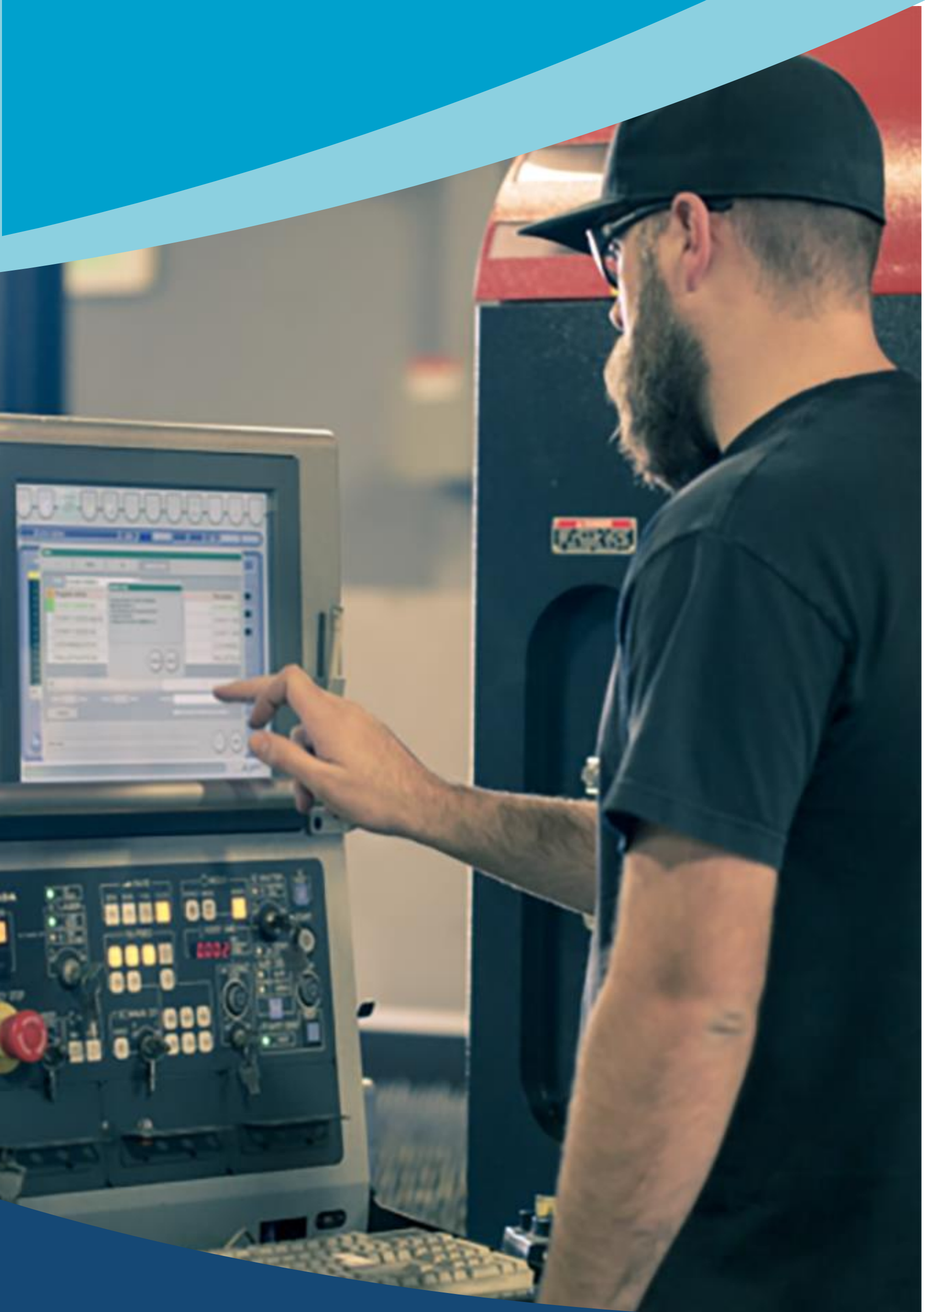

A FEW WORDS ON INTELLIGENT ROUGHING

An article by Louis Martineau, Consultant
Manufacturing Solutions at SolidCAD
(Canada's reseller of HSM CAM for Inventor
and Fusion 360)

INTRODUCTION

“Adaptive Clearing” is the name for Autodesk HSM CAM’s intelligent roughing routines. Introduced in 2007 under the HSMWorks banner, Adaptive Clearing is the archetype of what has become nothing short of a revolution in the world of CNC machining.

Prior to the advent of intelligent roughing algorithms, CAM software merely speeded up calculations that could otherwise be performed by hand. With intelligent roughing, CAM software truly leverages the power of modern computing to create high-performance tool paths that are more efficient than anything a CNC programmer could ever write. Put quite simply, if you aren’t using intelligent roughing to achieve near-net shape in a single, tool-sparring, operation, you are leaving money on the table!

Intelligent roughing was immediately recognized as a breakthrough technology and has been widely imitated since its introduction. Today, it is offered in some form or other on nearly all serious CAM platforms. However, over the years, Adaptive Clearing has been steadily refined to remain the gold-standard in performance and usability. Adaptive Clearing is also the most accessible, positioned by Autodesk not as an add-on module, but as the core, go-to strategy for its very attractively-priced HSM CAM.

While Adaptive Clearing was introduced more than a decade ago, there is a pretty good chance that you’ve never heard of intelligent roughing, let alone taken it for a spin. That’s because machine shops are often slow to accept change and tend to be set in their methods. But if you’re reading this text, you are likely looking for a competitive edge. Read on and I’ll explain how this technology works and why the return on investment is instantaneous.

Surrounding intelligent roughing are some pretty wild-sounding claims of decreased cycle times and prolonged tool life. This hype is mostly true. However, calculating the actual payoff can be a little tricky. More on that below. For now, if we wish to understand when and where these strategies really shine, then let's start by looking at how those benefits are supposedly derived.

Right off the bat, it needs to be said that intelligent roughing applies principles of high-speed machining (HSM) to decrease cycle time and prolong tool life. If you associate the words “high-speed machining” with a large investment in tooling and you are considering tuning out, know that it is absolutely possible to apply these principles and derive significant benefits without faster spindles, faster drives, and shrink-fit toolholders. Indeed, simply ramping up speeds and feeds within the capabilities of a traditional machining center can yield a surprisingly big payoff. This article is written in that context.

Let's first address the "roughing" part of intelligent roughing. This refers to trochoidal passes with a small and tightly-controlled maximal radial engagement. The beauty of such controlled trochoidal milling is that the tool is not allowed to make sharp turns into the stock, and therefore not given the opportunity to bury itself in a concave inner corner. Instead, independent of feeds and speeds. Why is this important? Because when a rectangular pocket is cleared with a familiar constant-overlap offset-roughing pocket strategy, tool engagement spikes every time the tool makes a 90-degree turn; it is precisely during such a spike (*i.e.,* a load spike) that tool breakage tends to occur. Indeed, constant overlap does not equate to constant loading. Therefore, the use of "load-protected" trochoidal tool paths for clearing pockets and slots increases tool life in the sense that mean-time between unscheduled tool failure is extended all other things being equal. In other words, for a given material removal rate (MRR), controlled trochoidal milling will be easier on tools than a traditional approach based on a larger radial engagement. This is benefit number one of this technology, and it happens to be independent of feeds and speeds.

Two other factors contribute to prolonging tool life when using intelligent roughing strategies. First, a small radial engagement translates into a lower duty cycle for the tool. In other words, with each revolution, each cutting edge spends less fractional time inside the stock and more fractional time cooling in the air outside the stock. This translates directly into a cooler-running tool. In fact, when using this approach, dry machining (with chip clearing handled by high-pressure compressed air rather than coolant) is viable in many materials. Secondly, a small radial engagement is normally balanced out by a large axial engagement (*i.e.,* peripheral, edge, shoulder, or side milling) so as to keep MRR at least comparable with a more traditional combination of larger radial engagement and lower axial engagement. If radial engagement is small enough and tool diameter sufficiently large to minimize deflection (end mills 3/8" diameter and up), then axial engagement can be increased to some 2 diameters. Cutting with more flute length means that tool wear is not focalized to the tool tip but is instead more uniformly distributed over the length of the tool.

Putting numbers onto the increased tool life resulting from controlled trochoidal milling with a small radial engagement and a large axial engagement is rather difficult and best addressed by empirical data collection and statistical analysis. Let's simply agree that prevention of load spikes, a cooler-running tool, and more uniform tool wear all translate into at least some economy.

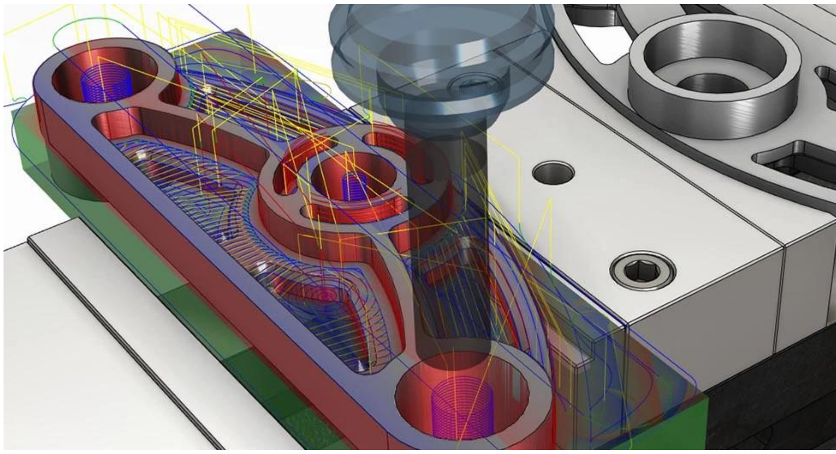

Let's now address the claim of reduced cycle time. The first benefit comes from the increase in feed rate and material removal rate (MRR) that is possible with controlled trochoidal milling. There are several contributors to increased feed rate.

First, if we know for certain that the tool path will not make any sharp turns into the stock and will not produce load spikes, then we can safely exceed the tool manufacturers' general recommendations, which are formulated to be conservative enough to handle the load spikes of traditional tool paths. In the same vein, general recommendations are based on tool deflection (that enemy of solid carbide end mills) predicted to result from load applied primarily on the tool tip. However, when using a strategy of small radial engagement / high axial engagement, load is distributed over a greater length of the tool, resulting in less tool deflection; again, recommendations can be safely exceeded.

Thirdly, as the cutting action produces less heat when using a small radial engagement, cutting speed (*i.e.,* Vc) can be increased significantly. Moreover, beyond a threshold speed (which varies across materials), heat actually decreases as Vc increases; this paradoxical phenomenon is a tenet of HSM and is particularly relevant to aluminum. Lastly, although it does not contribute to increasing MRR over a traditional approach of larger radial engagement and smaller axial engagement, it needs be mentioned that Vc must be increased to prevent rubbing and premature tool wear that result from the chip thinning effect of a small radial engagement; a straightforward equation is used to calculate the compensation required for any given radial engagement below 50%. It should be noted that all these contributors to increased feed rate are additive. Therefore, feed rates appropriate to controlled trochoidal milling should be expected to be substantially higher than those typical of a more traditional approach.

Now, does this increased feed rate really translate into increased MRR? That's a somewhat loaded question. To answer it, let's consider a quantitative example. <u>A good starting point for intelligent roughing strategies is a radial engagement (WOC) of 15% of diameter and an axial engagement (DOC) of 2 times tool diameter</u> (tool deflection and part geometry permitting). For a 5/8” endmill with 2 or more diameters of cutting length, this means .095" WOC and 1.250" DOC. This recipe does not require a terribly fast machining center (high RPM or high feed rates) or balanced tool holders to pay dividends. If this is a 3-flute solid carbide endmill sticking-out 1.5” from the tool holder, and we are cutting a deep pocket in some rigidly-mounted, hard 7075 aluminum with a 10,000 RPM / 15 HP spindle, then good speed and feed would be on the order of 7700 RPM (Vc ~1250 SFM) and 220 in/min for an impressive MRR of around 26 in^3/min while in the cut, and roughly 12 HP. (For reference, that's on par with a 3.0", 6-insert face mill cutting a .100"-deep, 2.0"-wide swath at 130 in/min.)

In scenario 2, the same tool is programmed with the familiar constant-overlap offset-roughing strategy using a more traditional radial engagement of 50% and axial engagement of 1 diameter. All other parameters being equal, a good speed would now be 5200 RPM (Vc ~850 SFM), and a good feed would now be around 80 in/min. This gives an MRR, while in the cut, of a little over half that of scenario 1. (For those keeping score, this 50% overlap tool path also produces *full-slotting* load spikes with every 90 degree direction change - again, constant overlap is not to be confused with constant engagement!) Increased MRR is therefore the first factor contributing to reduced cycle time and we can realistically expect some 50% higher *peak* MRR with intelligent roughing. It is noteworthy that in a softer material or with a smaller diameter tool, intelligent roughing would have maxed out our spindle speed and therefore lost some of its MRR advantage. Indeed, in 6061-T6, we would have had to “settle” for 10,000 RPM and 360 in/min (instead of the full 11,500

RPM and 400+ in/min), and the advantage over a traditional approach would only have been on the order of 25%.

That advantage in MRR, however, does not always translate perfectly into reduced cycle time: part geometry has an enormous impact on the result of any machining strategy showdown. For example, if only a rectangular pocket is being roughed out, then the mostly trochoidal tool path of the intelligent roughing strategy will require several air-cutting repositioning moves to finish cleaning out the four corners of that pocket, whereas the constant-overlap tool path will only require repositioning for its successive Z-passes. This will in part cancel out the increased peak MRR, reducing the all-important average MRR. This loss in productivity from the repositioning moves of trochoidal milling is at its worst when roughing slots. In the case of simple pockets and slots, it is probably fair to say, then, that the real advantage of intelligent roughing is increased tool life, rather than reduced cycle time. (Note that the negative impact of repositioning moves has greatly been attenuated with the advent of “both ways” Adaptive Clearing, wherein cutting in the conventional milling direction is allowed between main passes in climb milling – issues with surface finish are largely moot given that this is a roughing strategy with some stock-to-leave.) Moreover, if the pocket is not deep enough to run intelligent roughing at an axial engagement of 2 diameters, then MRR in the cut becomes less favorable still. Indeed, as axial engagement decreases, radial engagement should be increased proportionally to maintain productivity, but at the cost of feed rate and duty ratio; taking this argument to its logical extreme, then at some point, all advantage will be lost except for reduced tool breakage from a load-protected tool path.

That being said, controlled trochoidal milling can beat out a traditional roughing strategy in terms of average MRR anytime that it is used to rough out the outside contour of a part in addition to other part features, and that at least some of the cutting capitalizes on the considerably higher peak MRR. This now brings us to the notion of clearing an entire part with a single intelligent roughing operation.

The most dramatic reduction in cycle time can be attributed to the "intelligent" side of intelligent roughing. In Autodesk language, this refers more specifically to 3D-Adaptive Clearing; whereas roughing of the simple pocket described above would have been addressed using the 2D-Adaptive Clearing strategy, the roughing of multiple features or even of the entire part in a single operation would be addressed with the 3D-Adaptive Clearing strategy. Still using controlled trochoidal milling with a small and controlled maximal radial engagement and a large axial engagement, 3D-Adaptive Clearing differs from its 2-D counterpart in that it uses finer step-downs as well to clear floors and other flat areas that are at intermediate Z-heights, and to create staircases on 3-D ramp, radius, or organic features. Moreover, work-holding elements identified as such are magically avoided by the tool path.

This intelligence of 3D-Adaptive Clearing results in near-net shape from a single roughing operation. Depending on the specified size of the fine step-downs and the specified amount of stock-to-leave, semi-finishing can be skipped altogether, and final finishing of features can be undertaken immediately following roughing. Therein lies the most tangible advantage of intelligent roughing. And, as is the case with the load-protected tool path, this advantage is independent of feeds and speeds. As before, a good starting point is a maximal radial engagement (the Optimal Load parameter) of 15% of dimeter and a roughing step-down of 2 diameters. Fine step-downs can be set to something between .050" and .100" (depending on slope), so as to produce staircases on sloped features that can be finished in a single step-down with a radius end mill (using a 3D strategy like 3D-parallel) without any additional prep.

Intelligent-roughing to near-net shape equals savings not only in cycle time, but also in programming time. Indeed, programming a 3D-Adaptive Clearing routine can be achieved in just a few clicks if the tool path need not be limited to a specific boundary (although calculation time can be fairly long). This is an important consideration if programming time, normally much cheaper than machining time, happens to be in short supply in your shop.

We hope that this general discussion of Autodesk's intelligent roughing functions will convince you to give them a spin. Should you wish to do so, then continue reading for a brief discussion of some of the pitfalls associated with these strategies.

AVOIDING SOME PITFALLS OF ADAPTIVE CLEARING

1

3D-Adaptive Clearing can be a nightmare for your set-up man. Indeed, with thousands of blocks streaming faster than the eye can see and the tool dancing all over the part (and brushing up to clamps and other work-holding!) to create a near-net shape, there is almost no longer any point to running a reduced feed with a finger on the button when proofing a program. Instead, if the program passed all subsequent proofing steps and it is now time to cut a first part, then either cut a block of machinable wax or simply have faith and put that $125 end mill in harm's way. Over time, your team will hopefully learn to trust the software. Unfortunately, this is a case of "adapt or die".

2

Adaptive Clearing generates thousands upon thousands of blocks. If your controller is in any way light on program memory, don't forget to activate smoothing. Smoothing to .001" will cut memory requirements by more than half compared to default settings. Since Adaptive Clearing is a roughing/semi-finishing strategy, you will always be leaving extra stock anyway, and therefore smoothing even to a few thousands will have no negative impact. Note that even with smoothing activated, a typical Adaptive Clearing tool path on a small part will still run 100-150 kB (add 30-40 kB if you don't disable line-numbering). Therefore, if your controller doesn't have at least 256 kB program memory, you might have to split programs in two. Let's call 256 kb the practical minimum for getting the most out of Adaptive Clearing, and 512 kB a more comfortable amount. The reason for even bringing this up is that, whereas the cost of memory on a new machine is next to nothing and memory is now measured in GB, the cost of a memory upgrade can still be astronomical for an older machine.

In the same vein, smoothing can help with a controller that is on the slow side, with insufficient look-ahead to predict acceleration and deceleration requirements of fast-changing tool paths. Again, given that Adaptive Clearing should be used with some radial stock-to-leave allowance, a small deviation from the programmed tool path may be negligible.

If really desperate, smoothing can be combined with a slightly larger max. radial engagement (on the order of 25-30%) in order to further reduce the number of lines.

3

Be careful with stock size variability. Since such small radial passes are being taken, leaving an extra 1/8" at the saw can easily result in a "collision" at the point of initial entry into the stock: the tool may be feeding too fast to handle two or three times the expected radial load. To prevent this, one can program a safety margin in the form of extra stock during part set-up. The downside to this is some air cutting, and effectively no lead-in at the actual entry point, which can be almost as bad as the situation we are trying to avoid. The only real solutions are either to: 1) control the feed manually until the tool has cut a little bit of stock off all sides before letting it run free; or 2) program a traditional contour pass around the stock to ensure accurate dimensions before releasing the beast. Perhaps Autodesk will further refine 3D-Adaptive Clearing so that it can automatically reduce feed for the first pass around the stock when working from the outside in.

4

Keep in mind that tool life should not be considered independently of cycle time --longer intervals between tool replacements also contribute to decreasing cycle time, especially if we are talking about time between unscheduled tool failures (*i.e.,* tool breakage). By focussing too much on single part run-time, this very important benefit is often overlooked.

5

Be careful with thin-walled parts (or, more precisely, parts featuring walls with a large height to width ratio) that would normally be machined using an incremental inside-outside waterline approach, rather than very large step-downs. One approach is to specify a larger radial stock-to-leave setting for the 3D-Adaptive Clearing tool path. This is flawed, however, since the extra radial stock will be applied indiscriminately to all features and may therefore cancel out the benefit of intelligent roughing to near-net shape. We offer two workarounds for this situation. The first is to split a 3D-Adaptative Clearing routine into discrete routines for each full step-down through the Heights control, and then to finish all thin walls to final thickness between each of the full step-downs; in this way, the effective height-to-width ratio is decreased. For example, instead of leaving .100" of stock on either side of a thin, 1.250" tall, unsupported flange that is roughed out full height, you can leave only .050" if the flange is roughed and immediately finished in two roughing increments of .625" (remembering to increase max. radial engagement to make up for reduced axial engagement); this will leave .050" radial stock on all other structures, an amount more amenable to finishing in a single pass than .100".

The second workaround is to use tooling lugs to support walls. The idea is to position the lug in a strategic location for support and sufficiently far from the wall to allow the finishing tool to pass, yet just far enough that the 3-D Adaptive Clearing routine will not clear out all the stock between the lug and the wall but will instead leave a thin stiffening rib; finishing of the lug side of the wall is then performed in several Z-passes (each with 2 radial passes), leaving some rib to support each successive Z-pass. If the roughing tool is the same diameter as the finishing tool, then fine-tuning of the gap, and hence rib thickness, can be done through the radial stock-to-leave parameter.

6

Be mindful of the Minimum Radius parameter. The default setting leaves small radius corners incomplete for follow-up roughing with the rest-machining function of 3-D Adaptive Clearing. If your intention is to finish walls directly without a rest-machining step, then the default setting will leave a variable finishing allowance, causing your finishing cutter to be pulled into inner corners as extra stock is encountered and causing load spikes, potentially resulting in undercutting and compromised dimensional accuracy.

7

Depending on your requirements, you may find that trochoidal roughing with a low radial engagement leaves flat surfaces sufficiently smooth to warrant skipping a dedicated floor finishing step altogether. Indeed, it is tantamount to specifying a very large overlap in a constant overlap routine (with a possible difference that tool marks are more erratic-looking, if tool marks on a smooth floor are something that matter to you). If you decide to give it a try, program a micro lift for repositioning moves so as to avoid dragging the tool across the floor unnecessarily. If this works for you, then by all means go for it. However, be aware that should you find your Z offset a little too high, you then have to run the entire Adaptive Clearing routine again just to adjust your floor height.

8

Again, don't overthink applying principles of high-speed machining to non-HSM machining centers. Think of it as a ramping-up of your current parameters, or, better yet, as reaping the benefits of modern CAM programs and better computing power. Either way, you will be better utilizing your current machining potential and you *will* see results.

Feel free to contact Louis Martineau at louis.martineau@solidcad.ca if you require coaching on the implementation of Adaptive clearing in your machine shop.