EBOOK

Building a CAD PC: omnibus edition

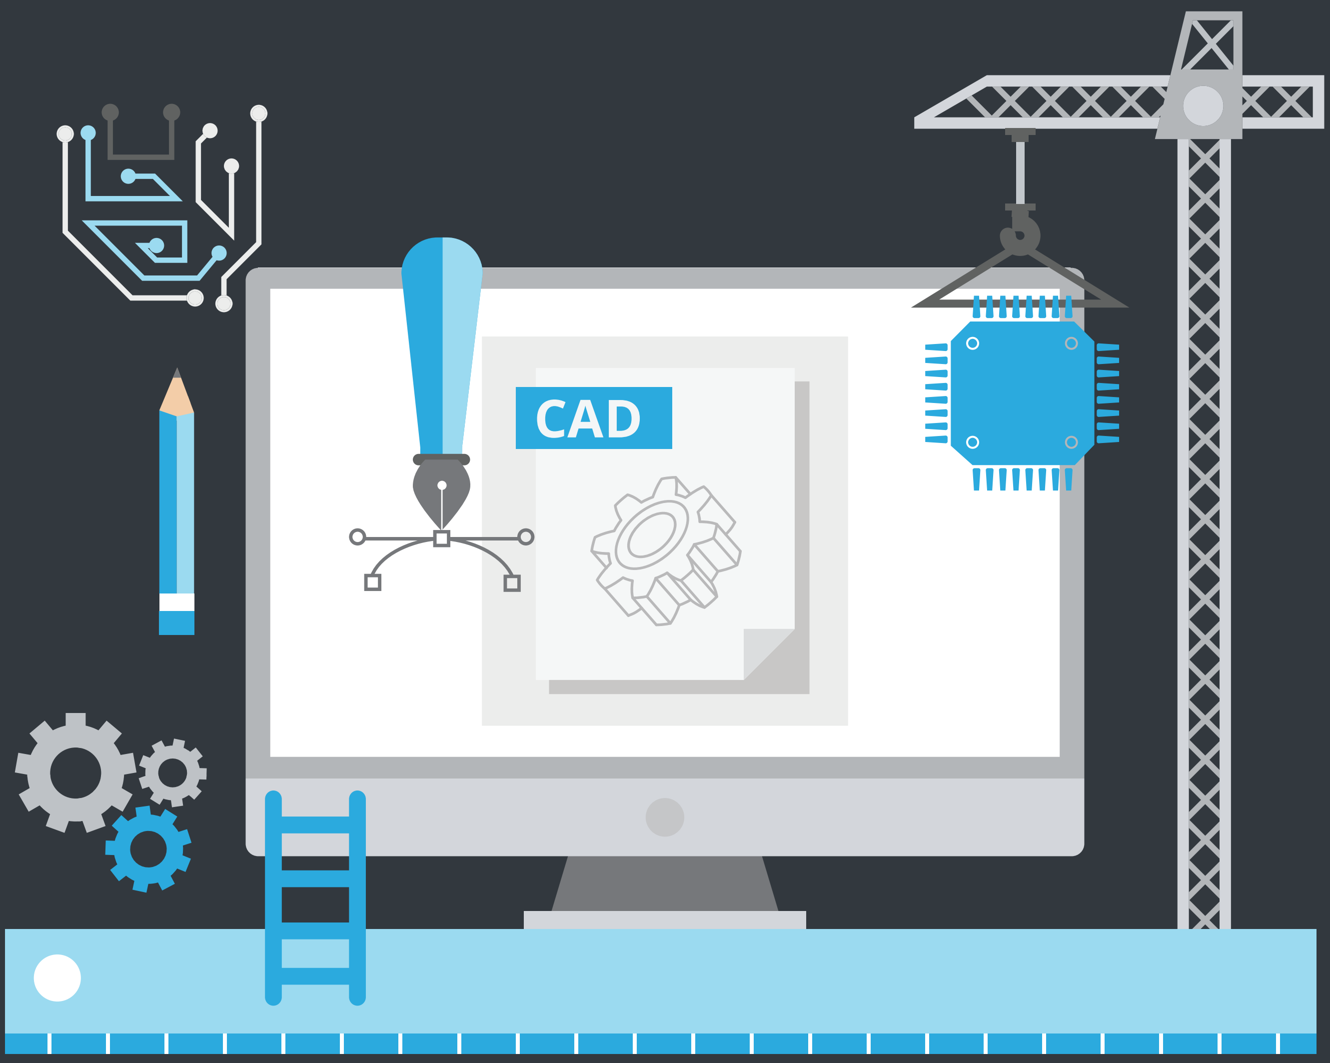

GRABCAD

Table of Contents

Introduction

You're here because you want to build an awesome CAD PC but would prefer not to cite Terry Crews when you explain your rig. What you want instead is a compilation of Ed Loptageui's *Engineering your own kick-ass CAD workstation* build series. Ed's bona fides include years as an aerospace engineer and technology evangelist. Check out his book. Let him know if you have questions. If you disagree, feel free to say so. Fair warning: you'll probably lose that argument.

Here come the 4K CAD monitors

Ed Lopategui

Chances are you've been staring at the same venerable CAD workstation display for quite some time, perhaps too long. This old relationship has probably carried on for years, surviving an upgrade here or a rebuild there. Like an old friend, you've been through CRT thick and LCD thin, teamed up to tackle some quality work and enjoyed a good time or two. But your dear friend is starting to look a little ragged. Now in 2016, you may have to gaze longingly at that anti-glare coated panel one last time, push that power button and bid adieu. Because here comes 4K.

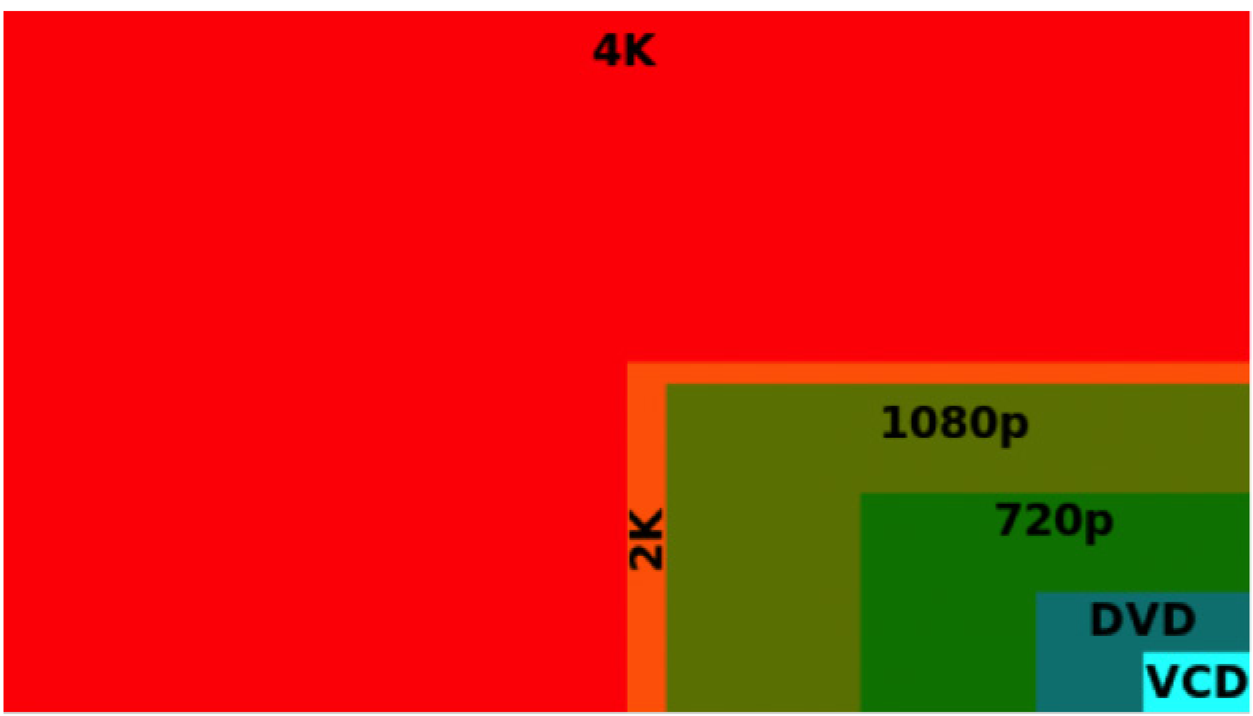

Resolution revelation

If compiled screen resolution statistic from various sources ring true, we're all long overdue for a display update. Approximately a third of us are stuck at standard HD resolution, i.e. 1920x1080 (otherwise famously known as 1080p) on a single display. Another third is enslaved by the less glamorous 1366x768 resolution perpetuated by laptops and less expensive desktop displays. The last third, more akin to the observant readership of this blog, has long recognized the utility of multiple displays for content authoring and CAD design, and typically carry dual 1080p displays for a grand total of 3840x1080 pixels of glorious screen real estate.

My particular setup, while comparatively rare, is a variation on the theme: 5840x1200 across three displays. This includes a positively Jurassic Hewlett Packard 2035 LCD panel, dating from the last ice age (~2004 or so) that dutifully serves as the tertiary display for Adobe Bridge sessions, streaming press conferences, or YouTube Beaglerush I2 runs without crowding my primary and secondary displays where actual work happens. While we've nursed setups like these for quite some time on the basis of economics and practicality, the price-performance case for 4K is almost ready to go mainstream.

4K is, sadly, not an engineering term

Just what exactly is 4K anyway? “4K” is what the salesman at Best Buy might remember to push in his pitch, but it’s less exacting than what engineers demand. 4K refers to “about 4000 pixels” but the reality, of course, is more complicated.

In film and video production the DCI 4K is the accepted standard, at 4096x2160 resolution. That’s a 1.9:1 aspect ratio (AR), which is a little oddly proportioned from the usual workstation display AR of 16:9 (1.78:1). If all that’s confusing, at least be glad that we’ve finally left the days of conventional television at 4:3, where cropping and/or letterboxed bars ruled the earth. If you ever owned a laserdisc player in the mid nineties, I feel your pain.

When it comes to computer displays, the relevant standard is UHD-1 (also referred to as Ultra High Definition TV), which checks in at 3840x2160 pixels. Yes, 3840 is less than 4000. There goes the whole 4K moniker, but it’s close enough for government work. If you’re more comfortable with the consumer television trend of labeling resolution according to the vertical dimension and scan type, then you can think of 4K as 2160p. “P” being “progressive scan,” of course. Thankfully no one in their right mind is using an interlaced “i” display anymore.

Desktop space: the final frontier

Whatever you want to call it, 4K affords quite a bit more resolution, roughly twice the number of pixels in a single display than what a current dual 1080p display setup affords. While displaying six simultaneous pages in Microsoft Word might have limited value, complicated graphics-intensive applications like CAD benefit tremendously from the additional space.

For example, a long time source of entertainment was creating a “whole enchilada” role (complete with an icon of said enchilada) in pre-ribbon Siemens NX, and turning on enough toolbars to reduce the graphics window to the size of a postage stamp. As ridiculous as it was, it really drove home the depth of the application. All that extra resolution adds room for multiple views, larger graphics windows, toolbars, dialog boxes, feature histories, expression tables, etc.

And for those of us who have a task bar that looks more like the unabridged history of computing, multitasking across all that space is a tremendous boon. And if you’re brave enough to step into the larger world of multiple 4K displays, it might very well be life-changing.

Attack of the upgrades

For quite some time, 4K displays weren't all that affordable. What's worse is that software support was iffy at best. Going 4K for a while was a splendid way to be consigned into a semi-hellish world of Lilliputian icons, illegible software menus, and OS scaling headaches. So the case for 4K was indeed a weak one.

But all that is changing. You'll find solid 4K displays on sale for $400 (albeit at 30hz refresh). Even the pricier IPS professional panels are quickly falling towards a cool grand, and OLEDs are finally starting to actually exist, albeit at world-breaking price points. What's more, direct software support is finally emerging. Native 4K support for Solidworks and CATIA is already available. NX 11 has support forthcoming. Solid Edge, AutoCAD, and the others likely aren't far behind. Make sure to check what versions of your favorite software natively support or otherwise behave well in 4K resolutions before you make the jump.

And therein lies the rub; the jump to 4K may trigger of wave of upgrades in both hardware and software. 4K adoption may give an instantaneous reprieve from downturns in PC hardware sales, where graphics horsepower has matured to the point that you can push pixels all day at existing HD resolutions. And yes, unless your graphics card is a recent model you may very well need one of those too. But don't complain too much: most haven't upgrading in a long, long time. You're overdue.

Warn your IT departments and significant others: 2016 may be the year of 4K. Don't fret too much, old LCD panel. We'll always have a jpeg of Paris.

APIs, GPUs, and drivers: CAD graphical conspiracy?

Ed Lopategui

Graphics performance is no doubt key to CAD productivity. Common recommendations involve running up the bill with professional level GPUs in certified hardware configurations. But is such a setup a wise investment? Hardware and software vendors tout how certified professional graphics cards are all that, and throw down the benchmarks to prove it. Many CAD hardware enthusiasts, however, contend that pro cards are perhaps an alien conspiracy designed to empty your pockets, and that consumer-grade gaming GPUs are up to the task at a fraction of the cost. The truth in the bewildering world of CAD graphics is complicated. But it's out there. You see a pattern emerging here, Scully?

Supernatural hardware shenanigans

The first point of contention in our little investigation: are professional GPUs running superior hardware? The answer to that is mostly no. Modern pro-level graphics cards share the exact same core hardware as their consumer counterparts and as far as the silicon is concerned, have equal graphics compute potential. We touched on this particularly irritating fact when making graphics card recommendations for our kick-ass CAD workstation.

There's minor variation with regard to binning, a common practice with mass-produced silicon, where specific chips are sorted based on their performance on internal quality tests during manufacture. Chips that well exceed the quality threshold are "binned" as more expensive professional parts on an assertion of reliability, while those that don't do as well, but still nonetheless still pass, are destined for lower-cost consumer markets. But the GPU cores are still the same.

Furthermore, the improved reliability of modern chip foundries make such distinctions practically irrelevant. In some cases, ECC memory is used onboard for pro-level cards, but that's not going to matter all that much either. Selling the same thing at two different price points sure does sound like something the cigarette-smoking man might use to pull your chain.

It gets worse

Believe it or not, in past years some professional level cards actually carried GPUs that were a generation behind their consumer equivalents. While that smells like the worst kind of conspiracy, most of this had to do with the longer lifecycle of CAD workstation hardware, where machines are typically leased/purchased on 3-5 year cycles. During that same time, hardcore gamers on the consumer side were pining after new and exciting hardware every 6-12 months.

Despite the fact that the silicon is the same, there's a deeper mystery. Take a pro card and a consumer card on the same workstation, fire up Specviewperf, run some tests (say for Solidworks, NX, or CATIA) and you'll find that consumer cards get destroyed in the benchmarks.

Every. Single. Time.

If there's no magical flux-capacitor in there to grant everlasting CAD goodness, then how, save for intervention by the supernatural, is such a thing possible? The truth, just like that Fight for the Future movie, may make you just shake your head in disapproval. The secret sauce is in the drivers.

Graphics driver secret police

Level two of the conspiracy is the graphics drivers which contain optimizations that specifically accelerate a variety of professional graphics applications, including CAD. The drivers are engineered to run only on the professional cards. Not because the hardware is more capable, but because they just won't let you.

Consumer graphics hardware is excluded from the inner circle.

Understandably, when confronted with such an unsettling truth, some enterprising CAD enthusiasts revolted, finding ways around such artificial limitations by soft or hard modding consumer cards to run an optimized driver. You can bet it wasn't long before the graphics card secret police closed the loophole and disposed of the evidence.

How can such shadowy evil be justified? Truth is, quite a bit of time and effort goes into creating and maintaining those driver optimizations. The process for creating certifications for each driver version, among the plethora of available hardware configurations and CAD platforms is onerous. So the cost of all that is passed to those that directly benefit –the professional users– while gamer cards are withheld from that privilege because they just didn't pay to play.

But that's still not all of the truth, the web of graphical intrigue goes deeper still, into very old battles over technology standards. Take for example benchmarks with AutoCAD,where the gaming cards actually prevail. How does something like AutoCAD manage to perform without optimized drivers? It's all about an old war over graphics application programming interfaces (APIs).

The API enigma

As we peel the onion back layer by layer, the fact that most CAD platforms are dependent on a graphics API called Open GL is revealed. It's a fitting joke that the API sounds like a bad 90s band, considering that the Open GL implementation for CAD hasn't changed significantly for about as long. Open GL's strength is in its extensibility, which was key in introducing CAD-specific rendering optimizations in a time when consumer-facing 3D graphics were in their infancy.

Open GL also featured prominently in the early evolution of 3D gaming, getting a notable boost from Quake ala John Carmack: full-time genius, the man behind Doom, failed rocket engineer, and now fully indentured Facebook employee. But chiefly this was from a time when games looked like some kind of neo-cubist artistic movement, with rough polygons big enough to land a plane on.

Gaming and consumer 3D has evolved significantly since. In the intervening years, Microsoft waged war on Open GL by creating, promoting, and evolving their own 3D API for Windows called DirectX. It's a war they handily won, and that consequently many have forgotten. That's why most consumer cards have crap OpenGL drivers, it's something that no one care about anymore, not even Carmack. It's a great topic to argue about at the local bar, after you've already flipped the table over arguing about Linux as a viable consumer OS and the host not-so-politely asked you to shut the hell up and talk about something else (like that's ever happened).

In 2010, Autodesk abandoned OpenGL and embraced DirectX. And now you know why AutoCAD works well with the consumer level cards.

With a DirectX rendering pipeline, consumer level graphics are more than enough for all but the most complex CAD applications. Expect more CAD platforms to follow in the future. Not to mention that for any browser-based CAD platform, the whole question is moot.

So will the very foundations of the professional graphics card market come crumbling down? Here comes the part that may make you happy then make you sad. Like when Scully and Mulder stare at each other for entirely too long.

The hard truth

CAD models just aren't complex enough to justify professional graphics anymore. Respectable 3D capability is becoming a commodity, as even common handheld devices start pushing enough pixels to handle the average design.

CAD is no longer near the tip of the spear for graphics rendering.

That may sound like an insult, or worse, grounds for a fight. While some of the most complex CAD assemblies still justify the current market segmentation, the average model is no comparison to mind-boggling developments in other corners of the professional graphics market. CAD is competing with the people who are busy rendering procedural fire or are cranking out the next Star Wars. CAD is overshadowed by James Cameron raising the bar on giant cat people, or Pixar generating photorealistic landscapes for The Good Dinosaur.

Open GL extensions for professional graphics functionality still matter to a point, but just not for CAD for much longer. The frustration will lie with the CAD software providers, many may be slow to adapt to this reality. But with every tablet that's used to run CAD, or for every CAD user who just runs a consumer card and doesn't miss the professional option much, that reality will get ever closer.

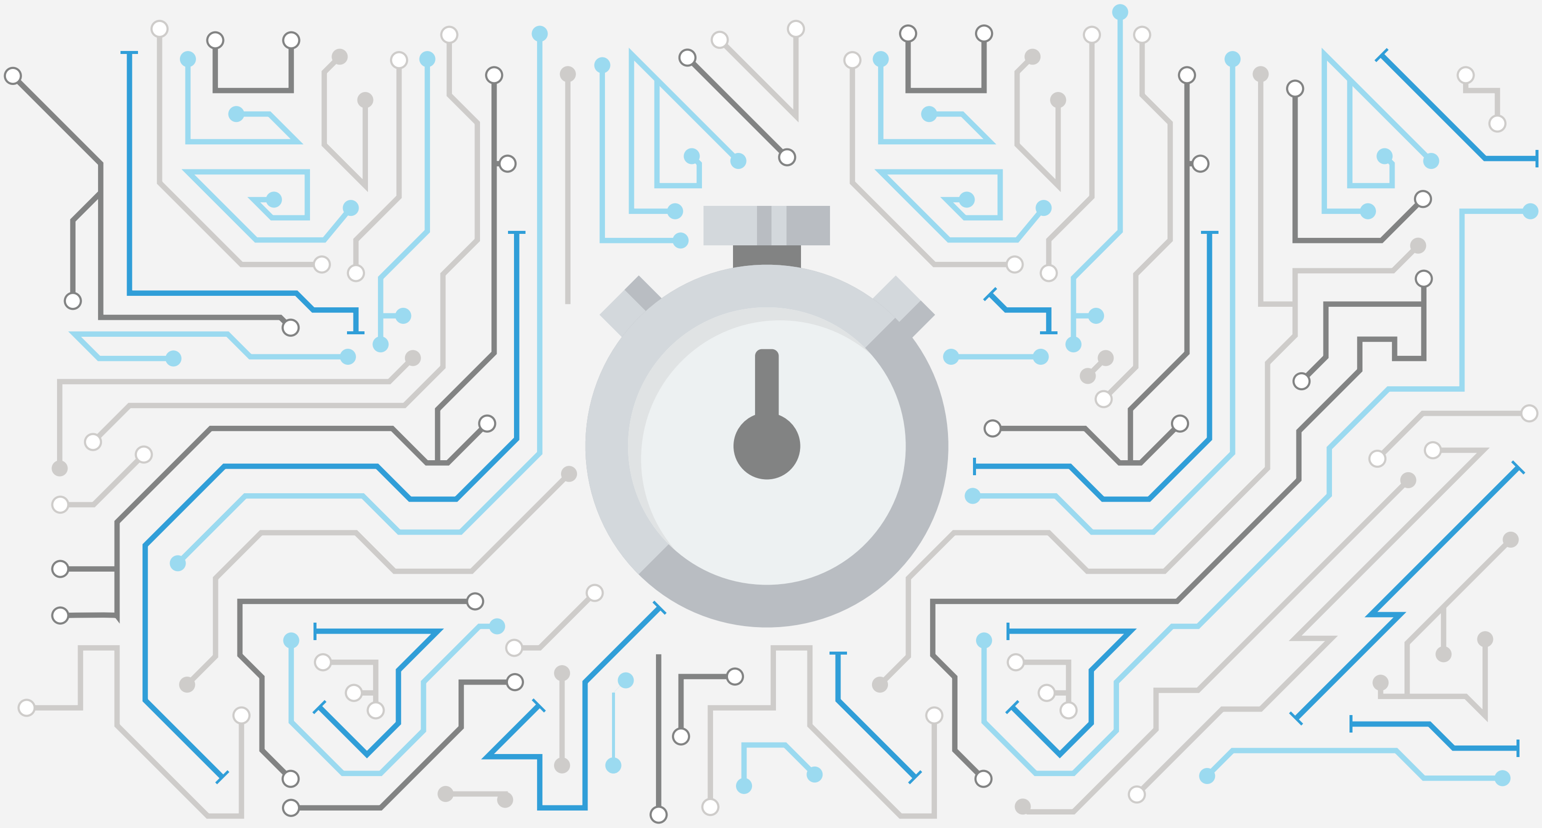

Overclocking for CAD: let's get dangerous

Ed Lopategui

You've heard the stories: buy a cheap CPU, throw some insane cooling on it and boost it into the stratosphere. Such is the promise of overlocking, pushing CPU processors past their design specs in the name of edging out extra performance for low, low prices. But these aren't the heady days of the Celeron 300A, where the megahertz flowed freely. With the invention of improved configuration aware microarchitecture, overclocking's gone a bit corporate these days. There's still some fun to be had, but more importantly, does it still make sense for your CAD workstation?

Lottery fine print

There are two rules in Overlocking Club:

1st rule: you need an AMD or K/X series Intel processor to overclock these days

The 2nd rule of overclocking club is... yeah, you can't overclock without an Intel K/X or AMD CPU.

If you wondered why we spec'd our machine with a K series processor and a Z97 mobo, well now you know.

Overclocking is often referred to as a lottery of sorts – luck is certainly a dominant factor. You're essentially gambling that CPU manufacturing variances have put your particular processor at an advantage over its identically engineered brethren. You're hoping you got the special one. And since microprocessor production is so much more efficient now than ever before, expect those variances to be comparatively rare.

Don't forget the inherent danger as well: you can smoke your chip into oblivion if you're exceptionally reckless. While committing involuntary chiptricide is uncommon, you may unwittingly reduce your processor's useful life by pushing it too far. There's no shame in backing out at this point, well, unless you don't like being called chicken.

Gear Up

If you're still with us, there's some prep work to be done before you start bumping multipliers and tweaking voltages:

- **Install an enhanced cooling solution**: You're not going to overclock on your stock OEM cooler. Just don't. You need something beefier. Fortunately there's a wide selection of readily available CPU cooling devices, including forced air or even liquid cooling. A variety of water-cooled kits come nicely sealed and pressure tested for ease of installation. If you're feeling particularly frosty many enthusiasts are dabbling in liquid nitrogen, but that's more for giggles than for practical applications.

- **Install monitoring software**: You'll need something to keep tabs on your processor, in order to make sure it's running how you have configured it, and to understand when you're skirting against thermal limits. Core Temp, CPU-Z, and Real Temp are among your options.

- **Install stress testing software**: You'll need an automated stress test to make sure your enhanced performance is viable under load and is not prone to crashing every 5 minutes. Just because the machine posts does not imply a stable configuration. Prime 95, IBT, and AIDA64 are popular for exactly this purpose.

- **Update your BIOS/UEFI**: You'll be depending on your motherboard BIOS to make all your tweaks, so make sure you've got the latest and greatest up and running.

- **Establish a baseline**: Benchmark your rig in its stock configuration, both in an idle state and at load. Before your doctor operates on you, he or she will establish a baseline, in case things go horribly wrong. Extend the same courtesy to your poor CPU, even if it doesn't have health insurance.

- **Research**: Lurk the internets for guides on your specific processor, which will give you common multiplier / voltage combinations and thermal loads that have proven workable in the past to define boundary conditions for your trials. You'll also want to note the upper limit on core voltage to provide a ceiling for your tweaking.

Let's get dangerous

Now is the time for your CPU to show its quality. Remember, swinging for the fences is unwise. Tweak in small increments. Test. Repeat. Gradual iterations will reduce the risk of toasting anything, as flaky stability serves as a leading indicator that you're reaching your particular CPU's limits.

First there's the lazy way to overclock. Most overclocker friendly boards feature overclock profile presets for your particular processor in the BIOS. It's basically automagic configuration based on a few simple inputs, and typically has a silly name like OC Genie, EZ Tuning Wizard, or Chipblaster+ (OK I made the last one up). If you're especially wary about fiddling with settings, this may be a good place to start. Just select a profile in the bios, reboot, and stress test while monitoring that core temp stays within reasonable limits. This is a great way to get overclocked without much hassle, though in most cases performance gains are modest.

If you're ready for prime time, you can dig right into manual configuration starting with the multiplier (aka the CPU ratio). The product of the core clock (typically 100 Mhz) and the multiplier determines the clock frequency. Hence a multiplier of 35 corresponds with 3.5 Ghz, 40 with 4.0 Ghz, etc. Change this value for all cores, reboot, and stress test. Keep an eye on core temperatures.

If you're solid, increase the multiplier again and repeat. Write down your configurations as you go. Eventually, the machine will lose stability and blue screen, then it's time to begin adjusting the CPU core voltage in small increments (typically .01-.05 V) until stability returns, stress testing with every change.

Once it's stable, bump the multiplier again and repeat. If you've overclocked a little too far, your machine won't ever stabilize no matter how much voltage you apply.

At this stage, crank your multiplier back 100 Mhz and change the voltage back to the last stable value for that multiplier. Stress test once more, and if your core temp looks fine than congratulations, you've pushed your processor as far as it will go. Depending on your cooling setup, however, you may reach a thermal threshold before you reach the ultimate limit multiplier/voltage limit.

Is overclocking worth it for CAD?

Overclocking might earn you a Ghz advantage in clockspeed, if you are both diligent and supremely fortunate. Ultimate performance gains typically are in the 5-10% range for CPU intensive CAD, which is not bad for "free." However, an overclocking setup is clearly a time-consuming process, and not without substantial risk. It's also not a sure thing, you may have gotten a dud in the batch, the folds a mhz above stock.

Unless you're overclocking for fun and nerd points, the time spent costs you as well. What seems viable for a student subsiding solely on ramen and mountain dew, won't make much sense for a successful engineer with disposable income. Cooking your CPU might help you stave off the next upgrade.

Truthfully, these days CPU margins are thin enough that coughing up cash for a better core is a almost always a better decision, especially if the microarchitecture and process size has changed. So do you still want to be an overclocker? The choice is yours.

CAD workstation performance tip: chuck those hard drives in the trash

Ed Lopategui

The age of hard drives has come and gone; the time of solid-state drives (SSD) is upon us. We made this point with our workstation recommendations a few months back, but a few of you were understandably skeptical. As engineers we have a special affinity for reliable technology, but such respect is not so easily earned. So are SSDs reliable? Let's end that debate and understand that it's well past time to chuck those hard drives. Yep. In the trash. See you in hell, magnetic platter.

All hail the Performance King

Even the fastest spinning hard drives are wholly outclassed by today's average SSD performance. A rather esoteric enterprise class 15k rpm hard disk might eek out 170MB/s with seek times hovering just under 4ms. You might be able to push that past 200MB/s using RAID 5 going downhill with an ample tailwind (provided you make the appropriate whooshing noises, of course). Even some of the newer drives that

use helium to reduce platter friction or cram unprecedented data densities with Shingled Magnetic Recording (SMR) only manage 150MB/s at 4.0ms access times.

In comparison, your average consumer SSD can easily manage 500MB/s transfers with just a .1ms access time. That's right, with no moving parts, there's no such thing as a seek time. SSDs are so fast in fact that the right setup can saturate your entire SATA bus, which has led to even more fanciful Non-Volatile Memory Express (NVMe) flavors of SSD that plug directly into a PCI Express lane on your motherboard to deliver an eye watering 5550MB/s. Ludicrous speed, indeed.

It should be no surprise that replacing a magnetic hard drive with an SSD is the single most cost effective upgrade to give your old CAD workstation a serious kick in the pants. Price per GB is of course higher for SSDs, but the cost continues to fall. 1 TB SSDs are already in the $300-400 range.

Don't SSDs degrade over time?

While SSDs have a clear performance advantage over conventional magnetic drives, there's still consternation about reliability. Is it true that each block of NAND flash memory can only be erased a finite number of times before it fails? You betcha. But what does this mean for reliability? To you, absolutely nothing. Let's understand why. First of all, it takes a really, really, long time for NAND block failure to become an issue. That's because modern SSDs use a variety of technologies to manage how your SSD ages including:

- **Wear leveling:** Technology that ensures read/write cycles are evenly distributed across all of the drive's memory blocks, ensuring that no single block gets overused and prematurely fails.

- **Over-provisioning:** The drive's advertised capacity is supplemented with additional data blocks to replace failed blocks. As blocks are marked unreliable, fresh ones are pulled from reserve until depleted.

- **TRIM:** Unlike conventional hard drives, data in NAND blocks cannot be directly overwritten. So in order to overwrite data, the data must be first read from its original block, then written to an entirely different block, which increases the number of write operations needed, a phenomenon known as write amplification. TRIM optimizes how blocks are treated with respect to permanently deleted files, by cutting down on the write operations needed for garbage collection (a process that frees blocks for overwritten data). The end effect reduces write amplification and increases the longevity of the drive.

But here's the real shocker, your magnetic drive degrades over time as well. Hard drives are subject to fragmentation, which reduces performance over time. Fragmentation is an irrelevant concept for SSDs as it takes the same amount of time to access any given memory block. Hard drives also develop bad sectors either due to heat, magnetic degradation or physical damage from vibration or sudden movement. SSDs are much more resilient to physical damage and produce much less heat, again, as there are no moving parts.

Yousa drive gonna die?

How long do SSDs last? The two SSDs used as my boot and primary program drives on my workstation (which typically runs 10-16 hrs/day) have worked flawlessly over 5 years. During that same interval, I've amassed a graveyard of dead magnetic hard drives on my server (though that does run 24hrs/day). Since anecdotal evidence likely won't satisfy your engineering sensibilities, check out this Tech Report study that took several SSDs and punished them with writes until they died. All of them lasted several hundred terabytes worth of writes, and others exceeded two petabytes. That's the equivalent of hundreds of years of average disk use.

So SSD longevity shouldn't be questioned. Now, it's important to mention that failure modes do differ. Hard drives tend to have a softer failure mode, where bad sectors mount at an accelerated rate before the drive becomes inaccessible. Well, most of the time. A catastrophic head crash can still cause insta-death. When SSDs fail, they tend to just stop and do their best impersonation of a brick. So if you're intimately concerned about the integrity of your data, you need to be making regular image based backups and monitoring your SMART data on a daily basis whether you have an SSD or not. If you're not doing so, well you're just asking for it.

Off to the magic cylinder

Now that you realize it's time to trash those clunky old magnetic drives, it's certainly understandable if you balk at the environmental impact of doing just that. In fact you shouldn't just chuck your drives in regular trash or feed them to wildlife at all unless you're some kind of monster. Most municipalities have electronic recycling programs for just that sort of thing. But there's a dilemma with that as well.

If you want to protect your privacy, you need to ensure data is erased. You can use something like Darik's Boot and Nuke (DBAN) to wipe information but that requires a working drive, and you still might be paranoid. The paid version of DBAN, Blancco, can erase data to a degree that is satisfactory with DoD 5220.2M. And if you're tin-foil hat paranoid, or just want to have a little fun, there's an infinitely more entertaining end game for your magnetic drive: wind chimes.

That also scores you some free rare-earth magnets in the breakdown process when you take the write head assembly apart. And no engineer should be without some spare rare-earth magnets.

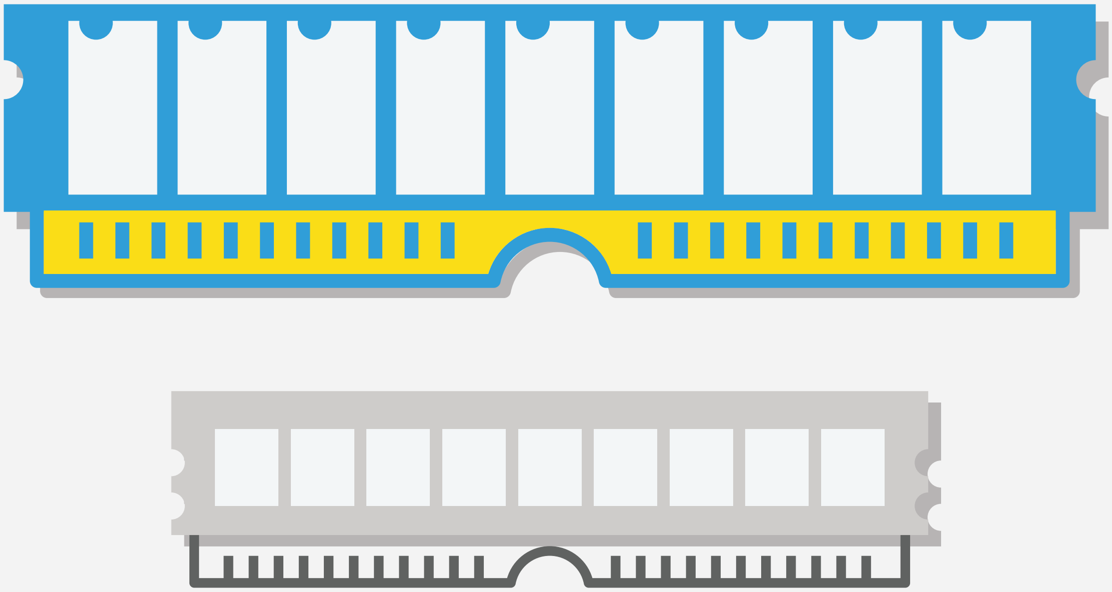

Why ECC RAM matters, but probably not for most CAD design

Ed Lopategui

When we put together a self-build kick-ass CAD workstation a couple weeks ago, we managed to raise a few eyebrows in our seemingly reckless disregard for random access memory. We didn't bother to spec a system with ECC RAM, and some of you questioned why populate a CAD workstation with "the cheap stuff?" To get to the bottom of this we're going to have to talk about billions and billions of bits, uncorrectable errors, hyperactive bananas, and supernovae. Quick, someone tweet at Neil deGrasse Tyson.

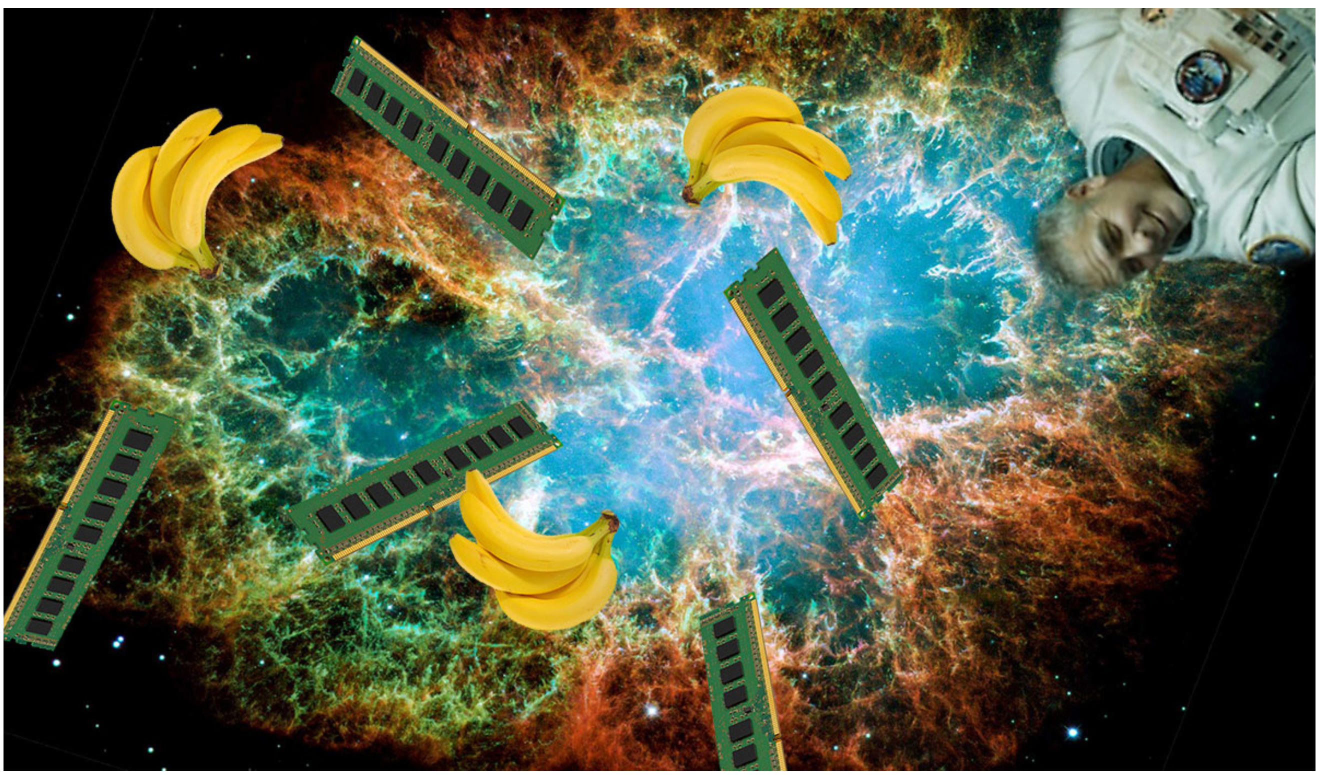

The universe wants to kill your CAD workstation

What causes memory errors? Technically speaking, everything. Computer memory errors can be divided into two classifications: soft errors and hard errors. Hard errors are simply explained: some kind of physical damage that causes one or more bits in memory to permanently misbehave. Soft errors are more esoteric: they are transient, instantaneous defects caused by the surrounding environment. In their simplest incarnation both hard and soft errors manifest themselves as a bit flip, meaning a single bit of binary information in the memory is altered, either from 0 to 1 or vice versa. Depending which particular bits get flipped among the 64 billion available on the average 8GB machine, the flip could mean a catastrophic crash or nothing at all.

Hard errors are attributed to internal component failure, power surge, or if you somehow manage to pummel your workstation with a particularly vengeful spinning helicopter kick. Soft errors, however, originate from more fantastic sources such as:

- **Cosmic rays**: A long, long time ago in a galaxy far, far away, a supernova ejects energetic protons careening across the cosmos all the way to your office on planet earth, and happens upon a DIMM modules in your workstation, temporarily freaking a bit or two out.

- **Radioactivity**: Trace amounts of radioactive isotopes like uranium-238 or thorium-232 occur naturally in the earth, and consequently are in pretty much everything, including the material used to make the memory chip itself. The alpha particles produced by decay of these trace materials can also flip bits. Fortunately, you needn't worry about the .0001 g of potassium-40 in that pair of bananas you had for breakfast because it only decays via beta radiation. In that case, your memory is probably safe.

Here comes ECC to save the day... mostly.

Error Checking and Correction (ECC) RAM is a step above your friendly, neighborhood memory. ECC technology can't prevent memory errors, but it can both detect and correct memory errors when they do happen, within certain limitations. Most ECC memory is engineered to detect and correct single bit errors, meaning one errant bit in a byte of memory. While typical ECC memory can detect two-bit and some multi-bit errors, it can't repair them. Such errors are uncorrectable. Certain exotic variations of ECC like IBM's Chipkill can wrangle multi-bit errors, but are a rather uncommon proprietary solution. One of the primary advantages of ECC is at least you know when and how many bit flips are occurring, with regular memory you haven't a clue. ECC allows you to truck along happily immune to single bit errors, and in this respect is clearly superior to non-ECC RAM. However, in order to understand ECC's value requires understanding memory error frequency and root causes.

Abort, retry, fail

How often do bits flip? For some time, the most often quoted benchmark was an old IBM study that claimed approximately one flipped bit per 256M of RAM per month of runtime. The more memory you have the higher the higher the chance you'll experience a bit flip. For someone working 40 hrs/wk on a workstation with 8GB of RAM that translates to about 7 or 8 flipped bits a month. More recently, Google conducted an exhaustive 2 and a half year study on their own server hardware that revealed some interesting insight into memory error rates. Some of the findings include:

- Error rates were highly dependent on hardware configuration, with some platforms showing errors in 20% of the DIMMs while other platforms exhibited errors in only 4% of the DIMMs. Google conveniently omitted naming any specific vendors, unwilling to throw any suppliers under the bus.

- Heavily utilized systems have considerably more errors, 2 to 3 times higher than less utilized systems. Google claimed their specific server utilization as sensitive, but you can bet these machines are being hammered pretty hard 24/7.
- Overall 8% of the DIMMs experienced at least 1 error per year. The rest didn't. At all.
- A DIMM that has experienced a correctable error is 9 to 400 times more likely to suffer from an uncorrectable error in the future.
- Because error rates had such a strong correlation with utilization, hard errors are likely the dominant root cause over soft errors.

The price of eternal memory vigilance

ECC is priced higher due to the extra error-correcting bits onboard and the fact that they are generally produced in lower volume as compared to their non-ECC consumer brethren. Depending on the size and particular speeds involved, the ECC premium can be anywhere from 5-100%. Adoption costs exceeds the RAM price differential, as you will also need an ECC capable motherboard, which in turn often requires a server class processor. Overall, you're looking at spending several hundred dollars to benefit from ECC memory protection with otherwise similarly performing hardware.

To ECC or not to ECC?

When it comes to most desktop CAD design, ECC largely doesn't make economic sense for a self-build. Right off the bat, you're spending money for an issue – correctable memory errors – that statistically will only affect 8% of your hardware, and only if the hardware undergoes a server-like utilization. At lower utilization, as is the case for most CAD workflows, error counts are 2-3 times less in the worst case.

But perhaps that's not enough justification for you. Know then that ECC is not a magic bullet, and requires a server-style maintenance philosophy to utilize effectively, otherwise it's a waste. The Google data indicates that modules with correctable errors are up to 900 times more likely to suffer from uncorrectable errors. You should have one of two reactions to this:

- **Holy crap, I should be monitoring my ECC memory!** Then you better read up on Windows Hardware Error Architecture

(WHEA) and keep some spare sticks around. Get ready to spend both time and money.

- **Wait, I have to monitor my ECC memory?** If you haven't bothered, and think ECC will save your bacon on its own, you're deluding yourself. You may have already suffered uncorrectable errors without noticing. If you're happy with uncorrectable errors, then you would likely be happy with non ECC RAM. You just spent your money for nothing. You're doing it wrong.

Finally, all of this assumes perfect software. While it seems really unpleasant to have a system crash because a star on the other side of the universe farted a million years ago, it's peanuts compared to how many crashes and problems you're going to experience because your CAD software is broken. Even in the case of a system crash, most file versioning and backup strategies are a more cost-effective investment. If you don't mind rebooting, you don't need ECC.

ECC only makes sense in server-like workflows such as FEM analyses or rendering, where a bump in the road costs hours of time. Well, unless you plan to design in space. But then you have a whole other set of problems, like how to keep George Clooney from eating the only piece of lettuce available.

Disclaimer: This article was written using non-ECC RAM and could be subject to eror. Oops.

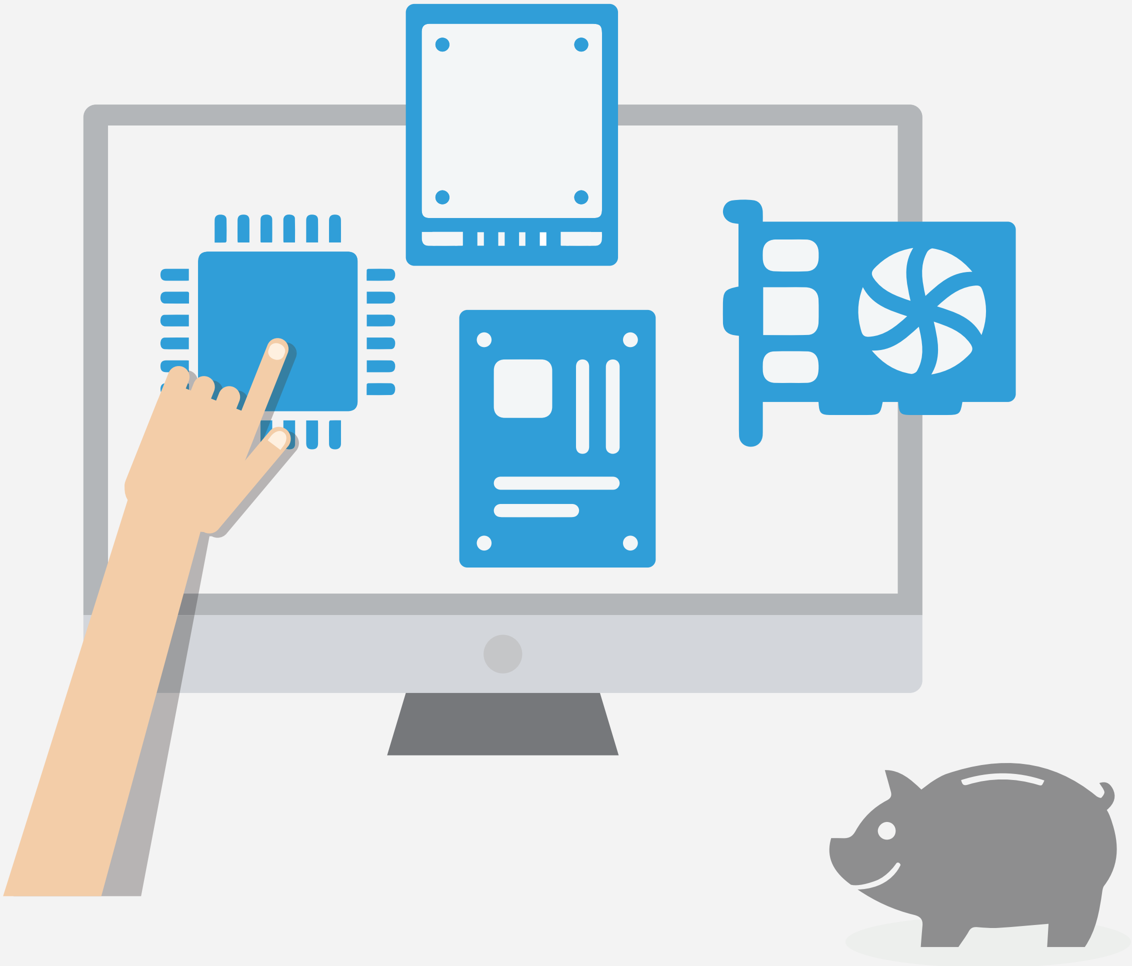

Engineering your own kick-ass CAD workstation build on a budget

Ed Lopategui

An engineer is particular about his/her tools, and there's no more important or personal an engineering tool these days than the CAD workstation. While you could just fork over your hard-earned cash for a turnkey CAD configuration designed for corporate sensibilities, you want a PC engineered your way. You want serious CAD power, but at reasonable prices. We can build it. We have the technology.

Since we're talking DIY builds, we're leaving laptops at the curb and focusing on desktops for now. While it's certainly possible to crank out a custom laptop, it's much like building a Hackintosh or a steam-powered giraffe: intellectually stimulating yes, but also ultimately pointless. We'll hit the highlights today, and can deep dive into your favorite topics in future posts.

Choose your own path in the certification labyrinth

What often causes the most confusion in spec'ing or upgrading a custom CAD rig is defining performance requirements that lead to choosing the right gear. While minimum CPU, RAM, and OS requirements are both easily discovered and rather universal, the moment you bring graphics cards into the mix, you plunge into the confusing quagmire of CAD hardware certifications. CAD software traditionally has been very sensitive to perturbations in the graphics pipeline, and the solution has been rather low-tech: vigorously test specific configurations. You'll discover that vendor certifications are based on specific graphics cards, on specific driver versions, and on specific operating systems (all of which are often not the latest and greatest). Furthermore, some certifications focus only on particular workstation models provided by named hardware partners. Considering how quickly everything changes, it's enough to make you go blind. Need a migraine?

Have a look for yourself:

Autodesk
Dassault
IronCAD
PTC

Siemens

So what's all this certification brouhaha mean for us DIY chickens? Exactly nothing. No one's necessarily going to certify your particular build, and that's the price of putting together your own rig. In many cases, you'll be <cue dark organ music> unsupported. Meaning if you run into problems with the software, you're likely on your own. Consequently, if you want to feel safe and warm, and make phone calls to tech support 3 times a day, then by all means, go get yourself a certified workstation instead. Last I checked, however, most of you are engineers. Let's not forget signs you're an engineer #3456: Warranties are meaningless. With that in mind, let's get to choosing components.

CPU: clockspeed ain't what it used to be

When it comes to CPUs, there was a time when all you had to do was pile on the Ghz, it was all about the Pentiums, baby. With Moore's Law under assault, and Intel stuttering on their ticks lately, the push toward ever-smaller die sizes and higher clocks are taking a physics beat down. Performance gains are more subtle, the real reason people aren't buying new PCs every year anymore. The good news is this means your wallet will thank you; getting a respectable CPU has never been more affordable.

Microprocessor performance, however, is now a non-obvious fusion of clock speed, process size, and architecture. Take, for example, the Intel i7-5820 which outperforms the i7-4790 on most benchmarks, despite a 700Mhz dearth in clock speed, and the same 22nm process, the difference being the Haswell-E architecture. The current sweet spot for value is around the i5-4690k, with strong single-core performance which matters most in a CAD universe that is still overwhelmingly single-threaded. Poor AMD has been playing second fiddle for the last 3 or 4 songs, and the nearest equivalent, the FX8350, doesn't fare as well. While you'll find most vendor CAD workstations are equipped with Xeon procs, the Xeon's larger memory address space and hyperthreading capability is going to be of marginal value for most of you, unless you're doing hardcore FEA or rendering most of the time.

Graphics card hunger games

So you've heard when it comes to CAD graphics, you must cough up the coin for the professional workstation cards (Nvidia Quadro or AMD Firepro). You can't use gaming cards, it'll unravel space-time, and possibly dissolve your kidney. Do not pass Go, do not collect 200 CUDA cores.

Professional cards are specially binned for enhanced reliability and benefit from dedicated driver support and optimizations for your CAD applications. The gaming cards are exempted from these optimizations, not because the hardware is necessarily incapable, but due to how the driver support has been coded. Truth be told you can run CAD (say AutoCAD or SolidEdge) on a Microsoft Surface Pro 3 with the rather homely Intel HD5000 (which is certified FYI). At that point most of the load is on the software renderer, but today's CPU's make it quite tolerable. You'd have a similar experience with a gaming card. It's not the best experience, mind you, but for most reasonably sized assemblies with the right graphics settings, it works. Going with a pro card is a smoother experience and certainly more costly, not because of hardware superiority, but just how the market's structured. Complicating matters, if you're an engineer, and have a PC, you're also no doubt a member of Gaben's glorious PC Master Race, and want to squeeze in some GTAV, Witcher 3, and Arkham on the side.

Ah, first world problems. What to do? Choose thy path:

- **Play it safe**: Get the fastest Quadro/FirePro you can afford. Play FTL and Sanctuary RPG

- **Combo Breaker**: Get two graphics cards. Multiple graphics cards with different drivers coexist just fine in Windows 7 and up. Get a lower end pro card (something like the Nvidia

Quadro K620) to ensure stability and use the optimized driver, and supplement with a gaming card to ensure happiness. Yes, you can have your cake and render it too.

- **Crazy go nuts**: Throw money and cares to the wind, pickup a TitanX with its floating point potential and general disrespect for normalcy, and blame any graphical glitches on solar flares.

Storage: save all the things

In the past SSDs were an optional enhancement if you had the extra loot. These days they're mandatory; anyone not using an SSD for their boot drive ain't right in the head. Early concerns over longevity and reliability of SSD technology are ancient history. They've also become extremely cost competitive, considering the significant speed advantage over magnetic hard drives. SSD's ramp up in cost as size goes up, the sweet spot being 250G.

If you need more space consider buying two identical drives: one for the OS, and one for your programs. It also makes good sense to supplement your SSD with a sizable magnetic drive, to store larger, less performance-sensitive data, like your totally legitimate 200G music collection.

Remembering about memory

Memory is plentiful these days, 8GB should be your minimum entry point for anything CAD related. 1600 Mhz DDR3 is the sweet spot for price to performance. If you have a few extra bucks, don't bother with getting faster RAM, just get more RAM. You'll also notice the vendor workstations come equipped with ECC RAM. Do you need it? Nope.avi

Don't forget the mobo

It's easy to overlook the motherboard, as it won't be responsible for dramatic performance gains. You'll need to pay attention to memory slots (4+), maximum addressable memory (32GB+), number of SATA ports (6+) and available PCI 2.0 and 3.0 slots to accommodate your graphics card(s). I'm partial to Asus, but EVGA and MSI make solid boards too.

You have the power

You need some electric juice for all this PC glory. Typically, a power supply in the 700-1000W range is what you need. Larger PSUs run more efficiently at lower loads, plus it buys you margin for that day when you decide that popping a couple TitanX's isn't nearly as insane as it sounds right this moment.

A place for your gear

Part of the fun of building your own rig is injecting a bit of your own personality. The best place to do so is via the case, a welcome improvement over the lifeless, boring black boxes indicative of a purchased workstation. Personally, I tend to err on the land-a-plane-on-it case size. Because you never know when you'll decide to cram 6 drives, 4 video cards, and a liquid cooling reservoir in it.

So let's lay it all out on the table and compare what we've got to the purchased equivalent:

A reasonable self-build (July 29, 2015):

CPU:	Intel Core i5-4690K LGA 1150	$239.99
Mobo:	Asus Maximum VII Hero Intel Z97 ATX	$209.99
Graphics:	Nvidia Quadro K620	$159.99+
	EVGA Nvidia GTX 980 Superclocked	$487.99 (With $20 Rebate)
SSD:	Crucial MX200 6G Sata III	$99.19
HDD:	Seagate Barracuda 2TB 7200 RPM	$76.99

RAM:	Corsair Vengeance 8GB (2 x 4GB) DDR3 1600	$45.99
Cooling:	Stock	
Case:	Cooler Master Cosmos SE	$169.99
PSU:	EVGA 220-G2-0750XR:	$129.99
OS:	Windows 7 Pro	$134.95
Keyboard:	Use what you got.	
Mouse:	Use what you got.	
Total:		**$1755.06**

Computer:	Dell Precision Tower 5810	
CPU:	Intel Xeon E5-1620	
Mobo:	A Scooby-Doo mystery	
Graphics:	Nvidia Quadro K620	
SSD:	256G SSD	
HDD:	2TB 7200 RPM HD	
RAM:	8GB (2x4GB) 2133Mhz DDR2 RDIMM ECC	
Optical:	Optical: DVD+/-RW	
Cooling:	Stock	
Case:	Generic Dell Sadness	
PSU:	Whatever they threw in there.	
OS:	Win 7 Pro / 8.1 Pro	
Keyboard:	Dell KB212-B Quiet Key	
Mouse:	Dell MS111 USB Optical	
Total:		**$2004 (After $668 Instant Savings)**

Final words

Building certainly has its advantages. We built something respectable without breaking the bank, managing to squeeze a bonus GTX 980 in there, and still saved $250 over the turnkey equivalent. Not bad, eh? You’re more than welcome to send me the difference in Steam gift cards.

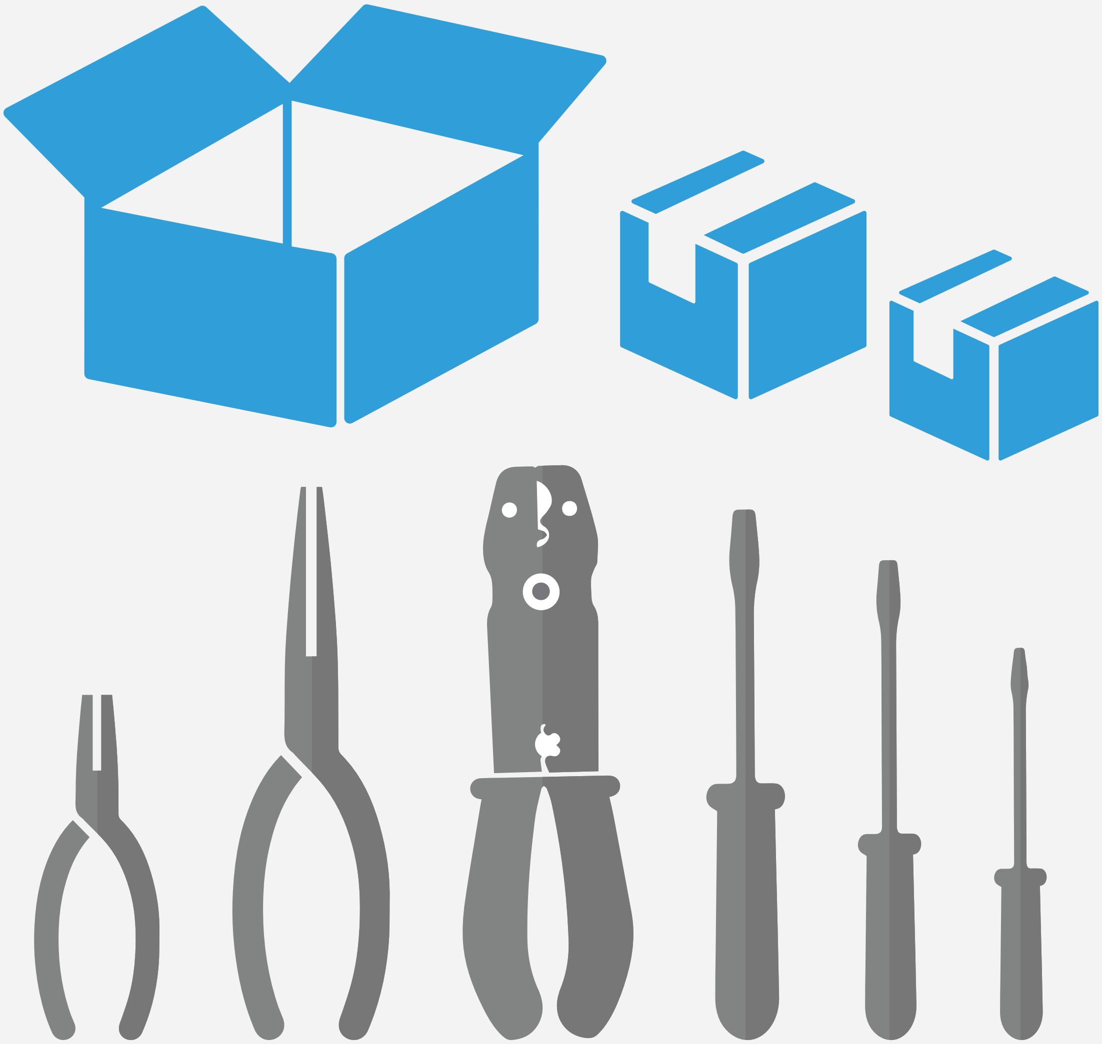

CAD workstation build tips from the trenches

John Cochrane

The sun rises; the small mountain of indistinct cardboard packages is illuminated by the morning sun as a glorious monument to technology and engineering. The deliveries have arrived in a steady procession over the last several days, solemnly and reverently brought to your doorstep by sentinels of commerce each demanding firmly that you sign here. With each new box containing carefully selected parts of your future DIY CAD workstation, the means to execute on your plan now lies before you. You've made some important decisions. Maybe you've followed some of our advice, or maybe you're rolled an infallible plan of your own. You're in the pipe, five by five. But now comes the moment of truth, and this machine's not going to assemble itself.

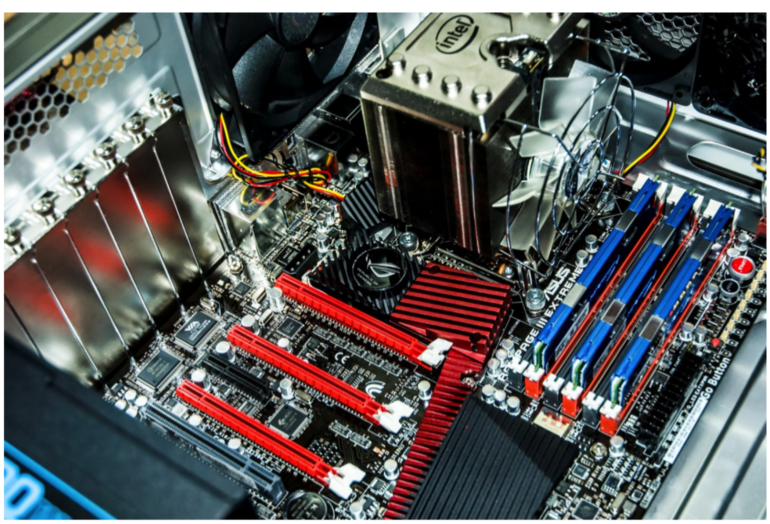

This post is not going to be one of those patronizing articles that shows you five dozen pictures about how to install your motherboard or plug in your video card. We're pretty sure you're fully capable of pushing Tab A into Slot B and can operate a Phillips screwdriver.

And if for some reason you can't, dude, you should have gotten a Dell. What we will offer is some practical advice that will save you some aggravation and optimize your troubleshooting. That's because when you undertake the challenge to DIY your CAD workstation, you proudly accept a harsh reality: you are your own tech support.

Forever.

That means taking certain precautions in case the build effort goes awry.

Preassembly

While it's always tempting to just go, there's some important groundwork you can do before even the first screw is turned. The following preparatory work can ensure a smooth build process:

- **Take inventory**: While this seems obvious, you'd be surprised how often you might be missing a cable or mounting rail. Make sure you have all the things.

- **Gather drivers and firmware**: Whatever is delivered on physical media with your hardware is likely already older than dirt, so toss all those driver CDs to the side. Visit the website of each OEM supplier and download the relevant drivers and firmware for every part in your build and keep them in one place for ease of access (say, a USB stick). Stay clear of beta versions. Some CAD platforms are especially sensitive to GPU driver versions, so check the certification basis (if there is one) to determine which driver versions are supported. Certified drivers are almost always significantly older than the latest available driver, it's best to go with the certified driver until you're more confident that your build is running smoothly.

- **Find a suitable spot for the build**: Find a decent work surface where you can comfortably work on the rig. Not in the middle of the carpet, in the kitchen sink, or on top of your dog. Clear the area of liquids and other contaminants, such as Mountain Dew and Cheetos. Have your tools readily available, and if you're using a table where you care about the finish, lay down a mat on top so you don't scratch it into oblivion. Finally, make sure your work area is reasonably secure. Finding a large cat or small child curled up on your motherboard just because you took five to grab a hoagie is the last thing you want to deal with.

- **RTFM**: Yes, I know reading is hard. But at the very least, skim the settings in the mobo manual - and decide how you're going to configure them. You'll thank me later.

Assembling Voltron

Now you're ready to start the exciting stuff, installing hardware and catching the sound of the thing finally roaring to life (Note: it's probably not going to sound like a roar, the days of the FX5800 are long gone). But before you get too caught up in the moment, keep to the systematic approach:

- **ESD Precautions**: Electrostatic discharge (ESD) can kill your newly acquired toys in an instant, but shockingly, the potential danger of ESD is sometimes overblown. While you could blow $5 on a wrist strap, it's a hassle to be tethered to your rig during the build. And you'll look like a total dork to boot. You're better off blowing that $5 on a zucchini-spice double macchiato with a quarter twist, and simply make sure you keep your arms in contact with exposed metal on the chassis whenever handling PCBs, and carefully handle everything by the edges. Also, just don't do any assembly in your socks on a shag carpet while petting furry animals.

- **Build in stages**: While you might be tempted to install absolutely everything at once, it really helps to establish some component level testing up front. Get the machine to Power-On Self-Test (POST) with just CPU, minimum memory, and GPU, with no OS installed and use synthetic benchmarks on a bootable USB stick to validate that everything's peachy. Nothing's harder than diagnosing a build issue when anything might be the culprit. This baseline will help you eliminate variables for any bumps down the road.

- **BIOS Update**: As soon as your machine POSTs, it's the best time to go ahead and flash the BIOS firmware with the latest and greatest that you downloaded earlier. BIOS bugs can cause all kinds of bewildering behavior, so it's best to minimize this potential failure mode. Some of the fancier mobos allow you to carry two different BIOS versions simultaneously and switch between them, but largely this is unnecessary unless you happen to do a lot of tinkering.

- **Firmware**: As you add additional components like SSDs, make sure to update them with the latest and greatest firmware before you even partition them; sometimes firmware updates are destructive to data.

- **Record your settings**: Record how you've set your BIOS in case you choose to tweak it down the line. Some BIOS allow saving your settings to nonvolatile RAM or externally.

- **A note on hyper-threading**: years ago hyper-threading caused some performance issues with software with poor SMP implementation, most CAD included. How hyper-threading has been implemented in recent processor design is more resilient, so it's best to leave this setting on in most cases.

Break in

Once you have everything assembled and running, you might think you're in the clear, but there's a few things best accomplished before you get to work on the new rig:

- **Clean up your cables**: Don't be that guy/gal who leaves their internal power and I/O cables looking like a rat's nest. Take the time to use the case's cable management (if any) and use twist ties to pack up loose ends and service loops. It's better for airflow and you'll be much happier the next time your crack the case open for changes/upgrades.

- **Wait to overclock**: Before you go to warp speed, make sure the ship holds together at impulse power. Seriously. Develop a baseline before creating problems for yourself.

- **Driver updates**: Install all your hardware drivers and make sure they're working correctly. Start with the motherboard drivers first, and then move on to graphics and whatever else you may have onboard.

- **Install your OS**: Now's your chance to choose the network name for the machine. So many LOTR locales, Disney characters, and space marines to choose from.

- **Benchmark and backup the clean configuration**: Before you load all your favorite resource-hogging, poorly-coded software, take the time to both benchmark and backup your workstation. Never again will the environment be so clean, unless you start from scratch with reformat and a new OS install. It'd be nice to have the option to revert to this clean configuration at any point.

- Burn-in: After you've gone through everything else and have your software installed, considering using a looping benchmark test or software script to run continuously for 24 hours, just to make sure you don't have any transient heat issues, etc.

With everything running in tip-top shape, you've finally laid the right groundwork to your workstation's long, productive life. Now it's time to get to work.

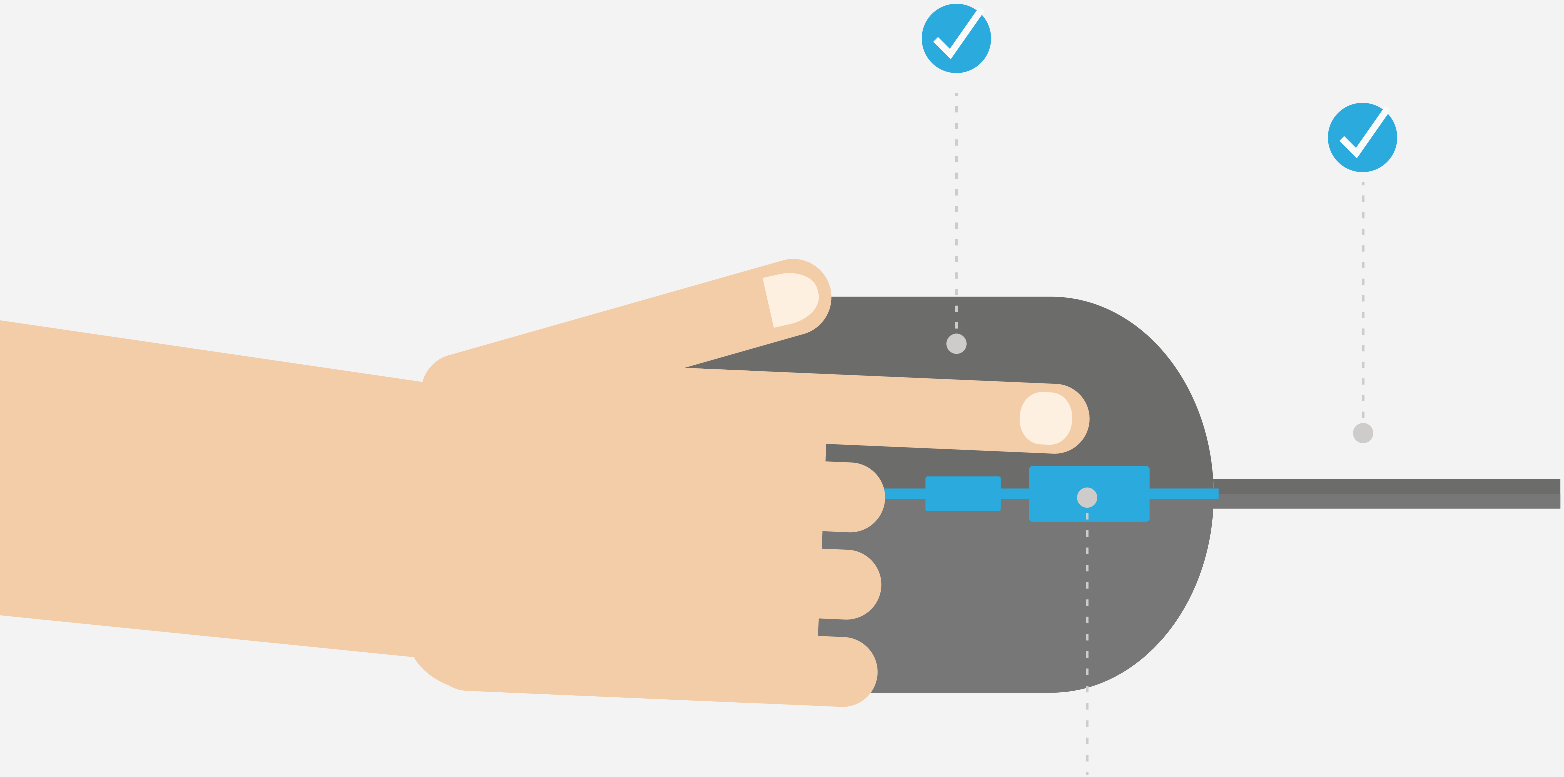

Our 3Dconnexion CadMouse review: precise clicking and high resolutions

John Cochrane

For today's design engineer CAD is the major software tool, many hours are spent at the workstation hammering away at the mouse and keyboard, it would seem that anything that helps with the ergonomics of this task would be a major benefit.

3Dconnexion, high precision, and you

3Dconnexion is already well known in this area for their range of 3D motion controllers (SpaceMouse, etc.). I have been a long-term user and fan of these devices since my days using Unigraphics and SolidWorks. 3Dconnexion, previously part of the Logitech organization, became an independent company in 2011, allowing it to focus 100% of its efforts on developing specialist input devices for CAD professionals...enter the CadMouse.

The mouse provides the following features:

- high precision
- separate wheel and middle mouse buttons
- zoom in and out side buttons positioned for the thumb
- customizable top button

On the software side the customization controls are integrated into the software for controlling the SpaceMouse, this allows all the buttons to be customized for the user's preferences and for specific CAD applications.

The mouse itself was very much plug and play with a wired USB connection – clearly a Bluetooth connection would be a good enhancement for the future. I did find that I had to set exceptions in my virus software to enable the configuration software to work.

Extreme resolutions and tiny icons

For me, the biggest problem I have with today's design and CAD software is with the extreme resolutions afforded by today's monitors not being considered (read: coded for), resulting in tiny icons in the UI and feature tree. These tiny objects are really hard to click on accurately using an "every day" mouse. The extra precision afforded here really helps with this problem. I found this made a real difference to the usability of SolidWorks on my Dell M3400 running on a 4000 x 3000 Dell monitor. I also found the extra precision of the scroll wheel and the dedicated middle mouse button to be very helpful in terms of less finger slips, misclicks, and strain on the tendon in my middle finger.

The thumb buttons I found I used less and when I did I found them quite hard to click with any accuracy, I had the feeling for me that the mouse was physically a little too big. I was left wondering if it would be worthwhile for 3Dconnexion to sell the device in different anthropomorphic sizes?

And the verdict

The top button was well set up with default button customizations, as the software seems to have a good knowledge of the CAD system you are currently running. It even adjusts the button set depending if you are in the part, assembly, or drawing functions of the CAD tool. Of course, you can also build your own customizations using lists of all the existing commands in your CAD tool, keystrokes, or even create macros.

So, for me, the CAD Mouse was a very helpful addition to my workstation, which made using SolidWorks more comfortable and efficient by way of the mouse's extra precision and the separation of the middle mouse button and scroll wheel functions. For relatively small cost, if you are already a space mouse user I'd say try the CadMouse.