is the fastest way to make a hole of all these methods.

In order of preference, the Helix is the easiest on the tooling, followed by ramping, followed by plunging. Avoid the latter where possible.

Travel in the Cut

Stay off-center when Face Milling

If your cutter is wider than the cut, don't position the cutter right on the centerline--keep it on one side or the other. This loads one side of the cutter and makes for a more stable cut since the cutting forces will tend not to move around the way they do when the cutter is centered. Ideally, position on the side (based on spindle revolution) that results in climb milling where the chip starts large and gets smaller. This is left of center for a clockwise spindle rotation. A centered cutter has a large average chip thickness throughout, but an offset cutter can start chips large and have them get smaller, which is more favorable geometry as we discussed under Cut Entry and Exit.

Going Around a Corner

Going around corners is really tough on tools because the tool engagement angle, and hence cutting forces, can radically increase in the blink of an eye. This diagram helps explain:

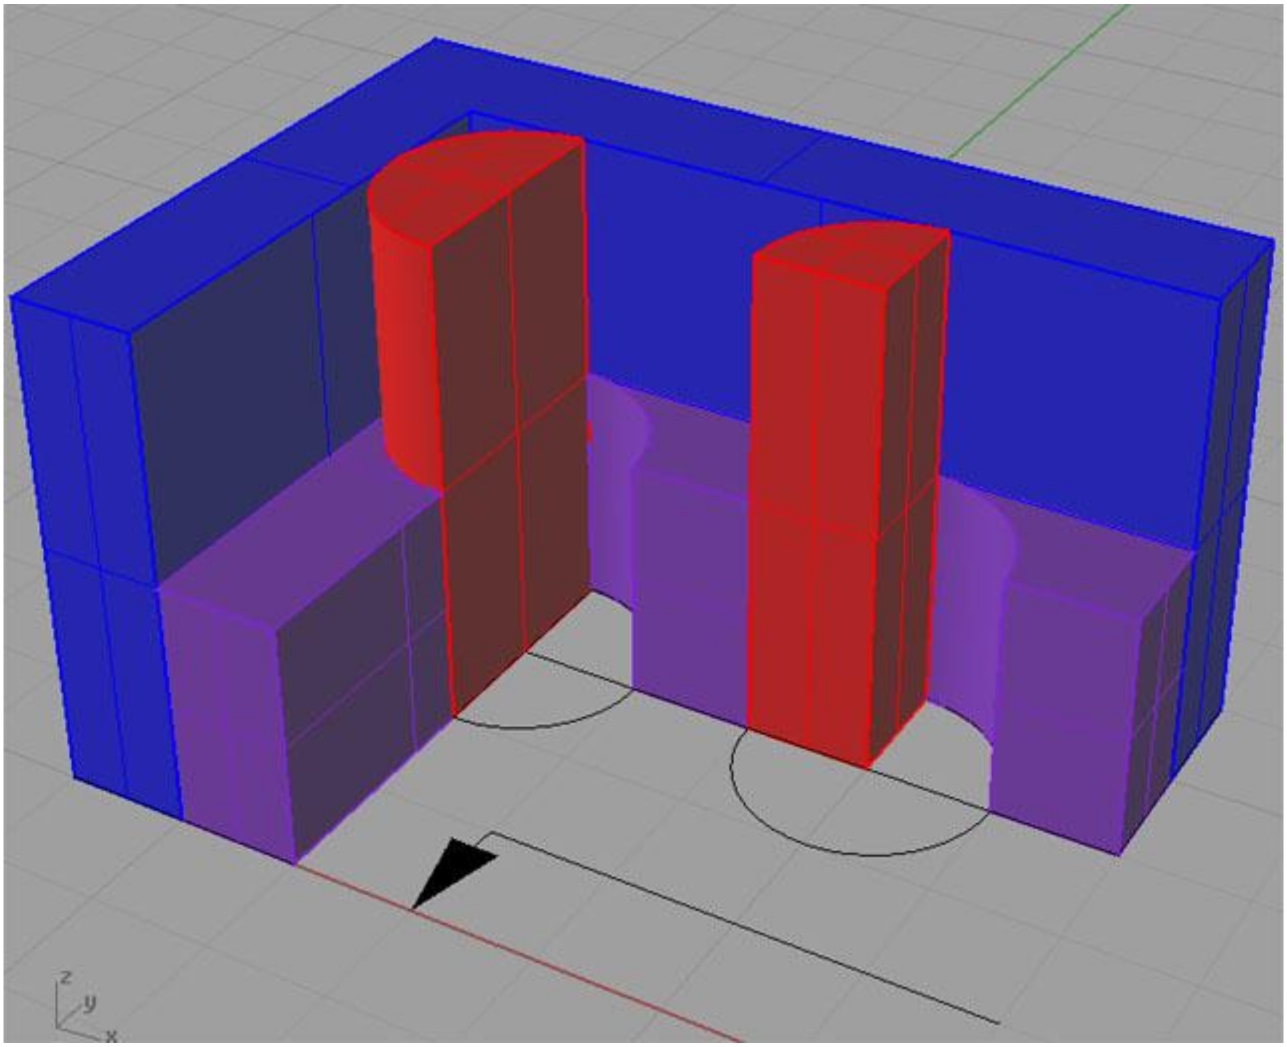

Entering a corner doubles the cutter engagement...

As the cutter enters the 90 degree corner, the Tool Engagement Angle, which is the angle of the tool's circumference that is cutting, doubles. Cutting forces spike commensurately and tool gets hammered. This tells you why special HSM tool paths that control the engagement angle can run so much faster. A manufacturer's normal feeds and speeds have to assume there will be corners. As such they're dialed down to take the huge spikes in cutting force. If the manufacturer could be assured the tool would only be used in a straight line with no corners and no cutting force spikes, they could offer much higher speeds and feeds.

CAM programs provide several methods of avoiding problems in corners.

Adjusting Your Feedrate

Every machinist has had the experience of hearing the machine "squeak" out a little chatter as the tool goes around the corner. Most have broken tools in corners when they pushed too hard. At some point, machinists learn to "ride" the feedrate override. That is, they would dial back the feed by twisting the FRO dial as the cutter came to a corner. Eventually, utilities and CAM programs became available that would automatically slow down the feedrate in corners to keep the cutting forces constant.

This is not a bad solution compared to the alternative of plowing into the corner full speed. It will let you tune up your programs to run with higher feeds and speeds. However, pretty soon better ways were developed. The disadvantage of simply adjusting the feedrate is that machines are limited in how fast they can accelerate and decelerate. Inevitably, this means g-code has to slow the tool sooner and longer than is optimal.

Trochoidal Paths and "Peeling" Corners

What if instead of slowing down for corners you simply arranged the toolpath so there weren't any corners? This is the secret of all modern HSM toolpaths. The first approaches tried involved techniques called "Trochoidal Milling" and "Corner Peeling."

Corner Peeling invovles taking a series of arc moves to slice down gradually into the corner instead of plowing into it:

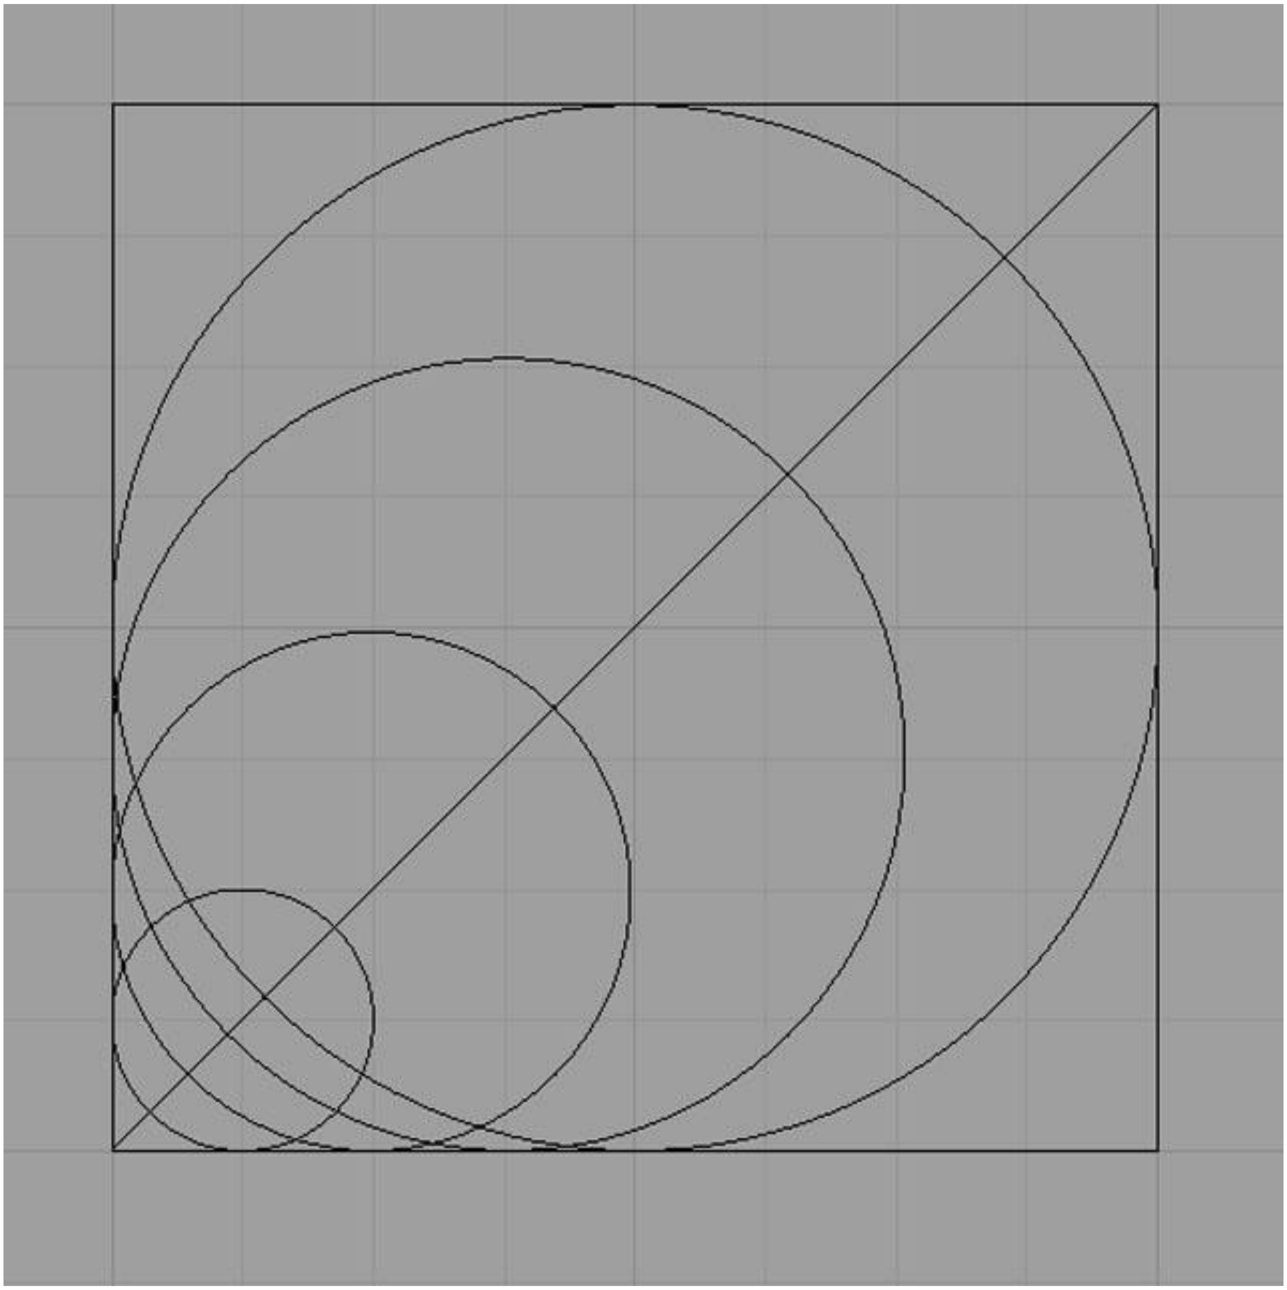

Corner Peeling involves using toolpath arcs to clean out a corner...

Trochoidal Milling was first introduced to clear slots. Imagine rolling a disc in a straight line along its edge, with a pencil stuck in the side of the disc. The tip of the pencil describes a looping figure called a "Trochoid" as it rolls. Now imagine cutting a slot with an endmill. Normally, the endmill just plows along the slow, full width. But suppose you used a smaller endmill, and suppose you made it loop around and around so that the loops were the full slot width.

That's a Trochoidal Toolpath, and you can run it a lot faster than you can run the full width endmill.

Constant Engagement Angles

The problem with Peeling and Trochoids is they're special case "hacks" for corners rather than general purpose ways of controlling tool engagement angle. Eventually, CAM vendors solved the problem of how to generally create toolpaths that keep constant tool engagement angles. The productivity when roughing with such paths is huge. Here is a typical path created by GibbsCAM's Volumill for a pocket:

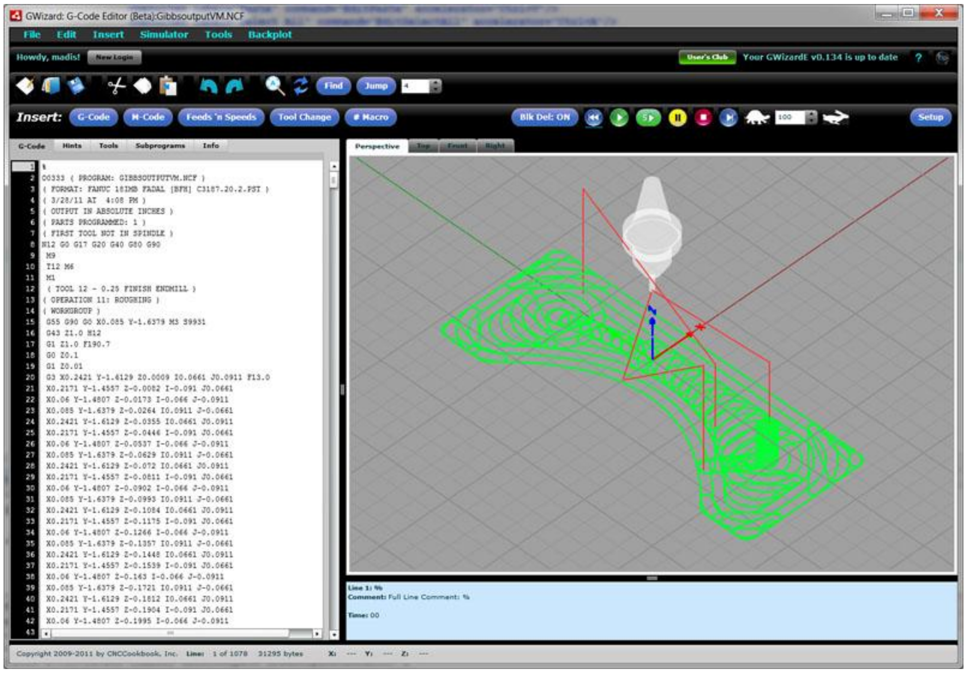

Constant Engagement Angle Path as displayed by G-Wizard Editor...

We'll have more to say about these sorts of toolpaths in the article devoted to HSM (High Speed Machining).

CAD/CAM Artifacts NURBS and Curve Fitting

One last thing I wanted to touch on related to the quality of g-code and toolpaths produced by CAM programs. These programs are working at a bit of a disadvantage from some perspectives. The reason is that they're typically building on layers of approximations, each of which has the potential to introduce errors and roughness in the final result. Let's start with the CAD model we're getting the geometry for our toolpaths from. Drawings and models consisting strictly of straight lines and sharp corners are relatively easy to create accurately. But when dealing with curves, all sorts of issues can crop up:

- Are we approximating the curves as a series of many line segments?

- Can we go beyond simple arcs to use more complex representations of the curves?

- Can we represent curves using NURBS (Non-Uniform Rational B-Splines)?

That's the order from lowest to highest quality representation of curve geometry. Suffice it to say that the further down the list your CAD operates, the better potential it has for delivering smooth accurate geometry.

Next, we need to consider what the CAM program is prepared to consume. For example, can it take in your CAD program's native files or did you have to convert to some other format? Consider STL, for example. STL is a common 3D format that represents 3D objects with thousands of small triangular facets. Clearly there is the potential for such an approximation to diverge from the smooth flowing curves you had in mind when you created the model. It's always best if your CAM program can import the native geometry format of your CAD. BTW, just because it happens to be integrated with the CAD program doesn't necessarily mean that's the case, although there is a high likelihood it is. You also have to consider whether the internal representation of the geometry used by the CAM program reverts to something less sophisticated than your CAD produces. The latter is something you're unlikely to ever find out for sure, but you can suspect it if you notice the CAM program acting like it ignores or approximates some of the geometry.

The last step is your CNC Controller. Basic g-code has line segments (G00 and G01 moves) and arcs (G02 and G03). That's it. It's only the higher end controllers that can directly operate on NURBS. So even after you've gone to great trouble to avoid losing fidelity throughout the CAD to CAM chain, your CAM program must still break things down to something your machine's controller can deal with.

Complex "curvy" toolpaths for tasks like 3D profiling can wind up with lots of operations as the curves are simulated as many very small line segments (and possibly arcs). In fact, the volume of data can overwhelm the controller, whether because it can't store it all and requires drip feeding, or because it can't process the moves fast enough to maintain the commanded feedrates. Your CAM program may also output changes of direction that are too frequent and sudden for the acceleration capabilities of your machine.

There are a host of ways that information, accuracy, and surface finish can be compromised on the road from drawing to finished part. How do you avoid these problems?

First, you can fiddle with the settings of your CAD and CAM. Settings like tolerance are pretty important for these situations.

Second, some CAM and other utility software has the ability to go back into the g-code and tune it up to run better. This is done using a process called "curve" or "arc" fitting. In essence, the idea is to convert a series of moves into a single smooth arc, provided the arc is within a specified tolerance of the original series of moves. Doing this can dramatically reduce the size of some g-code programs as well as making them run much better on machines that have trouble keeping up with all the little moves.

Another Take on Artifacts: Facets from STL

It's good to have a lot of choices for file formats and to know what the different formats strengths and weaknesses are. Depending on the job, some are better than others. Consider this part which is showing marked faceting where there should be smooth curves:

Faceting shows the g-code used lines where arcs might have been better...

The photo is from a CNCZone thread. The machinist had recently switched from using the DXF format for CAD to STL files. That's when the faceting began. The trouble is the STL file format has no way to represent a smooth curve or even arcs. It converts everything into triangles:

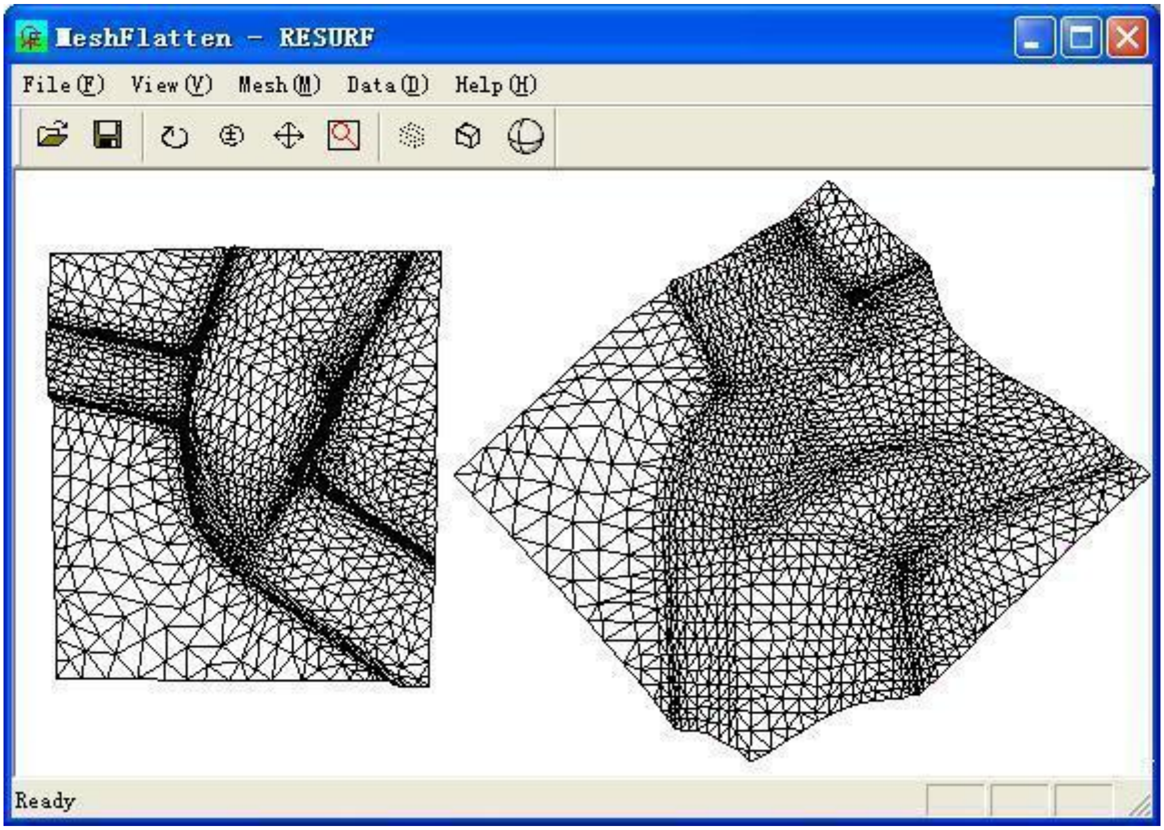

An STL mesh from MeshFlatten software...

There's not necessarily anything wrong with this, except that you have to be aware of it. Your CAD and CAM software will allow you to specify the tolerances--how closely must the mesh match the idealized 3D part? With a small enough tolerance, the facets will disappear. The flip side is unless your CAM software is pretty clever, you'll be forcing your machine to make hundreds or even thousands of tiny little straightline moves for those facets. A clever CAM will spot the opportunity to use arcs to simplify the code. You can even purchase software that goes back into your g-code and looks for opportunities to substitute arcs.

The thing is, with this particular part, the desired result is very amenable to 2 1/2D programming. STL is more of a 3D format. Why go crazy profiling a bunch of triangular facets in 3D when you can just cut a simple profile with arcs and line segments in 2 1/2D?

How Fast Can Your Machine Controller Go?

As we've seen, a lot of changes can be made that result in smoother cutting. You want the smoothest toolpaths possible. You want fewer operations, for example, don't use lots of line segments where a single smooth arc will fit the tolerances required. Some of this is just good cutter dynamics, and some of it is trying to overcome the controller's shortcomings.

Older controllers can't necessarily process as many blocks per second as is needed to follow all those little line segments. Let's say you're approximating an arc with a series of short segments. The arc has a radius of 1/2" and your tolerance on the job is 0.001". Using G-Wizard's Chord calculator, we can see that each segment is about 0.001" long and covers an angle of 0.1146. Let's say we need to do 1/4 of a circle for a corner, which is 90 degrees, or 785 of these little segments.

Now let's assume we're cutting aluminum with a 3 flute 1/4" cutter at 40 IPM. Those little segments are supposed to be cut in 0.01 seconds. Put another way, your controller has to execute 78,000 blocks/second to keep up. Typical modern controllers can run 2000 - 4000 blocks per second. What happens is your cutting slows way down in the corner because the controller can't command the servos fast enough to keep up with the program.

This has been a somewhat extreme example, but programs call on their controllers to go faster than they're able more often than you'd think, especially if you're running an older controller.

World's Fastest Controller?

MTI (Miceli Technologies) claims to make the world's fastest CNC controller. It runs at 50,000 blocks per second. Because of its greater ability to run programs at the speeds intended, Miceli claims they will run jobs 25-75% faster, and in fact they guarantee your machine will run the Mercedes Test Geometry twice as fast as it did before the Miceli control was installed.

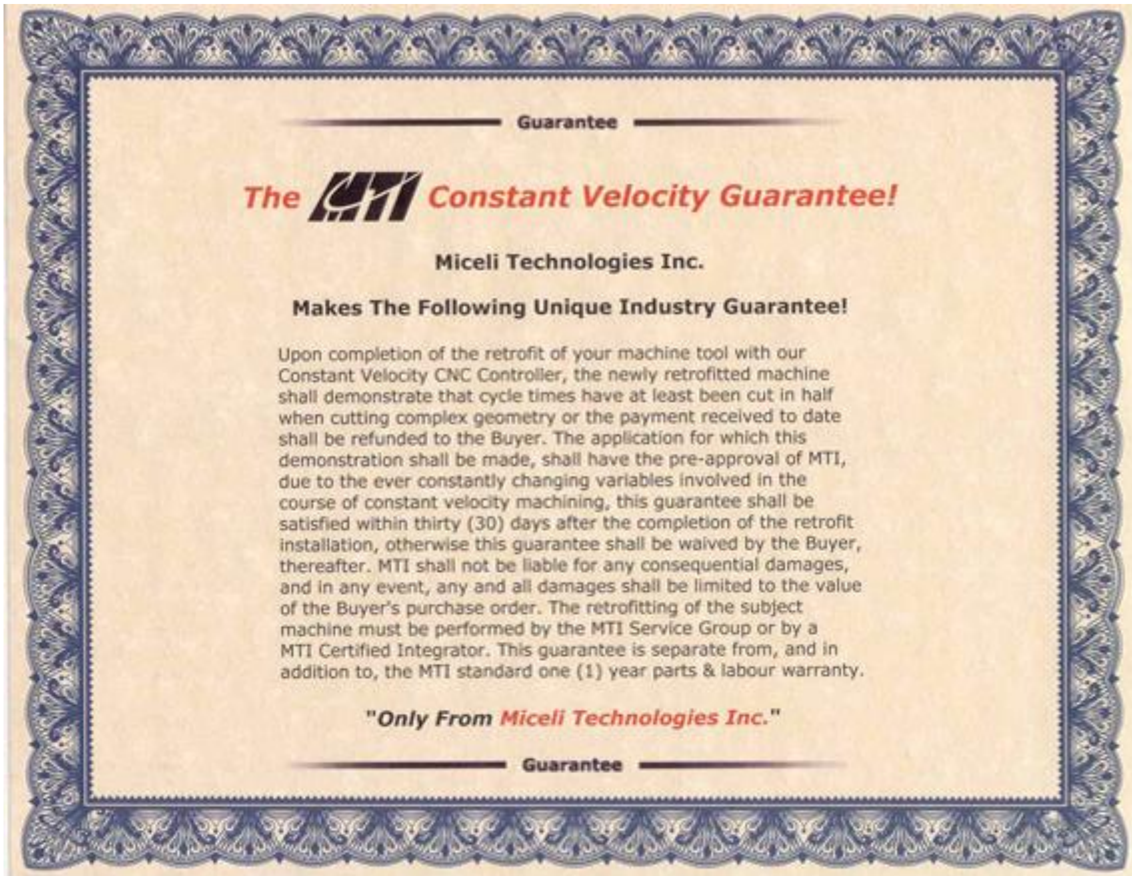

Guarantee

The MTI Constant Velocity Guarantee!

Miceli Technologies Inc.

Makes The Following Unique Industry Guarantee!

Upon completion of the retrofit of your machine tool with our Constant Velocity CNC Controller, the newly retrofitted machine shall demonstrate that cycle times have at least been cut in half when cutting complex geometry or the payment received to date shall be refunded to the Buyer. The application for which this demonstration shall be made, shall have the pre-approval of MTI, due to the ever constantly changing variables involved in the course of constant velocity machining, this guarantee shall be satisfied within thirty (30) days after the completion of the retrofit installation, otherwise this guarantee shall be waived by the Buyer, thereafter. MTI shall not be liable for any consequential damages, and in any event, any and all damages shall be limited to the value of the Buyer's purchase order. The retrofitting of the subject machine must be performed by the MTI Service Group or by a MTI Certified Integrator. This guarantee is separate from, and in addition to, the MTI standard one (1) year parts & labour warranty.

"Only From Miceli Technologies Inc."

Guarantee

Wait!

There's much more to the CNCCookbook Feeds and Speeds Cookbook on our web site here:

http://www.cnccookbook.com/CCCNCMillFeedsSpeeds.htm

Come check it out!