SAMSUNG 840 PRO SSD
New controller. Competition killer!
PG. 84

GEFORCE GTX 660
The best deal going for gamers
PG. 74

BATTLE OF THE SUPERCOMPUTERS
Inside the global tech arms race
PG. 36

MAXIMUM PC

MINIMUM BS • DECEMBER 2012 • www.maximumpc.com

BEST BANG FOR YOUR BUCK!

PICKS FOR EVERY PART OF YOUR PC—BACKED BY BENCHMARKS!

PART BY PART, WE IDENTIFY THE PRICE/PERFORMANCE SWEET SPOT PG.22

GPU
1080p GAMING IS CHEAPER THAN YOU MIGHT THINK

MOBO
WHICH CHIPSET AND FEATURES OFFER BEST VALUE?

CPU
KNOW HOW MANY CORES YOU REALLY NEED

$7.99US
12>
Future MEDIA WITH PASSION
0 72440 38753 9

inside
DECEMBER 2012

On the Cover
Photography by Mark Madeo

FEATURES

QUICKSTART


Sayonara, SM!

R&D

LETTERS

IN THE LAB

MORE +

inside
DECEMBER 2012

On the Cover
Photography by Mark Madeo

FEATURES

QUICKSTART


Sayonara, SM!

R&D

LETTERS

IN THE LAB

MORE +

MAXIMUM PC

EDITORIAL
Editor-in-Chief: Katherine Stevenson
Deputy Editor: Gordon Mah Ung
Editor: Josh Norem
Contributing Editors: Alex Castle, Nathan Edwards
Contributing Writers: Brad Chacos, Alex Cox, Alan Dexter, Justin Kerr, Tom Halfhill, Paul Lilly, Thomas McDonald, Jamie Middleton, David Murphy, Quinn Norton, Jonathan Roberts
Copy Editor: Mary Ricci
Podcast Producer: Andy Bauman
Editor Emeritus: Andrew Sanchez

ART
Art Director: Richard Koscher
Contributing Photographer: Mark Madeo

BUSINESS
Vice President, Consumer Media: Kelley Corten, kcorten@futureus.com
Vice President, Sales & Business Development: Nate Hunt, nhunt@futureus.com
National Sales Director: Anthony Danzi, adanzi@futureus.com
Associate National Sales Director: Isaac Ugay, iugay@futureus.com
Regional Sales Manager: Christina Grushkin, cgrushkin@futureus.com
Account Executive: Austin Park, apark@futureus.com
Advertising Coordinator: Heidi Hapin, hhapin@futureus.com

Marketing & Sales Development Director: Rhoda Bueno
eCommerce & Fulfillment Director: Lisa Radler
Consumer Marketing Manager: Jong Lee
Newsstand Director: Bill Shewey

PRODUCTION
Production Director: Michael Hollister
Production Manager: Larry Briseno
Production Coordinator: Jose Urrutia
Senior Print Order Coordinator: Jennifer Lim

FUTURE US, INC.
4000 Shoreline Court, Suite 400, South San Francisco, CA 94080
Tel: 650-872-1642, www.futureus.com

Chief Operating Officer: Rachelle Considine
Vice President & Chief Financial Officer: John Sutton
General Counsel: Anne Ortel

SUBSCRIBER CUSTOMER SERVICE
Maximum PC Customer Care,
P.O. Box 5159, Harlan, IA 51593-0659
Website: www.maximumpc.com/customerservice
Tel: 800-274-3421
Email: MAXcustserv@cdsfulfillment.com

BACK ISSUES
Website: www.maximumpc.com/shop
Tel: 800-865-7240

REPRINTS
Future US, Inc., 4000 Shoreline Court, Suite 400,
South San Francisco, CA 94080
Website: www.futureus.com
Tel: 650-872-1642, Fax 650-872-2207

Future
MEDIA WITH PASSION

Future produces carefully targeted magazines, websites and events for people with a passion. We publish more than 180 magazines, websites and events and we export or license our publications to 90 countries across the world.

Future plc is a public company quoted on the London Stock Exchange.

www.futureplc.com

Non-executive Chairman: Peter Allen
Chief Executive: Mark Wood
Group Finance Director: Graham Harding
Tel +44 (0)20 7042 4000 (London)
Tel +44 (0)1225 442244 (Bath)

Gordon
Mah Ung

WILL YOUR NEXT GAME 'CONSOLE' BE A PC?

GAME CONSOLES aren't old—they're practically walking fossils.

Don't believe me? The Xbox 360 has at its heart a Radeon X1800 and its CPU has been pretty much abandoned by every consumer computing company. The PlayStation 3 isn't in any better shape, with what is the equivalent of a GeForce 7600 graphics part. Both machines' graphics parts are about six generations older than anything available on the PC at this point.

I bring this up not to rub specs in the face of game consoles, but because of what's happening with Valve's new Big Picture Mode. To see more about the mode, read this month's Quick Start and Lab Notes section. In a nutshell, it's a 10-foot UI on the 900-pound gorilla of PC gaming: Steam. The upshot is that I seriously think Valve is making a run at the living room. If it's as successful as Steam is on the desktop, the PC may be the new king of "console" gaming.

I know, we'll get the usual naysayers who talk about how consoles "just work" and you don't have to pay as much for the console as you do for a high-end gaming machine. Well, buddy, you don't need a $5,000 gaming machine to play games at an "HD" resolution of 1920x1080. And remember, the vast majority of console games today run at 1280x720.

You could take the midrange Alienware X51 for about a grand and get decent living room gaming performance. Alienware isn't the only vendor into console-size systems, either. Falcon Northwest has its nifty and typically powerful Tiki, too. I suspect we'll see other similar designs in the months to come from other OEMs.

And think of what a midrange or entry-level gaming PC gets you over a console today: True 1080p gaming, independent developers, free-to-play models, DirectX 11.1, PhysX, MMOs, and of course, better graphics.

There's much work to be done by Valve in this area, such as addressing the mouse-and-keyboard play for first-person shooter games and also the one area console's dominate in: split-screen gaming.

Valve has some time, though. The Xbox 720 isn't expected to ship until next year at the earliest and the PlayStation 4 may not come until 2015. One thing I do know: When both the Xbox 720 and PlayStation 4 ship, the PC will still be ahead.

Gordon Mah Ung is Maximum PC's deputy editor, senior hardware expert, and all-around muckraker.

↘ submit your questions to: comments@maximumpc.com

THE NEWS

From PCs to HDTVs, Valve Wants to Rule 'em All

Valve's 'Big Picture' mode for Steam could be the beginning of something much bigger (figuratively and literally)

PC GAMING is evolving right before our very eyes. It always has been, but every so often, a major mutation disrupts the status quo and changes the PC gaming landscape forever. The last time it happened was a decade ago. A company called Valve envisioned a future where digital delivery would be the norm, not the exception, and so Steam was born. Fast forward to today and there are now over 1,500 titles available on Steam, along with 54 million active users. Way more often than not, Steam is how PC games are bought, installed, and played.

Now, the visionaries at Valve are hoping to once again turn the PC gaming industry on its head. They want you to think outside of the traditional desktop experience, the one in which you sit at your desk wielding a keyboard and mouse to control the action fed to your 24-inch monitor from your PC tower. Like they did so many years ago, the minds at Valve are thinking of ways to transform PC gaming.

Some of it is born out of perceived necessity. Valve CEO Gabe Newell has serious concerns with Windows 8. His disdain for the platform stems from the possibility that Microsoft might be planting a walled garden around its operating system, shunning platforms like Steam in favor of its own Windows Store. But beyond self-preservation, Valve has discovered a genuine interest in hardware innovation.

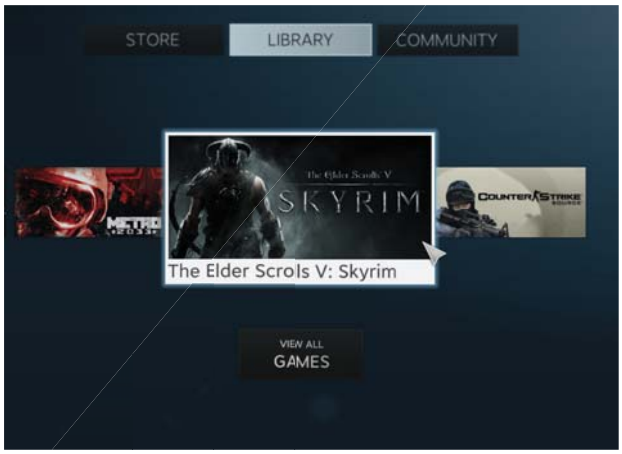

Big Picture is Valve's brand-new 10-foot Steam UI reformatted for viewing and playing your library of Steam titles on your HDTV.

At Valve, employees pick the projects they want to work on (a revelation that became public when an internal employee handbook was leaked to the web), and in a recent blog post, Valve programmer Michael Abrash discussed his passion for wearable computing. According to Abrash, "20 years from now [wearable computing] will be standard, probably through glasses or contacts, but for all I know through some kind of more direct neural connection. And I'm pretty confident that platform shift will happen a lot sooner than 20 years... quite likely as little as three to five."

Another interesting, and potentially more viable development taking place at Valve is extending Steam's reach into the living room. With "Big Picture," Valve has introduced a 10-foot user interface (UI) that's meant to be viewed and manipulated on your HDTV from your couch. Quite literally, this is a potential game changer, both for PC gaming and for consoles.

Currently being tested in beta form, Big Picture encroaches on the console's territory, harnessing the dynamic power of the PC over your living room television set. The snazzy UI gives you easy access to your entire library of Steam games. It implements a Dasiywheel interface for typing with a gamepad, if you choose to use one, and introduces what Valve claims is the world's first first-person web browser for the TV, and it "doesn't suck," according to Valve.

Big Picture could be the first step toward a Linux-based Steam Box, essentially an HTPC wrapped in a user-friendly UI. Valve is already working on porting Steam to Linux, and in a recent job listing for an Industrial Designer, the company said it was "frustrated by the lack of innovation in the computer hardware space." If Valve were to build its own hardware, it would create an insurance policy should Microsoft screw things up with Windows 8. Valve also stands to gain a new audience. Console gamers, numerous though they might be, are treated like second-class citizens, forced to game on dated hardware, sold gimmicky add-ons, and presented with a modding scene that doesn't hold a candle to the one that exists in Steam.

When you think about it, why wouldn't Valve launch a Steam Box? **–Paul Lilly**

Office 'Cloud' Pricing Revealed

Microsoft has announced pricing for its new cloud-based version of Office dubbed Office 365. Home users can expect to pay $99 annually for a five-user license, so the whole family can share a spreadsheet, Word doc, or a presentation that's stored in the cloud and accessible from any PC. We can already see the benefits: "Jimmy, do your homework!" "But Mom, the cloud is offline!" Ahem. The annual subscription also nets you an extra 20GB of SkyDrive storage (in addition to the 7GB everyone gets for free) and 60 minutes of Skype world calling per month. The subscription comes with Word, Excel, PowerPoint, OneNote, Outlook, Access, and Publisher. **–JN**

Technicolor Announces Certification Spec

Monitor color calibration is one of the most misunderstood aspects of computing, but Technicolor is hoping to put some clarity into the situation.

The company is pushing a new monitor certification process that is supposed to let the consumer see a movie or video the way the creator intended it to be seen. The process involves getting a monitor, laptop, or all-in-one certified by Technicolor and then running an installed application that automatically adjusts the color space based on the content specifications. Launch a web browser, for instance, and the color space will shift to a web browsing mode. Launch a Blu-ray player and the app will adjust colors, as well. **–GU**

What We Know About Haswell, So Far

Intel's new Haswell processor could be one beast of a processor when it pokes its head out from underground sometime next year. Intel says Haswell will have versions cutting power consumption to 10 watts—a 40 percent improvement over today's Ivy Bridge units.

More importantly, Intel has done enough tweaking to its newest 22nm processor that it'll qualify as a legitimate "tock" in performance rather than a "tick." Tocks, as you should know, indicate big performance leaps, while ticks offer smaller increments. Ivy Bridge offered a 10 to 15 percent bump over Sandy Bridge—not big enough to qualify as a tock.

So what other features qualify Haswell as a tock? Much-improved parallelism, a new AVX2 instruction set, and, of course, much improved graphics.

Among one of one the more exciting features will be the new Transactional Synchronization eXtensions. TSX would let a program access a table without having to lock the entire table to prevent other threads from creating inconsistencies, Intel says. The workaround to locking an entire table has been to create multiple smaller tables from the same information, but that adds overhead and complexity to the process. TSX promises the performance of smaller tables with the complexity of a large table.

Haswell will also be an important CPU for Intel as it is the first CPU that was designed under the company's mandate to make its chips more power efficient. Among the parts will be 10-watt Haswell CPUs that move the south bridge, or PCH, into the CPU. As is Intel tradition, Haswell will introduce a new socket, LGA1150, to support the new chip. LGA1150, as you probably have guessed, will make your LGA1155 board immediately obsolete. **–GU**

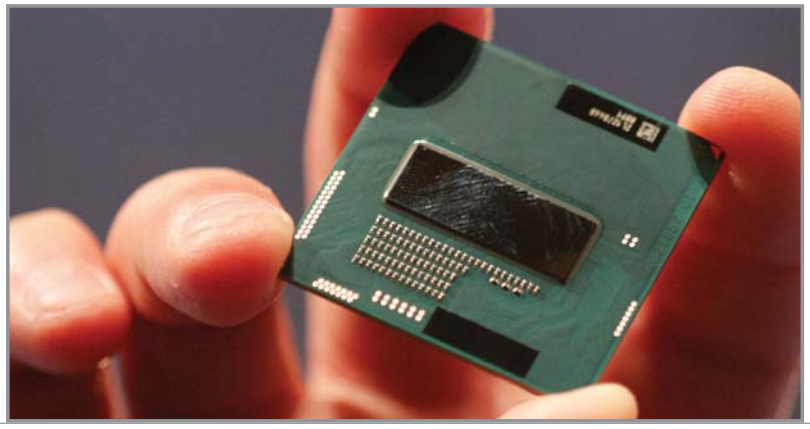

Tom Halfhill
Fast Forward

ERROR-PROOF PROCESSORS

SOMETIMES, CRYSTAL balls are made of silicon. To glimpse the future, look at extreme microprocessor designs, even if you'll never buy one. In this case, trickle-down theory really works.

Consider Fujitsu's new SPARC64 X, the company's tenth-generation implementation of the SPARC architecture. Never heard of SPARC? Don't feel bad. Although it was a pioneering RISC architecture, now it is clinging to life.

SPARC was introduced in the 1980s by Sun Microsystems, which Oracle acquired in 2010. Fujitsu and Oracle are the only major SPARC vendors, and their processors are found mainly in their own servers and supercomputers. One of Fujitsu's supercomputers has 88,128 SPARC chips and was recently the world's fastest.

SPARC64 X has 16 dual-threaded cores, 24MB of cache, and is targeting 3GHz. Only the core count beats Intel's Xeon server processors, but the usual specs aren't what grabbed my attention. It was the numerous error-checking circuits—thousands of them.

All server processors support error-correction codes (ECC) on their memory interfaces to verify data reads and writes. Some processors also have ECC or parity protection on their caches and registers. Only a few processors provide error protection for their internal data pathways, peripheral-I/O interfaces, and other critical components.

SPARC64 X does all that and more. It blankets almost the entire chip with 53,000 error-checking circuits. Among other things, they can trigger an automatic instruction-retry mechanism if something goes wrong. Without any software intervention, the chip can re-execute a program instruction that was derailed by a transient soft error. These errors can be caused by electromagnetic interference or even cosmic rays, and they are becoming more common as transistors keep shrinking.

Not every microprocessor needs this extreme level of error protection right now, but my crystal ball tells me that someday, they will.

Tom Halfhill was formerly a senior editor for *Byte* magazine and is now an analyst for *Microprocessor Report*.

Thomas McDonald
Game Theory

HOW TO RUIN A PORT

I'D BEEN hearing great things about Dark Souls since its release last year on Xbox and PlayStation. I didn't get a chance to play it the first time around, and when I heard it was coming to PC in a port with the irresistible title of Dark Souls: Prepare to Die Edition, I figured I'd just wait.

Big. Mistake.

Dark Souls serves up a healthy portion of hate. Hate for gamers who wait for special editions. Hate for keyboards and mice. Hate for PCs. And, most definitely, hate for PC gamers. Nothing else explains the lazy, lackluster port, which some experienced reviewers are already placing in their all-time top 10.

Graphics are locked into a crummy 1024x720 and they look dim and dull. A user came up with a mod that unlocks 1920x1080. There are also third-party fixes for the ugly low-res text fonts and user interface, and a mouse cursor that floats in the middle of the screen. Lone fans can do something in a few days, for free, that the developers couldn't do?

But the real cauldron of pain is the control system. You can't play with keyboard/mouse; you need a console controller. As an extra special "up yours," From Software leaves the A/B/X/Y/directional icons from the Xbox on the screen, in all menus, and on the interface prompts. This makes the tutorial useless, since it says things like "Left stick + B: Roll." Even worse is the weird, kludgy way the camera follows along, making the game extremely difficult and unpleasant to play.

Dead Souls was probably a good game on consoles, but that's irrelevant. This may well be the worst port I've ever played. I keep reading reviews that make excuses for this steaming pile just because the reviewer liked the original.

Screw that. I like consoles fine, but this is a PC game. Reviewers need to try this little thought experiment: If this game had landed on your desk cold, with no console experience behind it, what score would you give it?

Me? I never played the console game. I have no problem saying that Dark Souls: Prepare to Die Edition is a disaster. Do better next time, or *don't bother.*

You can follow Thomas McDonald on Twitter: @StateOfPlayBlog.

Intel Moves Beyond Touch

In an effort to make computers more perceptual, Intel has released a new software development kit and joined with Creative to let your PC literally see you wave at it or give it the finger.

The new Creative Gesture Camera features a dual-array mic, 720p sensor, QVGA depth sensor, and has enough resolution that it can tell if you are holding up five fingers or one. The SDK includes speech recognition, facial analysis, and can track objects in 2D and 3D. One example of the technology: Clench your hand in front of the camera as though you were opening a door knob to open a door knob onscreen in a game. The SDK is free, but the camera will run developers $150.

Intel hopes to bootstrap interactive computing models and applications on the PC and envisions a day when a standard monitor or laptop will come with the features built in directly. **–GU**

One in Five U.S. Internet Users Can't Stream 2Mb/s Video

According to a recent study conducted by video-hosting company Wistia, almost one in five U.S. Internet users are unable to reliably stream HD video over their connections. Even more depressing is the bar Wistia used to make the HD-capable determination: Compression technologies allow a 720p signal to squeeze down a 2Mb/s connection, and that's something 18 percent of U.S. Internet users simply can't do.

As one of the largest video-hosting companies around, Wistia is in a unique position to collect and analyze the viewing habits of millions of U.S. customers, and their conclusions are based on several weeks of data. Another finding: 10 percent are hovering in the sub-1Mb/s range.

Broken down geographically, Illinois, Mississippi, New Mexico, and Montana are the most hobbled areas, with nearly 40 percent of people incapable of 2Mb/s speeds. **–JK**

Increase Your Wi-Fi Range by 7,500 Feet

Amped Wireless aims to solve the problem of weak Wi-Fi signals with its new AP20000G dual-band Wi-Fi access point ($170, www.ampedwireless.com). According to the company, this high-power device will extend the range of your Wi-Fi coverage by up to 7,500 square feet.

The AP20000G works its Wi-Fi-extending mojo with a 620MHz processor, premium high-power Wi-Fi 600mW 2.4GHz amplifiers, high-power 5GHz amplifiers, and two detachable dual-band high-gain antennas. Just connect it to your home router (wired or wirelessly) and the AP20000G will do the rest, Amped Wireless says.

In addition to improved Wi-Fi coverage, the AP20000G sports a USB 2.0 port for file sharing, support for advanced security protocols, up to eight Wi-Fi networks for guest access, plug-and-play setup, and four Gigabit wired ports. **–PL**

Quinn Norton
Byte Rights

ACTA-ING UP AGAINST TPP

AT THE BEGINNING of 2012, protestors filled the streets of Europe to stop ACTA—the repressive copyright treaty secretly negotiated around the world. They damaged ACTA badly, though it's not clear that it's dead, and we've signed on with every intention of treating it like law. But maybe more importantly, hundreds of thousands of Europeans woke up the world to the fact that people do care about their digital rights.

Well, most of the world. Not the people already secretly negotiating the "Son of ACTA"—the Trans Pacific Partnership, a trade agreement suspected to be an even more oppressive legislative turd that American rights-holders are thrusting on the world. But no one was sure, because not even Congress was allowed to see it—that is, until draft text was leaked to the EFF. Even after the leak, the secrecy continued, with trade representatives refusing to answer questions about the leaked text, and refusing to give over any other material to the public.

But from the leak, it's bad. Many of the provisions of ACTA are back, and aggressively worse. TPP reaches into your computer, regulating temporary copies of copyrighted content, like cache. This version of TPP would make writing tools that interfere with DRM illegal, gut fair use, and turn what are currently civil matters into crimes. Because TPP only currently covers Pacific Rim nations and bits of Asia, we can't count on the Europeans to save our butts this time. No, if we want to stop a treaty that could compel our ISPs to examine our content, restrict how we use our software, and could even have implications for building controls into our hardware, we're going to have to do it ourselves. Will Americans pour into the streets to protect our rights, like the Polish did by the thousands in January? Or will we let them go?

Quinn Norton writes about copyright for Wired News and other publications.

Lucid Demos External Graphics

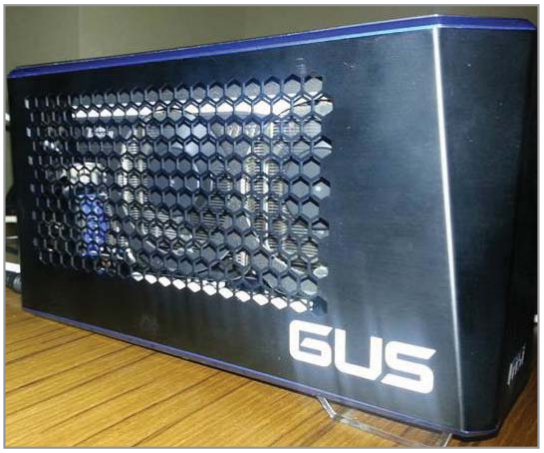

External graphics for notebooks has been a long experiment in failure, but LucidLogix thinks it might finally have a game-changer.

The company says it is working on a graphics solution that would be able to seamlessly use an external graphics card to run games that wouldn't normally run very well on a notebook with integrated graphics.

As you can tell by the name of LucidLogix' product—Thunderbolt External Graphics—you will need a Thunderbolt port, which would be used to pass data to an external box. Inside that box would be a standard PCIe graphics board powered by an external brick. Such a technology would give even the thinnest Ultrabooks with integrated graphics some teeth in more intense games.

LucidLogix' expertise in this area comes from its Virtu technology that's used by many motherboards to run discrete graphics over the integrated graphics. The company didn't announce a time frame for product availability, but one thing clearly stands in the way: Very few notebooks sport Thunderbolt ports today. **–GU**

Report: Minimum PCs Are Stressful

A study commissioned by memory-maker Crucial has revealed something we've known all along—people aren't satisfied with their PCs' performance, and it's stressing them out. According to the study, 94 percent of those surveyed have been unhappy with their rigs recently, with complaints ranging from slow-loading programs to system crashes, unresponsiveness, and memory issues. Respondents went so far as to equate the level of stress their PCs cause them to dealing with airport security, financial problems, doing their taxes, and arguing with their spouses. It should also be noted that although 94 percent is a high number, this study was conducted before the Windows 8 "Metro" UI was in wide circulation.

The study also revealed a mildly shocking (to us) statistic—that 63 percent of computer owners have never even attempted to install memory in their computers by themselves, even though we've run like 20 features on this very subject over the years. **–JN**

Rumor: New AMD Performance GPU on the Way

Now that AMD has finished launching its 7000-series GPUs, info is bubbling to the surface about its upcoming 8000-series cards, which so far includes a flagship "Tenerife" card with 3GB DDR5 memory and a 384-bit memory bus, and two midrange cards dubbed HD 8870 and HD 8850. The high-end HD 8970 will supposedly be more than twice as powerful as the HD 7970 and will feature 2,304 stream processors with a TDP of 250W. The midrange cards will reportedly offer performance equivalent to Nvidia's flagship GTX 680 and GTX 670, at roughly half the price ($280 and $200, respectively). **–JN**

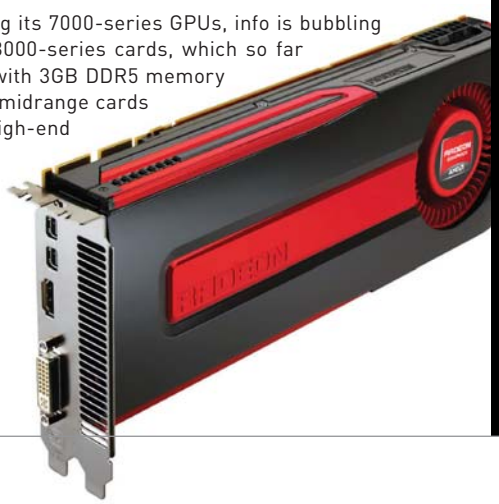

THE LIST

NINE DEAD AND DYING TECHNOLOGIES

9 OPTICAL DRIVES

You've complained that the optical drive isn't dead, and we'd agree—but it's being read its last rites and, we suspect, seeing a long tunnel of light leading to where ISA and VL Bus are waiting for it.

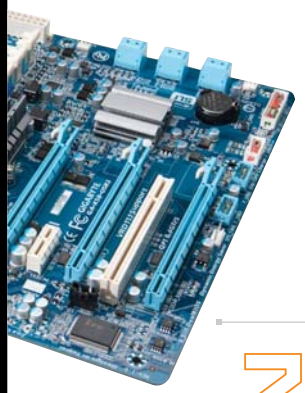

8 PCI

For Pete's sake, what are you even doing hanging around here? Go, already! And take your sidekick PS/2 with you!

7 PS/2

You're *still* here?

6 FIREWIRE

We say good day to this ill-conceived, port-blowing tech. Yup, superior engineering, indeed.

5 4/3 ASPECT RATIO

We wish this taller aspect ratio hadn't gone away in notebooks, as it works better for document handling. But alas, 16:9 is the standard.

4 START MENU

It's official, and no amount of whining will make Microsoft put the Start Menu back in place. So, adiós SM, we'll remember you fondly.

3 SOUNDCARD

The DACs are cleaner, the algorithms more advanced, and more importantly, they sound better. But you don't care, or you just plug in a USB gaming headset anyway. Sad.

2 ETHERNET

We hardly believe it, but with Ultrabook popularity and ever-increasing Wi-Fi speeds, it won't be long before the RJ45 joins RJ11 in the dust bin of history.

1 KEYBOARD AND MOUSE

If you just tried to zoom in on this page by pinching, you can see where we're going with this. The world is getting all touchy, feely and one day the mouse and keyboard will only be used for—gaming.

HEAD TO

BY BRAD CHACOS

XBMC vs. Plex

Attention, would-be cord cutters: If you're going to tell the cable man to shove it, you're going to want a full-featured media center app to make browsing your digital movies, music, and pictures as pretty and painless as possible. Two of the top no-cost contenders are the open-source XBMC and Plex, a partly proprietary fork of XBMC that focuses on streaming media to multiple devices. Which is the blockbuster and which is the dud? Let's find out.

Round 1: Setup

XBMC identified the music, movies, and TV shows in our massive 200GB collection in minutes, quickly dishing out accurate file details, episode summaries, and album/show art with next to no fuss. Plex, on the other hand, was a nightmare. During our initial setup of the Plex Media Server we ran into a crippling bug that wouldn't let us add *anything*. After stumbling around the Plex forums for two hours we managed to fix the issue.

Plex's plodding browser-based media manager took over three additional hours to scan our media, and when it was done, the video library was full of incorrect information. Note: Before you scan your library, rename your media according to specific (yet for some reason, unadvertised) Plex conventions (bit.ly/UFbNjA). Also note: Plex's newbie documentation sucks.

Winner: XBMC

Round 2: Device Support

Both services offer fairly robust desktop PC support, with full-blown Windows, Linux, and Mac offerings. XBMC also supports iOS devices and Apple TV, but they need to be jailbroken, while an Android app is currently in its beginning stages. A Raspberry Pi fork, Raspbmc, is also available.

You'll want to install Plex's Media Server on a central PC or server, but the service also has client apps available for a wide array of devices. There are $5 Android, iOS, and Windows Phone apps and you can find Plex clients baked into Google TV, the Roku Channel Store, and various LG and Samsung products. Plex also supports basic DLNA streaming to the PS3, the Xbox 360, and WD TV Live.

Winner: Plex

Round 3: Multi-device Setups

This round is no contest. The entire Plex ecosystem is based around a central media server that stores your media library and dishes it out to the various Plex clients on demand. Any changes you make to files in the Plex Media Server immediately show up in the Plex clients, and the media server transcodes video on the fly to ensure it can play on your mobile gadgets and connected devices. It's wonderful—once you get it running.

While Plex gets by running thin clients and a beefy media server, XBMC requires a full install on every machine, and frankly, its network-sharing functions suck in comparison to Plex's. Theoretically, XBMC's open-source configurability means you could get it to do everything Plex does, but it would be a heck of a headache.

Winner: Plex

Round 4: Interface

Plex is a fork of XBMC, so it's no surprise that the Plex Media Player (not the server) bears a striking resemblance to XBMC's vertical interface. Both are very pretty, functional, and sport a great 10-foot interface that's perfect for watching videos on your TV.

We love how easy it is to create custom content collections with Plex, which allows you to assign movies, music, and pictures to named groups that are then accessible from Plex's home screen. Though Plex's browser-based server media manager is slow, it's much easier to tweak content details on your laptop than on your TV screen.

On the other hand, XBMC's developer base has whipped up dozens of skins that let you tweak and tune the app's look, and the click-to-install process is remarkably easy to use.

Winner: Tie

HEAD

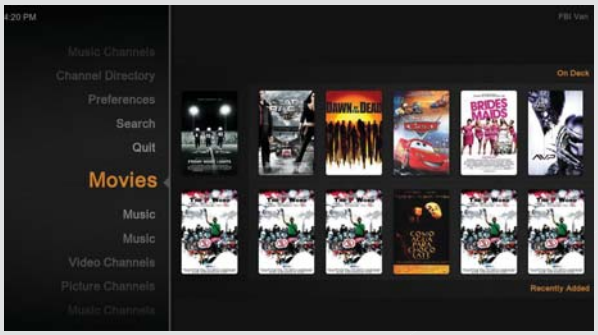

Plex's On Deck function is awesome, but its automatic media identification isn't: All of those videos are incorrectly ID'd.

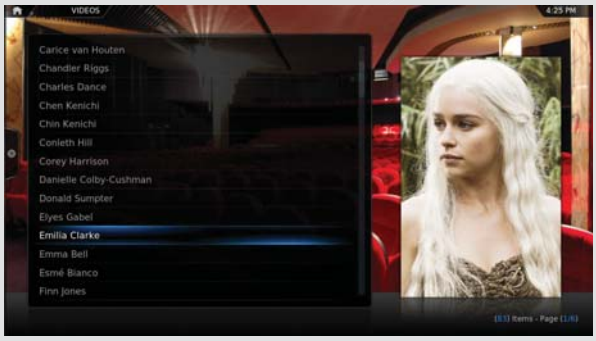

XBMC correctly identified our videos and used the information to provide episode summaries and filters by actor, genre, studio, and more.

Round 5: Extra Features

XBMC is built for HTPCs; it plays physical discs, supports a wide variety of remotes, and offers a deep add-on repository that greatly enhances XBMC's overall functionality. As noted previously, XBMC's media-scraping functionality just plain works, though its DLNA functionality often doesn't.

Plex's deep feature list is skewed more toward people who stream content. Since all media is streamed from the central sever, Plex includes an awesome On Deck function that lets you pause a show on one device, then pick up where you left off on another device. Additionally, Plex's streaming Internet channels are more numerous, stable, and of a generally higher quality than XBMC's video add-ons. You can also send web videos to Plex for future watching... but Plex can't play DVDs

Winner: Plex

And the Winner Is...

In this tale of the tape, the ability to stream content to several devices in an increasingly multiscreen world gives **Plex** the win, though XBMC's DVD and remote support make it the better option for HTPC purists. We got the best of both worlds by installing the PleXMBC add-on (bit.ly/RyzP0V) for XBMC, which lets you stream content from your Plex Media Server to XBMC. Woohoo! The plugin can't handle Plex's Internet channels, though.

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> RAM Errors
> RAIDing Different Drives
> Ivy or Sandy Bridge-E?

Errors on the RAM-page?

Five years ago, I built a computer that has performed without problems. I put in an Intel DP965LT 775 motherboard, Intel Core 2 Duo E6700 processor, EVGA GeForce 8800 GTS videocard, two Western Digital 500GB drives, two Toshiba DVD drives, and two sticks of Crucial 1GB DDR2/667 memory. By today's standards, the machine is ancient, but it has been fine for my needs.

Earlier this year, I decided to boost my memory by adding two 2GB Crucial DDR2/667 sticks to bring my system memory from 2GB to 6GB. The system showed that I had 6GB of memory. Then I started to notice that even though it was noticeably faster, I was having system problems. Thinking my problems might be related to heat, I started running SIW (System Information for Windows) and found the core temperatures to be (in my opinion) high. Most readings are between 58 and 65 C, with an occasional peak of around 80 C. (This is an air-cooled system with stock fans.) I then removed the two new sticks of memory and dropped the system memory back to 2GB. The problems went away, but the temps remained in the 58 to 65 C range. I then took out the original two sticks (2GB total) and put in the two new sticks (4GB total) and the problems returned. Thinking my problem might be the new memory, I ran Memtest86+ overnight twice, with no errors reported the next day. But the problem only occurs with the two new sticks.

I'm running Windows 7 and the problems I'm encountering are that the MMC or Windows Explorer or Windows Desktop stops working, or I can't start Outlook 2010. This happens within one to four hours after a cold boot. Most times it's within the first hour. If I then do a reboot, there are no more problems until the next cold boot. I have run the system many days after a reboot without an error. My question is, do you have any ideas what is wrong and how I can fix this problem? I'm guessing it's not a temperature issue since it will run without problems for days after a reboot.

—Howard Lewis

THE DOCTOR RESPONDS:

Although most motherboards should be fine with any RAM of the right form factor, speed, and capacity, you should always check the RAM you use against the motherboard's Qualified Vendor List if you're having memory issues. The Doctor isn't sure what model of Crucial DDR2/667 you're using, but your motherboard's QVL does list several 2GB DIMMs from that manufacturer at that clock speed, so they should work. Intel also specs the board to run up to 8GB of RAM, so capacity should not be an issue.

Given your trouble, though, it sounds like the new RAM is the problem. You haven't said whether you're running 32-bit or 64-bit Windows, but given the age of your system it's likely you're running 32-bit, which can only address a maximum of 4GB RAM, with some of it consumed by the address space for devices. Running 6GB with 32-bit Windows 7 is bit of a waste but shouldn't be the problem.

You've tried running the old RAM alongside the new RAM, as well as just the new RAM, but you should also test each individual new piece of RAM to determine if just one stick is bad. If the RAM checks out, it's very possible that the BIOS on the board isn't reading the RAM's Serial Presence Detect chip correctly. The SPD tells the board what frequency, voltage, and timing to run the modules at. Normally when you mix and match RAM, the board will set all the RAM to the speed of the slowest RAM you have installed, but it's possible it's just not reading it correctly. The Doc has seen that numerous times on boards of the LGA775 vintage, but usually on motherboards that play RAM-timing settings fast and loose. Intel boards tend to err on the side of stability, but an old BIOS might misconfigure the RAM. Consider updating the BIOS—with only the two old DIMMs in place. If you are still having issues, go into the BIOS on your board and manually set the RAM timing, voltage, and frequency.

As for your temperatures, they look within reason for an E6700 with a stock cooler.

RAIDing Dissimilar Drives

Is it OK to set two SSDs in RAID 0 if they're using different controllers? I have an older Vertex 2 120GB SSD using the SF-1200 controller and I would like to add a second, more for the extra space for my boot drive than for the performance increase. Would there be any issue using a newer SF-2281 SSD for this? Some people say you should always find an identical component, some say it doesn't matter. I'm looking for a more informed opinion.

—Ethan K.

THE DOCTOR RESPONDS:

Technically, you can do it. Should you? Nope. Different SSD controllers utilize their storage in different ways and have different speeds, so we recommend only RAIDing drives with the same model,

↘ submit your questions to: doctor@maximumpc.com

capacity, and firmware. Since any drive failures in RAID 0 cause you to lose the whole array, why add more potential problems? In your case, rather than buying another old Vertex 2, we recommend you get a higher-capacity 6Gb/s SATA drive to use as your boot drive and put the Vertex 2 to use as a games or programs drive.

LGA1155 or LGA2011?

Doc, my box currently uses a 2.5GHz Core 2 Duo with 8GB of RAM and RAID 0 1TB hard drives. I want to upgrade, but should I build on LGA1155 or LGA2011?

I play games at the highest settings and do some photo and video editing, as well. I expect to up the video editing soon, as I'm expecting a baby boy in December. Ideally, I would like something I could upgrade later with a new processor and videocard. My budget is under $2,000.

—David McGuan

THE DOCTOR RESPONDS: First, congratulations. The birth of your first child will be the happiest day of your life. Second, you seem like a good candidate for LGA2011 and with your budget of $2,000, you can build a very nice rig, like the Performance PC from our Blueprints section (page 95). This lets you drop in a hexa-core Sandy Bridge-E chip down the road or even an Ivy Bridge-E chip when they are expected to hit the market early next year. The GTX 670 is a heck of a card for gaming, too. The Doc will say that the current LGA1155 Ivy Bridge chips are very good processors and quite capable of handling video editing chores. If you actually split your gaming/content creation at around, say, 70/30 percent, you would be absolutely fine building a machine using a quad-core Ivy Bridge processor. The only real disadvantage to building on LGA1155 is that next year Intel is expected to introduce the CPU code-named Haswell on a new LGA1150 socket, which will be incompatible with LGA1155. As the Doc said, at least LGA2011 gives you one more processor upgrade in its life. But will you really want an Ivy Bridge-E chip—even with six cores—if Haswell can slice and dice it? Ah, the eternal tech treadmill....

A Sandy Bridge-E machine like our Deluxe rig is a good sub-$2,000 rig for gaming and video editing, and is upgradeable to hexa-core later.

GET THE BEST BA FOR THE

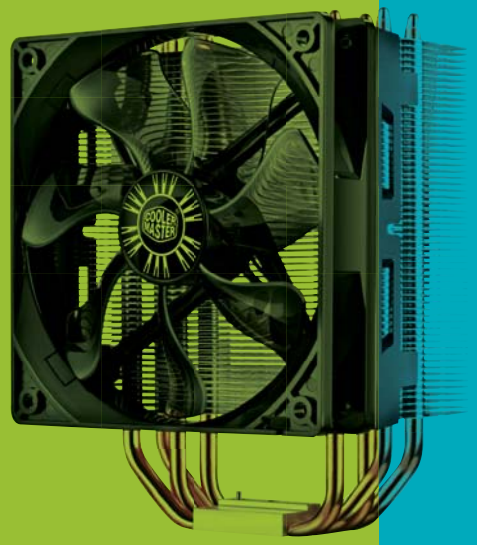

WE HIGHLIGHT THE HARDWARE THAT GETS YOU THE MOST PERFORMANCE PER DOLLAR SPENT—NOW WHO CAN ARGUE WITH THAT?

BY THE MAXIMUM PC STAFF

NG BUCK

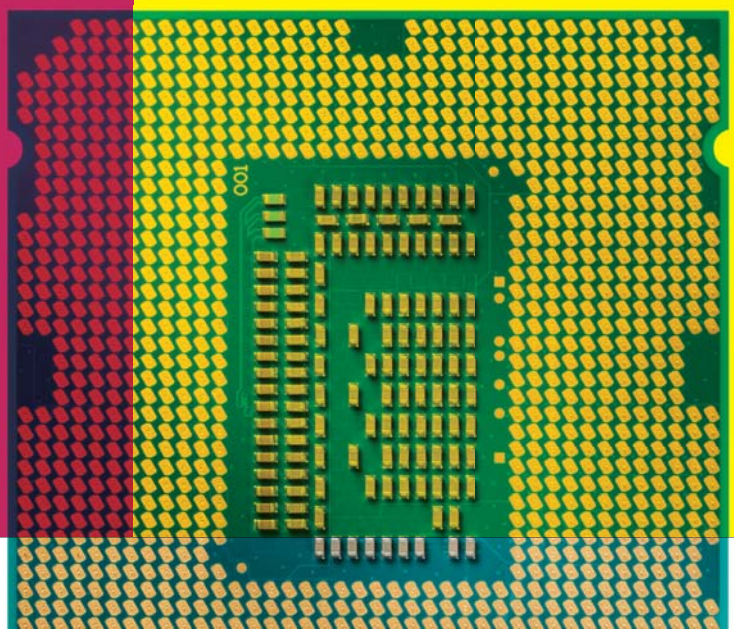

We all know that, generally speaking, buying the newest top-end part gets you the most performance. But in most cases, the premium you pay for that part covers a whole lot of other stuff as well that has no bearing on frame rates or video encoding times. We're talking about the added cost of covering research and development, product marketing, lower production yields, etc. That high price also includes a vanity tax, if you will—the extra charge incurred by folks who simply want to have the latest hardware, hot off the fab, for bragging rights.

Splashing out on the very latest gear is all well and good, if you're rolling in cabbage. But for most of us, the job of choosing what part to buy is much more nuanced. The object is to find the hardware that offers the greatest value—the best price-to-performance ratio.

That's our objective here—to examine each of the major PC component categories, run the relevant benchmarks where warranted, and zero in on the coveted sweet spot. What we've come up with is a list of parts that are capable of meeting a typical enthusiast user's needs for the best price possible.

CPU

Intel Core i5-3570K

PAY FOR THE THREADS YOU NEED, NOT THE THREADS YOU WANT

Make no mistake, we love cores and threads. But when you start talking about bang for the buck and sweet spots, eight-core and six-core CPUs don't fit into the equation (well, unless you consider AMD's Bulldozer, but that's another matter). As power users who like to regularly drop in new CPUs, we have a personal fondness for the LGA2011 platform, but frankly, LGA2011 carries a price premium that the vast majority of users don't need to pay.

Really, when you talk about bang for the buck, it very much gets into the complicated question of just what kind of bang you're after. Are you a 100 percent gamer (a rare duck, in our book) or someone who spends most of his or her time pushing pixels professionally?

Since you pay for every thread in your system, you want to make sure they're actually working for you. Take a gander at our benchmark chart, which compares a six-core Core i7-3930K to the quad-core Core i7-3820. To see what impact threads have on the benchmarks, we turned Hyper-Threading on and off for both processors. The value of the threads are there—but only if you have multithreaded apps that can use them. For the most part, today's apps are optimized for quad-core or lower.

Based on our thread experiment, we believe we can make a case for Intel's $230 quad-core 3.4GHz Core i5-3570K chip sans Hyper-Threading. Ivy Bridge offers at least a 10 percent performance advantage over Sandy Bridge and the 3570K is unlocked. Keep in mind, if you run thread-heavy apps such as encoding or 3D rendering, you should pay for Hyper-Threading or move up to a six-core part, but the sweet spot today remains four cores.

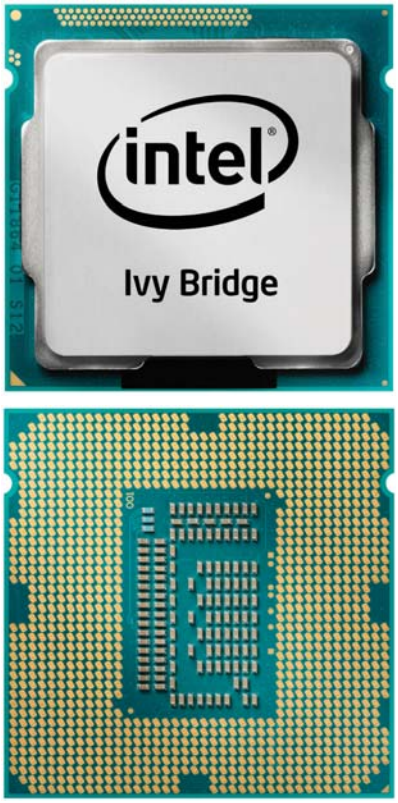

This quad part offers Ivy Bridge performance and is unlocked for overclocking.

WHERE'S AMD?

We know, why no AMD? To tell you the truth, we love AMD parts. The K6-2 was a spry chip. The Athlon saw David slay Goliath, and Athlon XP repeatedly made a mockery of the Pentium 4 during its run. The Athlon 64 and Athlon 64 X2 were such stunningly fast chips that we didn't see an Intel-based system for review for more than a year. See, we recognize AMD's glory days. But today, it's a different story. Intel firmly controls the high ground, middle ground, and would probably have run AMD out of the low ground if it wasn't so fixated on ARM-based CPUs today. Don't get us wrong, an AMD chip would be great in an entry-level box, but when the company's fastest eight-core processor can barely out-box Intel's last-generation quad-core (without Hyper-Threading enabled, mind you) it's hard to see a place for AMD processors except in the entry-level arena today.

BENCHMARKS

	3.4GHz Core i7-3930K with Hyper-Threading Enabled	3.4GHz Core i7-3930K with Hyper-Threading Disabled	3.6GHz Core i7-3820 with Hyper-Threading Enabled	3.6GHz Core i7-3820 with Hyper-Threading Disabled
Thread Count	12	6	8	4
Price	$583	$583	$294	$294
3DMark 11	X5,794	X5,703	X5,730	X5,669
3DMark 11 Physics	11,299	8,974	8,901	7,090
x264 HD 5.0 Pass 1 (fps)	85.8	51.9	62.8	51.6
x264 HD 5.0 Pass 2 (fps)	19	15.6	13.4	11.2
ProShow Producer 5.0 (sec)	1,553	1,500	1,565	1,546
STALKER Day (fps)	181	182	184.2	184.9
STALKER Night (fps)	206.5	211	206.5	206.5
STALKER Rain (fps)	240.7	238	240.7	236.4
STALKER SunShafts (fps)	159.2	158	161.2	160.4
Stitch.EFx (sec)	906	969	962	1,057
Cinebench 11.5 (score)	10.2	8.35	7.72	5.85
Valve Particle Benchmark (fps)	249	186	181	146
7Zip (MIPS)	29,748	21,655	21,122	15,767

We used an Asus Sabertooth X79 motherboard, 8GB of DDR3/1333, a GeForce GTX 690, and OCZ Vertex 3 SSD, and 64-bit Windows 7 Professional for testing both CPU configurations.

VIDEOCARD

GeForce GTX 660

IT'S A TIGHT RACE BUT THIS RELATIVE NEWCOMER TAKES THE CROWN

When you begin to compare price/performance ratios of GPUs at the high end of the scale, you typically see diminishing returns once you go beyond about $300. That's rarefied air up there, where 5–10 additional frames per second will cost you an extra $100 or so. The real battle for frames, frags, minds, and dollars exists below that price point, as the GPUs loitering in the mid-$200 region offer very good performance for literally half the price of the flagship videocards. This is the sweet spot—typically around $250 or so—where you can find 1080p gaming at 60fps for the least amount of money possible.

After examining the cards in this pricing segment, as well the cards in the pricing segment above it, we're granting the brand-new Nvidia GTX 660 best-bang-for-the-buck status. At just $230, this GPU is faster than the slightly more expensive (or same-priced, depending on where you shop) AMD HD 7850 by a respectable margin, and even beats the pricier $280 Radeon HD 7870 in some of our tests, making it the fastest GPU under $250 by a wide margin.

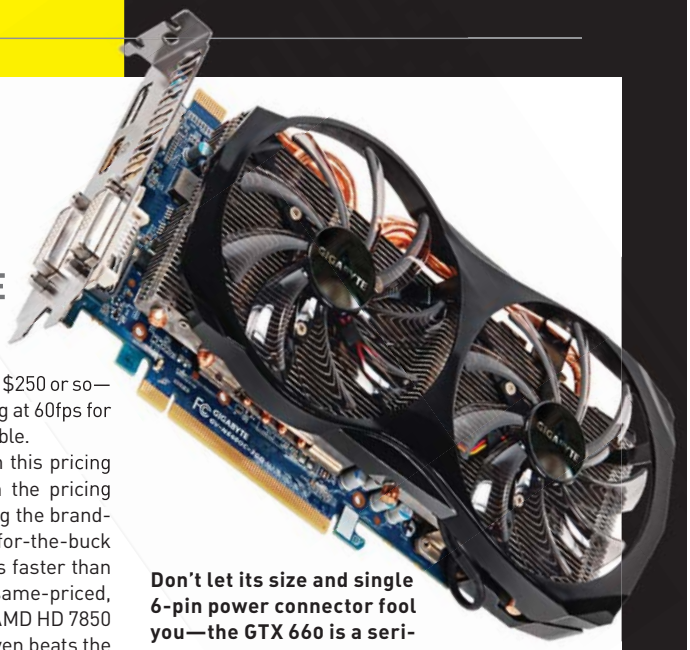

Don't let its size and single 6-pin power connector fool you—the GTX 660 is a serious gaming weapon.

So, why not spend $70 more and get a GTX 660 Ti? Looking at the benchmark chart, you can see that extra money nets you some performance gains, but in our judgment you don't get enough of a return on that investment. The GTX 660 Ti costs roughly 25 percent more than a GTX 660, and yet offers an average of 15 percent performance improvement in most tests. When you tabulate price-per-fps, the GTX 660 also has the advantage over the GTX 660 Ti in the majority of tests we run, making it the value leader. As icing on the cake, you only need one 6-pin connector instead of the two that bigger cards require, so you're saving money in the PSU department, as well.

BENCHMARKS

	Gigabyte GTX 660 OC Version	XFX Radeon HD 7850	EVGA GTX 560 Ti 448	XFX Radeon HD 7870 GHz	MSI GTX 660 Ti Power Edition
Price	$230	$245 (street)	$250	$280 (street)	$310
3DMark 2011 Perf	6,935	6,075	6,295	7,001	**9,118**
3DMark Vantage Perf	27,858	24,584	25,523	27,953	**31,575**
Unigine Heaven 2.5 (fps)	36	31	31.8	33.2	**39.8**
Shogun 2 (1080p, fps)	53.7	47	43.1	62.2	**71.1**
Far Cry 2 / Long (fps)	119	103	112.4	112.7	**133.3**
Dirt 3 (fps)	75.7	50	68.8	82.4	**95.1**
STALKER: CoP DX11 (fps)	42.8	34.7	38.4	42.6	**49.9**
Just Cause 2	60.3	51	55.8	59.6	**68**
Batman: Arkham City (fps)	76	60	64	72	**80**
Metro 2033 (fps)	22	23	22.6	**24.2**	22.3

Best scores are bolded. Our test bed is a 3.3GHz Intel Core i7 3960X Extreme Edition in an Asus P797X Deluxe motherboard with 16GB Corsair DDR3/1600 RAM and a Corsair AX1200 PSU. The OS is 64-bit Windows 7 Ultimate. All games were run at 1900x1200 with 4x AA unless otherwise noted.

MOTHERBOARD

GIGABYTE Z77X-UP4 TH

IT DOESN'T COST MUCH TO GO FROM STRIPPO MOBO TO FEATURE-RICH

When it comes to motherboards, you can seemingly spend as much as you want or as little. So, do you need a bells-and-whistles board like Asus's ultimate Z77 P8Z77-V Premium with its Thunderbolt and PLX chips for four-way GPUs, or is a $90 Z77 strippo board good enough? We find the sweet spot for a cost-conscious power user to be right about $190. Sure, you can get a microATX Z77 board for under $100, but it's usually stripped down to the barest of essentials, with thermals on the VRMs as an afterthought. It's like getting a new car with roll-up windows and manual locks. We think that by stepping up to about $190, you get such worthwhile amenities as Virtu hybrid-graphics support, multi-GPU support (both CrossFireX and SLI), and ports galore. Some good examples of bang-for-the-buck boards are Asus's P8Z77-V and Gigabyte's Z77X-UP4 TH. Both are feature-rich yet don't break the $200 mark.

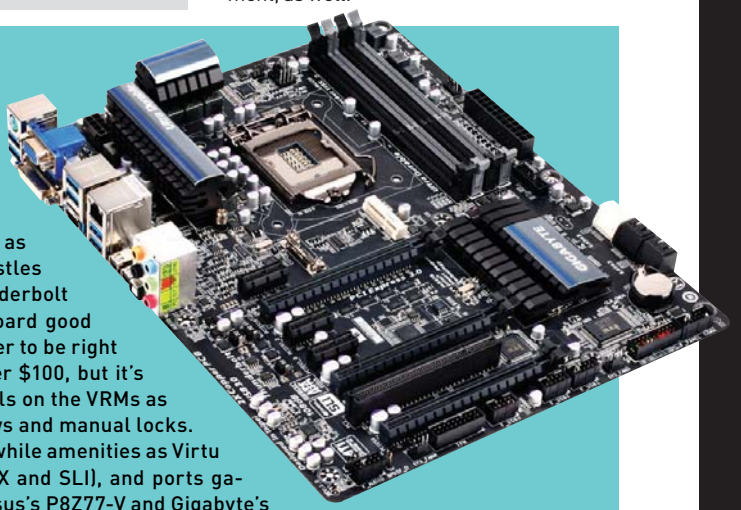

SOLID-STATE DRIVE

Crucial M4 256GB

400MB/S READS, 250MB/S WRITES, LESS THAN $200

Not the newest nor the fastest, but good speeds for a good price.

Right now the sweet spot for solid-state drives is a 6Gb/s SATA SSD in a capacity near 256GB (240GB for a SandForce-based drive). Why 256? 256 is big enough to store your OS, programs, and several games, and many modern controllers are optimized for that capacity—SandForce SF-2281 controllers, for example, have 16 lanes, and a 240GB drive has 16 NAND modules. Best of all, thanks to this year's massive price drops, SSD prices are under a dollar per gigabyte: You can get a 256GB Samsung 830 Series, one of our favorite SSDs, for under $230.

So, which SSD offers the best value? Crucial's M4 SSD is a favorite of system builders for its relatively good performance and low price. It's not *the* fastest 6Gb/s SATA SSD, but it's plenty fast by any standard, and it's attractively priced. For $180 at the time of this writing, you can get a drive with sustained reads over 400MB/s, sustained writes over 250MB/s, and good random-read and -write performance, as well. For $20 more you get the OCZ Vertex 4, with reads and writes in the 440MB/s range and higher random IOPS, and for $50 more than the M4 you can get a Samsung 830 Series drive, with 500MB/s-plus sequential reads, 400MB/s writes, but lower random IOPS—which is fine, unless you're running a really active database server.

The specific drive you get will depend on current pricing, and (as is the case with most components) you can get a great drive for a little more than the cost of a very good drive, but right now we think the Crucial M4 is the sweet spot to beat.

BENCHMARKS

	Crucial M4	OCZ Vertex 4	Samsung 830 Series
Capacity	256GB	240GB	256GB
Price	$180	$200	$230
Price Per Gigabyte	$.70	$.83	$.89
Controller	Marvell 9174	Indilinx Everest 2	Samsung
CrystalDiskMark			
Sustained Reads (MB/s)	404.5	440.9	506.4
Sustained Writes (MB/s)	257.3	446.9	398.5
AS SSD			
4K Reads (IOPS)	5,091	6,632	5,513
4K Writes (IOPS)	13,837	17,159	14,412
Iometer			
4KB Random Writes, QD32	56,087	65,111	35,329

Our current test bed is a 3.1GHz Core i3-2100 processor on an Asus P8 P67 Pro (B3 chipset) running Windows 7 Professional 64-bit. All tests used onboard 6Gb/s SATA ports with latest Intel drivers.

HYBRID DRIVES

A SWEET SPOT UNTO ITSELF?

On the surface it might seem like a hybrid drive is itself a sweet spot, giving you SSD speeds in your most common tasks with the capacity of a standard hard drive. And, well, that's true. There are two basic types of hybrid drives. A true hybrid, like the Seagate Momentus XT, consists of a hard drive with onboard NAND cache and an algorithm that mirrors the most frequently accessed sectors of the hard drive to the cache. The current Momentus XT is a 2.5-inch drive that tops out at 750GB of storage with 8GB of cache and is a good choice for laptop owners who don't mind the $60 price premium over a standard 750GB drive.

The other type of hybrid—roll-your-own—is more accessible for desktop owners: Intel chipsets since Z68 have shipped with Smart Response Technology to allow users to add a cache drive, and many SSD vendors now ship low-capacity drives with Dataplex caching software. This is a good option for people who can't afford an SSD large enough to be its own boot drive, or who don't want the trouble of a reinstall—but there are limitations. SRT can only use cache drives under 64GB and the Dataplex software that ships with most cache SSDs won't work if the system drive has a GPT partition scheme or is over 2.2TB.

Either type of hybrid will speed up the "feel" of Windows—it'll boot faster and frequently accessed programs will launch much faster, but you won't get the massive speed boosts when copying files, and tasks you don't perform often won't feel the benefit at all. It's not quite the best of either world, but it's better than just rolling mechanical.

At 4.6 cents per gigabyte, you can't afford to pass up the Barracuda 3TB.

HARD DRIVE

SEAGATE BARRACUDA 3TB

THE BEST DRIVE IS ALSO THE BEST VALUE

It's weird to say this, but the sweet-spot hard drive is the same as the Best of the Best hard drive. The 3TB Seagate Barracuda is the best combination of price, performance, and capacity, whether you're using it as your only drive or as the backup for a boot SSD. Its sequential read and write speeds of over 150MB/s make it the fastest 7,200rpm drive we've ever tested, and its price of $140 means its per-gigabyte cost is only 4.6 cents, making it the best value on the market right now. By contrast, 2TB "Green" drives, which spin at around 5,900rpm, are around $120 from WD and Seagate right now. Why pay more per gigabyte for slower storage? Cost-per-gigabyte goes way up as size goes down, too, especially for any drive under 1TB.

If your budget can't stretch to $140 and you don't have an SSD, you'll want to prioritize performance over capacity. Get a 2TB Barracuda (make sure it's one of the two-platter 7,200rpm ones with 64MB of cache). At $100, that's 5 cents per gigabyte. 2TB WD Caviar Black drives, on the other hand, are over $200 at press time.

If you do have an SSD and you only have, say, $85 or $90, get the 1.5TB Barracuda Green from NCIX US. It's one of the few drives over 1TB that are under $100 right now, though we hope that changes as the industry recovers.

Whether Seagate is engaged in a price war or it just fared better in the Thailand floods of last year, its drive prices are unbeatable at press time, and the 3TB Barracuda is by far the most screaming deal.

SPECIFICATIONS

	Price	Price Per Gigabyte	Spin Speed (rpm)	Cache (MB)
3TB Seagate Barracuda	$140	0.047	7,200	64
1.5TB Seagate Barracuda Green	$70	0.047	5,900	64
3TB WD Caviar Green	$148	0.049	5,400	64
2TB Seagate Barracuda	$100	0.050	7,200	64
2TB Seagate Barracuda Green	$100	0.050	5,900	64
1.5TB Seagate Barracuda	$82	0.055	7,200	32
2TB WD Caviar Green	$110	0.055	5,400	64
3TB HGST Deskstar	$200	0.068	7,200	64
1TB WD Caviar Blue	$83	0.083	7,200	32
1TB Seagate Barracuda	$86	0.086	7,200	64
1TB WD Caviar Green	$89	0.089	5,400	32

CASE

NZXT PHANTOM 410

THERE'S NO NEED TO PAY MORE FOR A SOLID ENCLOSURE

There are people who say that the case doesn't matter, that the sweet spot for cases is "as cheap as you can possibly get." These people are bad and they should feel bad. The sweet spot for cases is right around $100. At $100, you can get a great mid-tower or full-tower case with solid build quality, plenty of fans or fan mounting options, front-panel USB 3.0, toolless drive bays, and a nice paint job. Below that you'll sacrifice build quality or looks, and above that you start getting fancy, with water-cooling mounts, support for giant motherboards or a half-dozen fans, or premium materials, but $100 will get you a case with good air cooling that you can be proud of.

Our favorite cases at $100 are the Fractal Design Define R4 and the NZXT Phantom 410. The Define R4 is the more flexible of the two. It's more of a silent PC by design, with its acoustic foam on the inside of the side panels and its unused fan holes covered with sound-damping material until removed to add fans. The Phantom 410 looks stunning, especially in red or gray, and comes with much more fan support and more fans, flat out.

These aren't the only cases at $100 and below that are worthy of consideration, but they're two of our favorite mid-towers and they both happen to be around the $100 sweet spot.

8GB is truly the new 4GB when it comes to RAM.

RAM

8GB DDR3/1600 OR DDR3/1866

THE SWEET PLACE TO BE IS IN DUAL-CHANNEL MODE

How much RAM do you need, how many channels, and what speed? From our experience, and based on today's RAM prices and capacities, we recommend 8GB of DDR3/1600. Everyone knows that 2GB is not enough RAM today. And while 4GB is OK, why stop there when 8GB isn't much more? At press time, 8GB of DDR3/1600 fetched roughly $40 while 4GB is about $30.

Should you buy more RAM than that? Only you can answer that question. If your applications are RAM hogs or you tend to multitask, additional RAM can't hurt.

There's one more pesky problem, though: How many memory channels do you *really* need? Is it worth paying the extra surcharge for a quad-channel setup or is dual-channel enough? To find out, we took our trusty LGA2011-based zero-point, set it to bone-stock clock speed and tested it with 8GB of DDR3/1333 in quad-channel, dual-channel, and, well, single-channel, using various benchmarks. Some of the benchmarks are GPU-limited and others are CPU-limited to make a point that memory bandwidth won't make a difference in many of your applications. Stepping from quad to dual doesn't hurt much, but going to single will definitely take a bite out of performance, although not as much as you would suspect. In fact, upping the memory speed from DDR3/1333 to DDR3/1866 actually erased the difference in the apps that seem to like bandwidth the most: encoding.

Why does memory bandwidth make nary a difference in the vast majority of apps? Thank the huge caches in today's CPUs. The only time you really want to go all-out on memory bandwidth is when you're intending to game with integrated graphics.

BENCHMARKS

	Quad-Channel	Dual-Channel	Single-Channel
3DMark 11 overall	X5,794	X5,764	X5,634
x264 5.0 HD Pass 1 (fps)	85.8	83.4	74.4
x264 5.0 HD Pass 2 (fps)	19	18.9	18.8
SiSoft Sandra (GB/s)	34.2	17.6	9.0
Stitch.Efx (sec)	906	898	919
ProShow Producer 5.0 (sec)	1,553	1,586	1,732
Cinebench 11.5	10.2	10.2	10.02

We used an Asus Sabertooth X79 motherboard, 8GB of DDR3/1333, GeForce GTX 690, OCZ Vertex 3 SSD, and 64-bit Windows 7 Professional for testing both CPU configurations.

AIR COOLER

COOLER MASTER HYPER 212 EVO

THE BEST COOLER-FOR-THE-COIN WE'VE EVER TESTED

There are many unsolved mysteries in life, such as who built the pyramids, where socks go during laundry, and how is it possible that the Cooler Master Hyper 212 Evo costs only $35. It just doesn't make any sense. This simple cooling device has bested the majority of our skyscraper coolers since it was reviewed in January 2012, and yet it has cost less than almost every single one of its competitors, making it a rare gem in the world of PC components.

What about water cooling? We're not going to go there, simply because even though the all-in-one designs have radically simplified the installation and maintenance of cooling loops, their price-to-performance ratio is still not as good as what's available from a solid air cooler. In fact, in some cases the bigger air coolers perform better than their liquid brethren while costing less, so the sweet spot for now is definitely on dry land.

BENCHMARKS

	CM Hyper 212 Evo	Xigmatek Aegir	Noctua DH-14	Corsair H100 LCS
Price	$35	$60	$85	$115
Ambient Air	23.8	24.8	25.2	25
Idle Temperature	36.2	33.8	34.1	33.2
Burn Temperature	74	**71**	72.3	69.2
Temp Delta (Burn minus idle)	37.8	37.2	41.1	**36**

All temperatures in degrees Celsius. Best scores bolded. All tests performed using an Intel Core i7-3960 at 4.2GHz, on an Asus P9X79 Deluxe motherboard with 16GB DDR3/1600, in a Thermaltake Level 10 GT with stock fans set to High.

PSU

CORSAIR CX600

BUY JUST ENOUGH PSU FOR YOUR NEEDS

Our normal advice for an enthusiast constructing a system is to think long term, overbuild because you don't know what parts your machine will end up with, and consider how hard your PSU is working on those sweltering days of summer. And if we were making a recommendation for someone using our guidelines, we'd say buy a Corsair TX 750 V2 or Thermaltake Smart 730. Both are about $100, give you a nice long five-year warranty, and offer four 6+2 connectors for future beefy-GPU support.

But this story is about bang for the buck, and with today's sweet-spot GPUs offering pretty fantastic power envelopes, we surprised even ourselves by stepping down in PSUs. Corsair's $70 CX600 V2 provides 600 watts and two 6+2 GPU power connectors. The PSU has plenty of power to run a single-GPU machine, no problem, and even SLI if you keep the system properly cooled. Be aware, if you plan to vary from our sweet-spot GPU, the TX 750 V2 is recommended instead.

BEST BOX FOR THE BUCK

WHEN YOU PUT ALL THESE GREAT-VALUE PARTS TOGETHER, YOU END UP WITH ONE SWEET PC

1. The NZXT Phantom 410 easily accommodates all our parts, with room to spare for expansion. The Gunmetal paint job looks hella tough, too.

2. With an Intel Core i5-3570K onboard and the cooling prowess of the Cooler Master 212 Evo, you'd be a fool not to overclock.

3. The GTX 660 lets us connect to a 1080p panel and get 60fps with all settings maxed out.

4. Gigabyte's Z77X-UP4 is SLI-friendly, so commence saving your shekels for that second card.

Proof positive you can get a lot of bang for 1,200 bucks.
Cost reflects NewEgg and Amazon prices at press time.

SWEET SPOT HALL OF FAME

THEY CAPTURED OUR HEARTS WITH A COMBO OF KICK-ASS PERFORMANCE AND KILLER PRICE

Image Credit: Wikipedia

CELERON 300A

At the Bang-for-the-Buckers reunion, there will always be a seat reserved for Intel's Celeron 300A. The "Celery" 300A's claim to fame was an almost guaranteed overclock from 300MHz to 450MHz. You scoff today, but that difference let you get performance similar to a $670 450MHz Pentium II, but for only $150.

ANTEC 900

You're probably spoiled by today's $100 enclosures, but years ago, all a hundy would get you was a razor-blade-shaped box made from recycled roller skates. The P900 gave you incredible bang for the buck and nearly a full set of fans instead of the typical 10-exhaust-ports-but-just-one-fan routine.

NVIDIA GEFORCE 8800 GT

How great was the $250 GeForce 8800 GT? It was faster than the next card up the ladder and damn near as fast as the one a rung above that. It was so popular that Nvidia will still refer to a new card in its lineup as the "GeForce 8800 GT" of the litter.

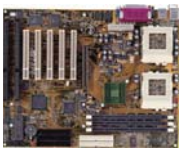

ABIT BP6

Dual-proc men (and women) have always gotten gouged on price—you literally had to pay for the privilege of running two processors. Abit's BP6 shattered that price lock by giving you a dual-proc board that let you run and overclock two Socket 370 Celerons—a config the Celery was never supposed to support!

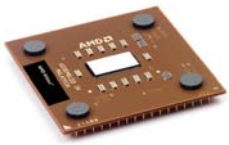

AMD CPUS

We're going to break form here by declaring AMD CPUs bang-for-the-buckers. Which one? Just pick any, because nearly all of them qualify. Yes, AMD is out of the hunt right now, but in general its chips have always been incredibly price competitive and have kept Intel's prices far more reasonable.

NVIDIA GEFORCE 4 TI 4200

Before the GeForce 8800 GT, there was the GeForce 4 Ti 4200. The 4200's hallmark was that despite Nvidia intentionally Nerfing the card to keep it from competing with higher-end cards, card vendors ignored the clock-speed directive and sold boat loads of them, all but killing its sibling, the Ti 4400.

COOLER MASTER HYPER 212 PLUS

Cooler Master's famous Hyper 212 Plus proved that an excellent cooler didn't have to cost a claw and a tendril. The amazing thing is that the $25 Hyper 212 Plus was as good as 90 percent of the pack that cost twice as much, and it came pretty darn close to the best performers, too.

INTEL CORE 2 QUAD Q6600

The sweet spot in CPUs always seems to be near the $250 mark. At $266, the 2.4GHz Core 2 Quad Q6600 was exactly the same as the $1,000 2.93GHz Core 2 Extreme QX6800 in cache and front-side-bus speed. Oh, and it overclocked just fine, thank you.

THE WORLD'S *FASTEST* COMPUTERS

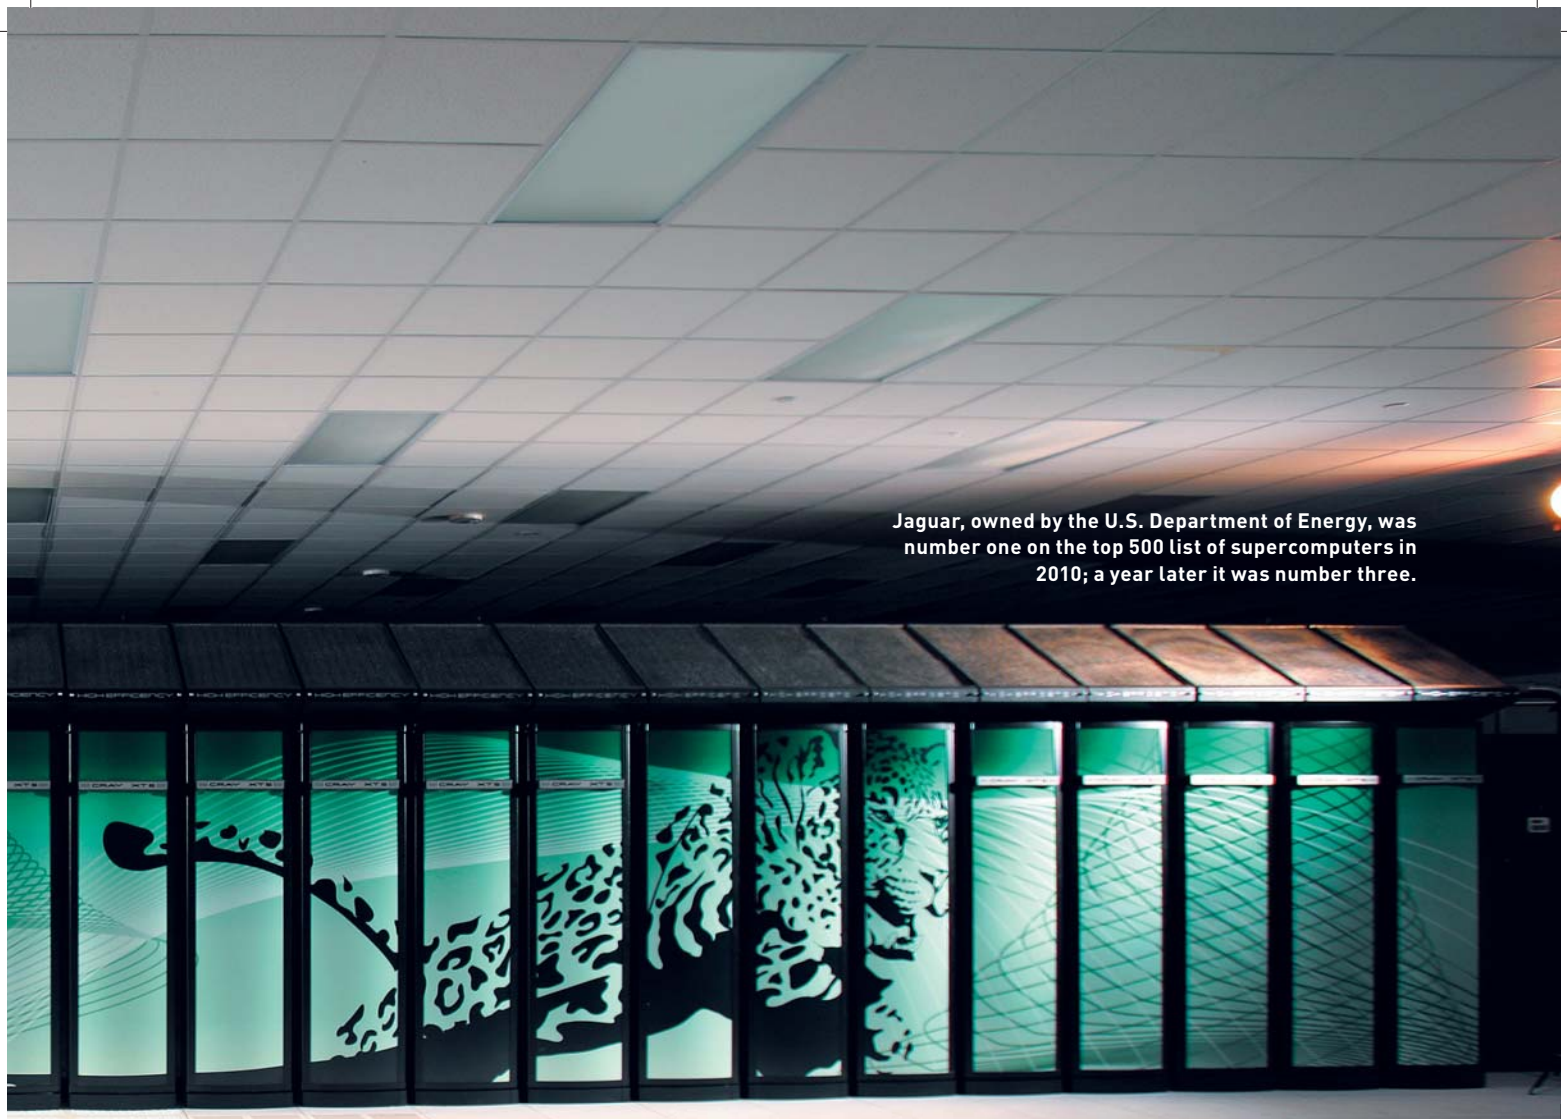

Jaguar, owned by the U.S. Department of Energy, was number one on the top 500 list of supercomputers in 2010; a year later it was number three.

BEHIND THE SCENES OF THE TECHNOLOGICAL ARMS RACE, AND THE CONSEQUENCES OF FALLING BEHIND IN GLOBAL SUPERCOMPUTING

By Jamie Middleton

Forget the space race to Mars—the real technological arms race is taking place in super-cooled bunkers across the globe as countries compete to take the crown of fastest supercomputer on Earth.

Supercomputers are big business. They provide the power to dramatically reduce research time, to accurately simulate and model real-life systems, and to help design products perfectly suited for their environment. Nations that invest in supercomputers can quickly reap the benefits in terms of technological advances, not to mention the kudos of being at the very forefront of technologies that are already changing the world we live in. There is a definite pride in being the best in the supercomputing sphere—a pride the United States enjoyed since 2004. It was, therefore, big news when China's Tianhe-1A, based at the National Supercomputer Center in Tianjin, took the supercomputing crown from America's Jaguar system at Oak Ridge National Laboratory in November 2010.

For years, the United States maintained its status as the leader in the field, and despite losing the top spot, it could take some solace from the fact that the Tianhe-1A was built from U.S. technology—specifically Nvidia GPUs. That all changed when Japan knocked the United States down to third place in June 2011 with its K Computer, built from Fujitsu's SPARC64 VIIIfx processors.

It had been seven years since Japan last held the title of fastest supercomputer on Earth (in 2004 with NEC's Earth Simulator), and when the K Computer took pole position, it did so in style. It harnessed more processing power than the next five supercomputers combined, despite the fact it wasn't even completed yet. It was commissioned by the Japanese Ministry of Education, Culture, Sports, Science, and Technology with the aim of breaking the 10-petaflop barrier (the K in its name refers to the Japanese word "kei," which means 10 quadrillion). In June 2011 the K Computer contained just 672 computer racks of its planned 800, but with 68,544 CPUs operational it still scored a LINPACK benchmark of 8.162 petaflops (quadrillion floating-point operations per second). By November 2011 the K Computer reached its goal, with 864 server racks and 88,000 CPUs it achieved a LINPACK benchmark of 10.51 petaflop/s (93 percent of its theoretical peak speed of 11.28 petaflops).

A POWER SHIFT

Suddenly, the East's dependence on Western technology looked shaky and America's dominance was under threat. Despite still owning the majority of the supercomputers on the top 500 list (255 in June 2011, to be exact), America's claim to being the supercomputer power was further questioned in October 2011 with the announcement of China's new supercomputer, the Sunway BlueLight MPP. This wasn't because it was the most powerful—it's thought to be about the 15th fastest machine in the world—but because it was built entirely from Chinese-made processors, making it the first Chinese supercomputer that doesn't use Intel or AMD chips. The arrival of the Sunway realizes China's dream to sever its reliance on U.S. technology. This isn't just rhetoric, but a stated ambition of the Chinese government. What's more, the Sunway's small size and relative energy efficiency compared to other supercomputers in the top 20 have surprised top computer scientists, and some commenters believe this is a sign that China is about to leapfrog its competitors with the next iterations of its own technology.

China's Sunway BlueLight MPP is built from ShenWei SW-3 1600 chips (16-core, 64-bit MIPS-compatible, or RISC, CPUs) that are thought to be a few generations behind Intel's current chips, but with the large funding base provided by the Chinese government, it's unlikely that new versions of these chips are far away. ShenWei CPUs are believed to be based on China's Loongson/Godson architecture, which China has been working on since 2001. Some people believe this family of chips was reverse-engineered from the U.S.-made DEC Alpha CPU. This may be true, but the nice thing about reverse-engineering products is that it allows you to apply hindsight—to redesign and remove all those hangovers and unnecessary artefacts from previous iterations—hence their better energy efficiency and smaller size.

As of June 2012, the United States had regained dominance on the list with IBM's Sequoia supercomputer, sporting 1,572,864 cores. But prior to that, the U.S. government had already raised questions regarding how the best supercomputer in the world is determined. These supercomputer world rankings in terms of power are all very well, but a lot of people, including President Obama, have pointed out that it isn't what your supercomputer can do, but what you do with it. They argue that supercomputing isn't just about raw power, but how that power is used, how much electricity it consumes, and how fast it is at bringing about real-world solutions to problems.

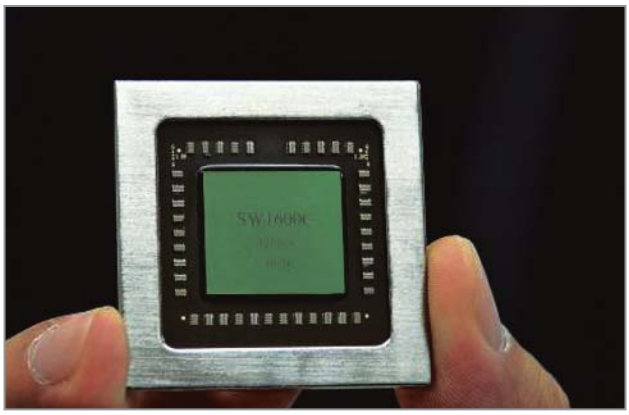

The ShenWei CPU may have been reverse-engineered from the U.S.-made DEC Alpha.

WHO DARES TO JUDGE?

So who judges which supercomputer takes the crown, and what criteria do they use? Most people are happy to be guided by the Top500 List of supercomputers (www.top500.org). The Top500 project was started in 1993, and ever since then has produced a list of the most powerful computing devices twice a year to "provide a reliable basis for tracking and detecting trends in high-performance computing."

Assembled by the University of Mannheim in Germany, the Lawrence Berkeley National Laboratory in California, and the University of Tennessee in Knoxville, Top500 uses the LINPACK benchmark to decide which machine has the best performance. Introduced by Jack Dongarra, the LINPACK benchmark tests the performance of a system by giving it a dense system of linear equations to solve, measured in floating-point operations per second (flop/s)—as in, how many calculations per second it can perform. The higher the flop/s, the better the computer.

It has been agreed that this isn't the best way to measure power. Not only is it measuring only one aspect of the system, but two machines with the same processor can get different results: The performance recorded by the benchmark can be affected by such factors as the load on the system, the accuracy of the clock, compiler options, version of the compiler, size of cache, bandwidth from memory, amount of memory, and so on. This, however, is true of most benchmarks.

The Top500's focus on just one measure led President Obama's Council of Advisors on Science and Technology to argue that the United State's position in supercomputing shouldn't be decided by this single metric. In a report in December 2010, it stated: "While it would be imprudent to allow ourselves to fall significantly behind our peers with respect to scientific performance benchmarks that have demonstrable practical significance, a single-minded focus on maintaining clear superiority in terms of flop/s count is probably not in our national interest."

REAL-WORLD BENEFITS

One real-world benefit of supercomputers is in product design. According to Dongarra, "Supercomputers enable simulation—that is, the numerical computations to understand and predict the behavior of scientifically or technologically important systems—and therefore accelerate the pace of innovation. Simulation has already allowed Cummins to build better diesel engines faster and less expensively; Goodyear to design safer tires much more quickly; Boeing to build more fuel-efficient aircraft; and Procter & Gamble to create better materials for use in its home products. Simulation also accelerates the progress of technologies from laboratory to application. Better computers allow better simulations and

The Cray XE6 uses 8- or 12-core Opteron 6100 processors, allowing for up to 2,304 cores per cabinet.

more confident predictions. The best machines today are 10,000 times faster than those of 15 years ago, and the techniques of simulation for science and national security have been improved."

Another person who believes in the need for supercomputing is Kevin Wohlever, director of Supercomputing Operations at Ohio Supercomputer Center. He has been involved with supercomputing for almost 30 years, working for vendors such as Cray Research, Inc. We asked him whether the supercomputer race between China and the United States is a matter of national pride, or if there's there something more important at play.

"It is both a matter of pride and something deeper. If you look back 10 to 15 years, the same questions were being asked with the Japanese-vs.–United States supercomputer race. There is obvious national pride. There is also vendor pride for the company that designs and builds the fastest computer. There are marketing benefits from being the company or country that has and provides the fastest computer in the world.

"There are also significant benefits from the proper applications running on top-end computers," he adds. "Modelling the financial markets and making trades just that split-second faster than someone else can make a lot of money for someone. There are also implications for everything from weather modelling and designing cars, planes, boats, and trains, to finding new drugs, modelling chemical reactions, and such. The right tools in the right hands can make a big difference in the level of security anyone can feel, whether that happens to be in their job or their life."

SCIENCE'S THIRD LEG

Gerry McCartney, vice president for Information Technology and CIO at Purdue University, also agrees that there needs to be more emphasis on developing supercomputers, and not just to be able to say you are the best: "I would say that high-performance computing—computa-

MAKING RANK: TIME FOR A NEW SCORING SYSTEM

EVEN JACK DONGARRA, the creator of the LINPACK benchmark, which is the current standard for ranking supercomputers, agrees that it isn't the best way to measure power.

Dongarra is keen to start using another benchmark he originally designed for DARPA, called the High Performance Computing Challenge. HPCC completes seven performance tests to measure multiple aspects of a system. As yet, it has not been adopted by the Top500. Another benchmark, the Graph 500 benchmark, has also been introduced by an international team led by Sandia National Laboratories. Its creators describe it as complementary to the one used by the Top500. However, because it produces just one number based on a single aspect of supercomputing, the Graph 500 may be open to the same criticisms as the LINPACK benchmark.

The Top500 project recognizes the limitations in the single algorithm it uses, and it does its best to prevent exploitation. The website carries a disclaimer that the Top500 list only looks at general-purpose machines—ones that can handle a variety of scientific problems. So any system specifically designed to solve the LINPACK benchmark problem or have as its major purpose the goal of a high Top500 ranking will be disqualified.

tional modelling and simulation—is now the 'third leg' of science and engineering, alongside the traditional legs of theory and experimentation.

"So having inferior computing is exactly like having inferior laboratories or research tools. It has a direct relation on the amount, quality, and speed of developments—both basic science and technological developments, which advance knowledge, and science; and technological developments that drive economic growth. There's a lot more at stake than bragging rights."

With this in mind, we asked McCartney, who oversees the biggest campus supercomputer in the world, how important he feels supercomputing is in the political sphere today. "Probably not important enough," he told us. "In the United States, as well as the UK, it's facing the same problem as science, medical, and other types of research generally. The economy is down, budgets are tight nationally, and in the United States, at the state level. We're looking for ways to spend less, not places to spend more. Certainly, we should find ways to leverage efficiencies and spend money wisely. But we need to think long-term in doing so. Developments today, like atomic-scale silicon wires, may not pay off in a big way for a decade. But they will never pay off if you don't have the tools to begin the process."

However, despite these monetary constraints, the United States is not standing still. IBM and Cray are already developing 20-petaflop machines, and despite President Obama's Council of Advisors on Science and Technology stating in December 2010 that, "Engaging in a [supercomputing] 'arms race' could be very costly, and could divert resources away from basic research," America's 2012 budget has earmarked $126 million for the development of exascale supercomputing, the next major milestone in supercomputer research. This compares with the mere $24 million set out specifically for supercomputing in the previous budget. Japan has earmarked $1.3 billion for development of its own exascale computer by 2020.

Japan's K Computer has the same power as a million desktop PCs and cost an estimated 112bn yen to develop.

THE ALTERNATIVES

You can build a supercomputer from solid-state drives, so why not make one from mobile phone chips? By using low-powered Tegra CPUs, typically found powering mobile phones and tablets, the Barcelona Supercomputer Center in Spain hopes to reduce the amount of electricity needed to run its new supercomputer. There's a deal with Nvidia in the offing, and it hopes to produce a high-performance computer by 2014 that runs with a 4- to 10-fold improvement in energy efficiency over the most energy-efficient supercomputers that currently exist.

QUANTUM COMPUTING: WHY DOES IT MATTER?

WE SPOKE TO MARK THOMPSON, senior researcher at the Center for Quantum Photonics at Bristol University, about the attraction of quantum technology for high-performance computing.

"A quantum computer can perform certain tasks extremely efficiently, performing calculations that would be impossible or impractical on any conventional computer, no matter how powerful it was. The computational power of conventional computing technologies is limited by the laws of classical physics, but the world is fundamentally quantum mechanical in nature, and a quantum computer can harness the strange and spooky properties of quantum mechanics—such as entanglement—to process information in completely new and powerful ways."

So how far away is quantum computing from reality? "Compared to the development of the microelectronic industry, we are probably somewhere at the stage between the development of the value amplifier and the first transistor—so we still have a long way to go. It is widely believed that a full-scale quantum computer (a computer that could run a quantum algorithm/program) will not become reality for at least 20 years. However, at the University of Bristol we are developing a range of quantum technologies, based on light in microchips, that will deliver new quantum applications in communications, sensing, and computing. Within 10 years it is anticipated that quantum technologies will be capable of performing calculations beyond the limits of conventional computers—for example, in the simulation of complex quantum systems."

Or how about building a cloud-based virtual supercomputer from the servers you have lying around idle? Servers that could be used as a supercomputer today, and something else tomorrow. A supercomputer for hire, if you will. Amazon did just that in fall 2011, building a virtual supercomputer out of its Elastic Compute Cloud (Amazon EC2), the worldwide network of data centers it uses to provide cloud computing power to other businesses across the world.

In fact, Amazon's global infrastructure is so vast that it could run a virtual supercomputer—one that's fast enough to be placed as the 42nd most powerful supercomputer in the world—without it affecting any of the businesses that are using its data centers for their own means. Nowadays, then, you don't need to physically build a supercomputer, when all you need to do is just hire enough virtual servers to get the computational power you desire to run your particular project, and for a fraction of the price it would cost to build your own.

This form of distributed supercomputing is something the Chinese could easily reproduce, and if they were to connect all their different supercomputing complexes together and then run them as a virtual supercomputer, they could easily outpace anything the United States could produce by a factor of 100.

However, these are just ways of employing technologies that already exist in a commercial form. The truly exciting research is based on technologies only just being tested in the lab. It's in the development sphere that various technology buzzwords and phrases are bandied about, all of which could lay the foundation for the next generation of supercomputing. These include developments such as optical information processing, atomic-scale silicon wires, and quantum computing.

OPTICAL PROCESSING

Purdue University is working on two of these technologies already. The first concerns optical information processing—where light instead of electricity is used to transmit signals, much like in a fiber optic cable.

However, although fiber optic cables can transmit a lot of data, the need to translate the data back into electronic signals at either end of the cable slows down the data transmission considerably (and opens the data up to cyber-attack). The optical processing dream is for the information carried by the light to be transmitted directly into the computer, and Purdue thinks it may have found a way to do just that.

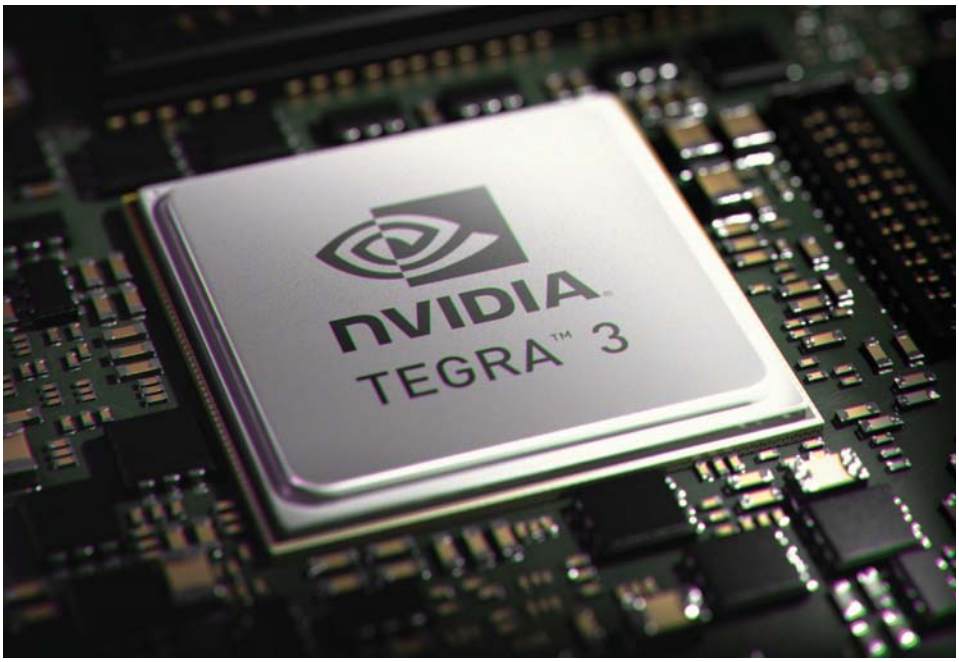

Using low-power processors, such as Nvidia's Tegra 3, is one method of increasing the energy efficiency of new supercomputers.

A team based there has succeeded in making a "passive optical diode" that transmits optical signals in only one direction. As associate professor of electrical and computer engineering Minghao Qi explains: "This one-way transmission is the most fundamental part of a logic circuit, so our diodes open the door to optical information processing." Made from two tiny silicon rings measuring 10 microns in diameter (about one-tenth the width of a human hair), they are small enough that millions can fit millions on a computer chip.

As well as speeding up and securing fiber optic data transmission, these optical diodes could speed up supercomputers entirely if they were used to connect the processors sitting inside the computer. As graduate research assistant Leo Varghese points out, "The major factor limiting supercomputers today is the speed and bandwidth of communication between the individual superchips in the system. Our optical diode could be a component in optical interconnect systems that could eliminate such a bottleneck."

Another breakthrough made by Purdue University, in conjunction with the University of New South Wales and Melbourne University, is the smallest wire ever developed in silicon. By inserting a string of phosphorus atoms in a silicon crystal, the research teams created a silicon wire one atom high and four atoms wide, with the same current-carrying capability as copper wires.

This could pave the way for the creation of nanoscale computational devices and make quantum computing a reality (See "Quantum Computing: Why Does It Matter?" on page 40). With size being such an issue in the world of supercomputing, this research shows what could be the ultimate limit of downscaling using silicon, leading to the smallest possible silicon-based supercomputers.

KEEP PEDALLING?

With all these new technologies being employed, are we already chasing computing power for the sake of just being the best rather than thinking of practical applications? We asked Purdue University's Gerry McCartney if supercomputing is going too fast for people to keep up

with. "You have to have people who know how to build, operate, and maintain these systems and, more importantly, how to work with researchers to make it as easy as possible for them to get their research done using these tools," he answered. "There's no surplus of qualified people. They're in great demand. As a university supercomputing center, we help teach students by involving them as student employees with important roles. These students end up with multiple job choices in the industry. But our motives aren't entirely altruistic. We're able to attract some of them to our expert professional staff after they graduate."

We asked Kevin Wohlever of the Ohio Supercomputer Center if there are problems with training people to know how to use these emerging computing technologies. He said it's more a problem with educating companies on what's available.

"I think that supercomputing is going too fast for some industries and companies to keep up with. The rate of change in technology overall is very fast. Very good and smart people are able to look at the new technology and take advantage of it to solve today's problems. I think the issue is that lack of acceptance and use of the new technology by the companies and people higher up in the organizations, or that have a nice profit stream from current technology, suppresses the use or widespread adoption of new technology. We do need to get some better standards in place to adopt new technology faster. I think we need to show that this is an exciting field to get into. It may not have the financial benefits of some technology or social technology companies, but there is still exciting work to be done."

> AMAZON'S GLOBAL INFRASTRUCTURE IS SO VAST THAT IT COULD RUN A VIRTUAL SUPERCOMPUTER

This excitement is evident throughout the supercomputing industry, and whatever the politics, it's clear that supercomputers are more and more necessary in our world of high-speed calculation. If the western world were to lose its dominance in supercomputing, it could potentially fall behind in many other areas of science.

However, this seems an unlikely occurrence—the race is on, and it's likely to be never-ending. When the prize is just to be the fastest computer in the world, there's always another contender just around the corner. The true benefits of this continued race will be seen by the consumers, as these supercomputing leviathans pave the way for the computer technologies that we'll find in our homes, cameras, tablets, and phones a mere few years later, not to mention the scientific breakthroughs that faster computing speeds will make possible. It's unlikely that any government will fail to recognize the importance of supercomputing, and even if that is just due to political reasons of pride and competition between nations, technology will still be the winner.

COMING SOON: EXASCALE COMPUTING

EXASCALE COMPUTING is a big dream in the supercomputing world—producing a machine with 1,000 times the processing power of the fastest crop of supercomputers. Most pundits are predicting the breakthrough by the end of the decade, but there are various real-world constraints that need to be overcome first.

Technically, there's nothing stopping someone from assembling the hardware to create an exascale computer today. The problems are how much power it would consume, its mammoth size, and the heat it would produce. Even today's high-performance computers would melt in minutes if you took away the cooling fans used in their ice-cold bunkers.

China's Sunway BlueLight MPP computer is leading the way in energy efficiency. Its power requirement is one megawatt. In comparison, its predecessor, the Tianhe, built using U.S. technology, consumes four megawatts when running, and America's Jaguar around seven megawatts. That's a considerable saving in both power and running costs for a machine that's about 74 percent as fast as the Jaguar, especially when you consider that each megawatt translates to about a million dollars a year in electricity costs. Running Japan's K computer at its June peak of 8.162 petaflops took 9.89 megawatts. That's roughly $10 million a year.

The leap to exascale may seem a big one to make by 2020, but it will be needed to wade through all the data that will be stored by then. Eight years ago there were only about five exabytes of data online. In 2009, that much data was going through the Internet in a month. Currently, the estimates stand at about 21 exabytes a month. That's not to mention all the sources of data stored elsewhere, all ripe for analysis—loyalty cards in supermarkets, credit card data, information from weather tracking satellites, tidal flows, and star mapping, to name but a few. Nowadays, data is being generated and stored everywhere, producing endless seas of statistics to wade through. Supercomputers can, and do, turn these swathes of data into useful information, analyzing trends, comparing data sets, finessing product design, producing hyper-realistic simulations of systems. Exascale supercomputing will allow us to analyze our world as never before.

To find better ways to achieve better data speeds or more energy-efficient supercomputers, research centers are starting to build their supercomputers out of more unusual components. For example, the San Diego Supercomputer Center (SDSC) has built one that uses flash storage (specifically, 1,024 Intel 710 Series drives), allowing it to access data up to 10 times faster than with traditional hard drives. At the computer's unveiling in January, Michael Norman, the director of the SDSC, said: "It has 300 trillion bytes of flash memory, making it the largest thumb drive in the world." It's a powerful thumb drive, though, ranking 48th on the Top500. In case you're wondering, the computer is called Gordon. As in Flash Gordon, not our own Gordon Mah Ung.

POWERLINE TO THE PEOPLE

By Alan Dexter

FORGET ABOUT MESSY AND HARD-TO-MANAGE NETWORK CABLING AND WI-FI THAT DOESN'T MEET YOUR NEEDS. POWERLINE TECHNOLOGY OFFERS AN EASY AND CONVENIENT WAY TO CONNECT ALL YOUR DEVICES

POWERLINE NETWORKING, like many traditional networking standards, aims to solve the seemingly simple problem of connecting devices that are physically far apart. The difference is, while traditional Ethernet networking entails cabling strewn throughout your home, powerline does it mostly invisibly, apart from the telltale adapters in your power sockets.

The modern conceit of powerline networking is surprisingly similar to that of wireless networking: connecting devices simply in the home. Indeed, as the performance and range of Wi-Fi increases, it has the potential to eat into the market of powerline networking, especially with the promise of newer standards, such as the soon-to-be-ratified 802.11ac. The beauty of a good powerline solution, though, is its simplicity—plug in an adapter next to your router and another beside your computer, television, Blu-ray player, or console, and you're good to go. There's rarely a need to configure the adapters themselves, and the range can be much greater than anything Wi-Fi can offer. You don't have to worry about the (admittedly overstated) possibility of others snooping on your data, or, more realistically, unwanted guests using your Wi-Fi because you haven't changed the default password.

Sound intriguing? Read on to see how you can get started using powerline adapters throughout your own network.

HAZ HOMEPLUG AV?

Just as Wi-Fi standards have advanced, powerline networking has also marched forward. The first serious consumer devices offered a theoretical maximum throughput of 14Mb/s, didn't adhere to any interoperability standards, suffered from interference from other devices on the same electrical loop (such as lamps, TVs, and microwaves), and were surprisingly expensive given their less-than-impressive throughput. But hey, even the most useful standards have to start somewhere.

This first generation of devices were soon replaced by faster models, offering a theoretical max throughput of 85Mb/s and a more robust infrastructure. On top of this speed increase, manufacturers began to introduce different models offering more flexibility—for instance, more ports or wireless extenders.

Interoperability soon became an issue, and the HomePlug AV standard was formed by the major players to make sure we didn't all drop powerline before it had a chance to shine. The standard dictated that any device with the HomePlug AV logo on the box must operate at predefined speeds, and be compatible with other devices at the same time. This is something that has continued, and if you're looking to mix and match a variety of adapters, you simply have to make sure they all carry the HomePlug AV logo.

The big news about the first strand of these devices was that the operating speed of the networks increased to 200Mb/s. In real-world terms, you'd do well to see anything like that performance (generally anything approaching two-thirds that speed at a useful distance is to be commended), but even 60–70Mb/s is sufficient for a single HD stream or for moving medium-size files around.

Powerline networking is a fast, reliable, and easy way to extend your network to areas where Wi-Fi fears to tread.

> THE IEEE 1901 STANDARD HAS SEEN THE SUPPORTED SPEED INCREASE UP TO 500Mb/s

GENERATIONAL DIFFERENCES

You can still find plenty of last-gen adapters on sale, and although we'd recommend avoiding the first two generations of adapters, anything displaying the HomePlug AV 200 certification is worth picking up, especially as the pricing can be so low—starter kits can be picked up for as little as $50.

This standard has recently been tweaked in preparation for the IEEE 1901 standard, which has seen the supported speed increase up to 500Mb/s. Again, you won't see such high transfer speeds between two machines; in testing at reasonable distances, you'll generally see less than 100Mb/s. However, the extra potential bandwidth means you can add other adapters to the same network (theoretically up to 255), so you can stream multiple HD movies around the home.

GETTING STARTED

If you're looking to buy your first powerline networking adapters, then a starter kit is often the best way to go, although the price difference over single adapters is nominal at best.

More interestingly these days are the number of options available to solve particular problems with individual setups. Pass-through power sockets, for instance, are available and ensure that you don't lose a power socket when installing a powerline network. This can be a problem if you're looking to include your television, consoles, or set-top box in your home network.

Another modern fad for manufacturers is to include three or four network connectors on adapters. Again, this makes them an ideal solution if you're looking to share your network with all the devices under and around your television, as you don't need to use a network switch to get the network to the right place. These units tend to be bigger than their single-socket siblings, but the reason for this is fairly obvious.

At the opposite side of this scale is the predilection for manufacturers to produce mini adapters. These compact units are ideal if your sockets are in a tight space or in a visible spot. They tend to command a premium over standard-size adapters, but they're worth considering if you have a socket that's low to the ground or is shared with a device that has a built-in transformer.

Another consideration for anyone looking to buy adapters is Gigabit Ethernet. As HomePlug AV is now running well over the 100Mb/s limit of standard Ethernet connectors, you would assume that manufacturers have shifted to Gigabit connectors. Surprisingly, this isn't the case. While this was less of a problem with the HomePlug AV 200 standard, which rarely saw real-world speeds close to 100Mb/s, it's more of an issue with HomePlug AV 500 devices. If you're looking to get the most out of your network, we recommend ensuring that any devices you look at boast Gigabit connectors.

POWERING ON

In the computing industry, there's rarely an ideal time to buy into a specific technology. After all, the next iteration is always just around the corner. For once though, Powerline is looking very good—the performance on offer is fast enough for multiple HD streams, and straight throughput for file transfers is excellent, too. If there's a technology around the corner that's likely to upset the powerline applecart, then it's the next iteration of Wi-Fi, 802.11ac, which boasts a theoretical throughput of 1.3Gb/s. Such devices are going to demand a hefty premium for some time to come though, so until that standard has proved itself, there's still definitely a place for powerline.

CISCO LINKSYS PLSK400

A good-value four-porter pack

POWERLINE NETWORKING continues to evolve in response to our needs. The first generation of adapters focused on taking a single link between two plug sockets, but as more devices become network-aware, our physical network needs increase. You could install a network switch in your current setup, but that's a bit messy when there are far simpler solutions available, such as the PLSK400.

The PLSK400 does the previous-gen standard proud with decent performance and a bargain price.

This isn't just a point-to-point pair of powerline adapters, as one of the devices in the kit boasts four 10/100Mb/s Ethernet ports. This enables you to share your network with up to four different units. The other adapter in the kit is a straightforward single-port unit, which is designed to plug into your router or network switch. The pair of adapters makes for an elegant solution.

When it comes to performance, the PLSK400 kit is predictably behind the most recent competition, although not as much as you might imagine. Indeed, the write performance in testing wasn't far off that of many 500Mb/s units, such as the D-Link kit below. Even so, seeing as this is essentially a last-generation product and is available at a bargain-basement price, slightly slower performance seems reasonable.

Overall, this Cisco kit shows that the outgoing 200Mb/s kit is still relevant, especially if you're a little strapped for cash. It offers incredible value for money if nothing else.

VERDICT **7**

Cisco Linksys PLSK400

$82, www.cisco.com

D-LINK DHP-501AV

Consistent throughput delivered with style

ROLLING IN at $105, D-Link provides an affordable entry point into the 500Mb/s market. Inside the formidably large box you'll find a pair of stylish powerline adapters, two cables to connect the units to your machines, a simple manual, and a setup CD that isn't really needed for day-to-day operations.

D-Link breaks from the boxy mold of most powerline adapters, and also uniquely places the Ethernet port on the side of the unit.

The units themselves are large, but they counter this with some design touches that elevate them above the boxy looks of the competition. Another neat addition on the design front is that the Ethernet port is on the right-hand side of the unit, which makes it slightly more accessible than the usual bottom placement.

For a unit that's so beautifully designed and perfectly crafted for the style-conscious tech aficionado, it seems insane that D-Link has decided to include bright blue CAT5 cabling. Given the side mounting of the ports, it's hardly discreet.

The odd choice of a 100Mb/s Ethernet port didn't affect the throughput too much in our real-world tests, where the powerline adapters were used to connect between floors and rooms. Here the D-Link managed a mixed set of performance figures—offering a fairly impressive read result of 64Mb/s, but lower write speeds of 51Mb/s. Indeed this actually makes the D-Link one of the slowest solutions here, but not by a huge amount.

Despite the odd choice of blue cabling, the D-link makes for a striking and affordable entry into 500Mb/s powerline networking. They aren't the fastest adapters, but they are consistent, which can be just as important when you're delivering HD video.

VERDICT **8**

D-Link DHP-501AV

$105, www.dlink.com

NETGEAR NANO 500 XAVB5101

Does small and steady win the race?

NETGEAR IS a respected name in networking, offering everything from routers to powerline hardware, and the Nano 500 is one of just three 500Mb/s starter kits that the company produces.

The design of Netgear's adapters is compact but pleasingly chunky at the same time. The tiny units are home to three LEDs that show power, Ethernet connectivity, and the state of the powerline network. This latter LED shows red for power connections of less than 50Mb/s, amber for up to 80Mb/s, and green for connection speeds above that. It's a simple way of checking that you're getting the optimum performance on installation.

We like how one of the LED status lights on the Neatgear adapter indicates the strength of our connection.

Netgear has sensibly elected to employ Gigabit Ethernet on the Nano 500, which means you're not going to be limited by the Ethernet side of the connection. The package comes with an installation guide, although somewhat unusually there's no install CD for checking your network.

When it came to testing, the Nano 500 managed competent, if somewhat uninspiring read and write performances across our network of 62.2Mb/s and 54Mb/s, respectively. On the bright side, the units were consistent, which can be just as important in real-world terms.

This set is a solid option for anyone looking to tap into the latest 500Mb/s powerline standard. And even though the performance wasn't enough to make it stand out from the crowd, the throughput is steady and the pricing is reasonable.

VERDICT 9

Netgear Nano 500 XAVB5101

$130, www.netgear.com

TRENDNET TPL-401E

Plain-Jane looks, attractive performance

THE TPL-401E is one of the latest additions to Trendnet's extensive powerline range, with support for the latest 500Mb/s connection standard. These are the more basic adapters, with each unit offering a single Gigabit connection.

Trendnet supplied two single TPL-401Es for testing rather than its starter twin-pack, the TPL-401E2K. This kit is available for about $95, saving you a few bucks compared to buying the units separately. Even so, they aren't exactly what you'd call expensive, and at $57 apiece, the single units are an affordable way to extend your network.

You can save a little jingle by purchasing the TPL-401E2K starter kit, as opposed to two separate 401E units.

The units are large, and extend a considerable way below the baseline of the power socket. These do look a little "budget" in comparison to some of the newer compact units that have started to appear, like the Nano 500, but the styling is fairly unobtrusive once installed.

On the performance front, we were pleasantly surprised by the peak transfer rates, although continued testing produced some surprisingly slow figures at times, as well. At their best, the TPL-401Es managed a headline-grabbing read speed of 71Mb/s, although they also slowed to an unimpressive 40Mb/s at times, with no notable changes to the load on the circuit (such as devices being turned on). Overall though, the units averaged a read performance of 62Mb/s when copying files from one machine to another.

The Trendnet TPL-401E is a respectable addition to the powerline networking party. The units aren't the prettiest available, and we'd personally prefer to see something a little more subtle, but they shine where it matters—in the performance and value-for-money stakes.

VERDICT 8

Trendnet TPL-401E

$115, www.trendnet.com

TP-LINK AV500 TL-PA511KIT

Today's tech at yesterday's tech prices

TP-LINK may not come to mind as quickly as some of the bigger network hardware companies, but if you're searching for powerline networking on a tight budget, you'll soon find its equipment popping up in your searches. This starter kit boasts a pair of identical 500Mb/s units that can be used to extend your network to hard-to-reach areas.

TP-Link's kit doesn't support the IEEE 1901 standard used by other manufacturers, such as Netgear, to support 500Mb/s connections, but TP-Link claims that its powerline units are backward compatible with all HomePlug AV kit running at 200Mb/s.

The AV500's contrasting color scheme is the opposite of unobtrusive, but we're drawn to its performance and price.

The adapters themselves aren't the classiest we've ever seen, with the dark sides and white faceplate sure to catch your eye due to their high contrast rather than striking good looks. The units are fairly compact, although not quite in the same league as the Netgear Nano 500.

In testing, the overall performance of the TL-PA511KIT surprised us, turning in read and write speeds as good as any units out there—easily on a par with sets that cost twice as much.

Conventional wisdom dictates that the latest hardware always costs a premium. Conventional wisdom therefore has a bit of a problem with this starter kit. For this price you might expect to get yesterday's technology, but everything is present and correct with these TP-Link adapters, including top-notch performance. The only thing not to our liking is the styling, and that seems like a reasonable sacrifice for an adapter kit that's half the price of some of the other options.

Thus, we're declaring TP-Link's kit the overall winner in this showdown.

VERDICT **9**

TP-Link AV500 TL-PA511KIT

$95, www.tp-link.com

DARE TO COMPARE: POWERLINE ADAPTERS

	Cisco Linksys PLSK400	D-Link DHP-501AV	Netgear Nano 500 XAVB5101	Trendnet TPL-401E	TP-Link AV500 TL-PA511KIT
Ethernet Interface	4x 10/100Mb/s	10/100/1,000Mb/s	10/100/1,000Mb/s	10/100/1,000Mb/s	10/100/1,000Mb/s
Powerline Speed	200Mb/s	500Mb/s	500Mb/s	500Mb/s	500Mb/s
Dimensions	4.6 x 3.2 x 2.1 inches	3.9 x 2.7 x 2.1 inches	2.6 x 2.1 x 1.6 inches	3.8 x 2.8 x 2.5 inches	4.1 x 2.2 x 1.5 inches
HomePlug AV Compliant	Yes	Yes	Yes	Yes	Yes
IEEE 1901 Compliant	No	Yes	Yes	Yes	No
Average Read Throughput	53.7Mb/s	64.0Mb/s	62.2Mb/s	61.8Mb/s	66.4Mb/s
Average Write Throughput	55.3Mb/s	51.1Mb/s	54Mb/s	53.6Mb/s	55.2Mb/s

iFixit presents:

AUTOPSY

THIS MONTH WE DISSECT...

Vizio Co-Star

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their products. To learn more, visit www.ifixit.com.

VIZIO CO-STAR

The Vizio Co-Star set-top box is the second attempt at a Google TV product, following the universally panned Logitech Revue. We haven't actually tested the $99 box's performance, but here's how it fares in the teardown benchmark.

MAJOR TECH SPECS:

- Marvell Armada 1500 1.2GHz dual-core processor
- Marvell 88DE2755 QDEO video processor
- Nanya NT5CB256M8GN-DI 1GB RAM
- Samsung K9GBG08U0A-SCB0 4GB NAND flash
- Marvell Avastar 88W8787 WLAN/Bluetooth/FM SoC

KEY FINDINGS:

• As soon as we were inside, we pulled off what looked like an EMI shield—only to find a thermal pad sitting on top of a sea of small resistors. It's not often that we encounter cooling for the back side of a motherboard, but Vizio took extra precautions since the Co-Star has no fans.

• A grand total of just five screws and two cables must be removed to free the motherboard.

• We're used to device innards being dominated by a single component, but those space hogs are usually batteries or power supplies. In the Co-Star, over a third of the real estate is instead inhabited by the aluminum heatsink.

• Why such a honker heatsink? The answer lies in air movement. Without fans to circulate air, there is no forced convection. Therefore, the Co-Star must rely on conduction and natural convection to keep the processor cool.

• Modders, take note: The Co-Star's silver bezel comes off easily and would probably look pretty slick when painted neon green.

• Some fun factoids on the Co-Star's remote:

Both keyboard bezel and rubber button cover easily detach from the rest of the remote, making it a cinch to clean off grime after prolonged use.

An abundance of glue secures the trackpad to the front of the Co-Star remote, but pretty much everything else comes apart with relative ease.

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

WINDOWS TIP OF THE MONTH

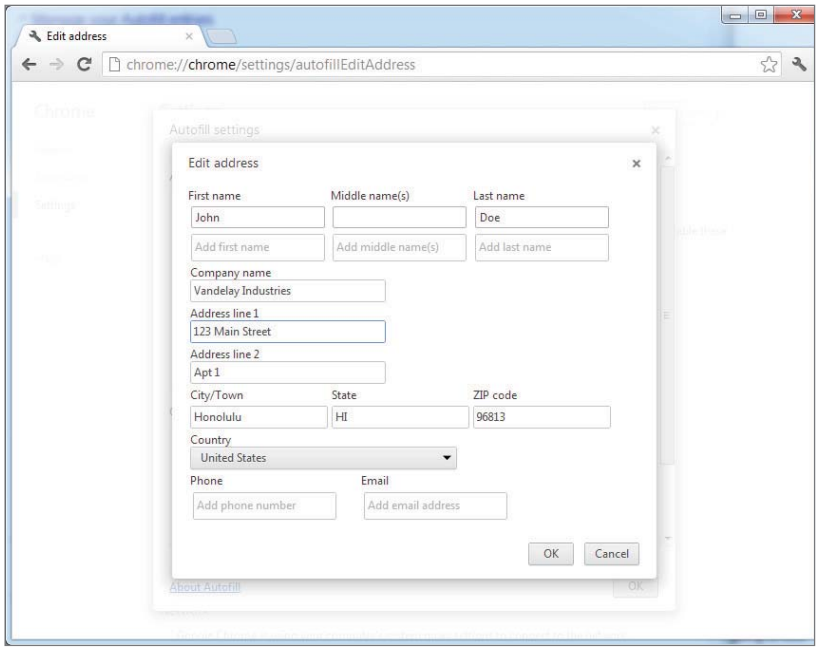

CUSTOMIZE YOUR CHROME AUTOFILL

If you're frustrated with Chrome getting your "Autofill" personal information wrong, or suggesting out-of-date address and credit card info, you can manually edit the information. You can find it in the Chrome Settings menu if you click Show Advanced Settings. Look under Passwords and Forms for the link marked Manage Autofill Settings.

MAKE - USE - CREATE

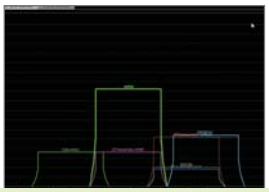

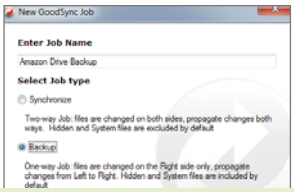

ALEX CASTLE
CONTRIBUTING EDITOR

DO YOU REALLY NEED SYNC?

LATELY I'VE been thinking a lot about cloud storage, and I've realized something interesting: For a lot of people, a file-sync program like Dropbox might not be the best choice. I've personally been a big proponent of sync, but it's got a few drawbacks.

For one, to use Dropbox, you have to have an extra process running on your PC at all times, and if there's one thing I love more than file sync, it's an empty process manager. Second, for a lot of people sync is simply unnecessary. If you only use one computer, or you don't have a lot of files you need instant access to from anywhere, you probably don't need it.

Of course, your data should still be safe in the cloud—I'm just saying that scheduled online backup could be a better choice. You can get much better per-gigabyte pricing with a non-syncing service like Mozy, Carbonite, or even Amazon's business-oriented S3.

↘ submit your How To project idea to: comments@maximumpc.com

Improve Your Wi-Fi Signal

YOU'LL NEED THIS

A WIRELESS NETWORK
Sorry, Ethernet users, these tips will only speed up a Wi-Fi network.

INSSIDER
This Wi-Fi network analyzer is available for free at bit.ly/22nvUo.

THERE ARE LOADS of ways to get a faster wireless network connection. You can knock down walls in your house so the signal can travel more efficiently. You can chuck your current network hardware and buy a whole new set. Or you could go down the easier, free route of making sure your network is operating on the right channel and tweaking its settings for maximum speed. It's like tuning a radio to a different station.

There's a simple reason you might need to change your network channel. When there are many other wireless networks in your local area, signals can clash, causing interference that makes network traffic slow down. There are several possible channels a wireless network can operate on. We're going to look at the traffic in your area and move your network to the least-clogged channel, ensuring you'll get the best results. **–ALEX COX**

1 GET STARTED To analyze your local network traffic in greater depth, we're going to use a tool called inSSIDer. It's a full, open-source network analyzer capable of some incredibly complex things, but we're really only going to skim the surface of its abilities. Download the installer from bit.ly/22nvUo and run it—just click Next repeatedly to get it installed.

» Run inSSIDer—you'll find it inside a folder called Metageek in your Start menu. Click the tab marked "2.4GHz channels," then look at the top of the screen for a button labeled Start. This sets inSSIDer hunting around your local area for networks. It'll likely find many more than Windows would normally bring up (**image A**). Make sure you click Stop when you're done.

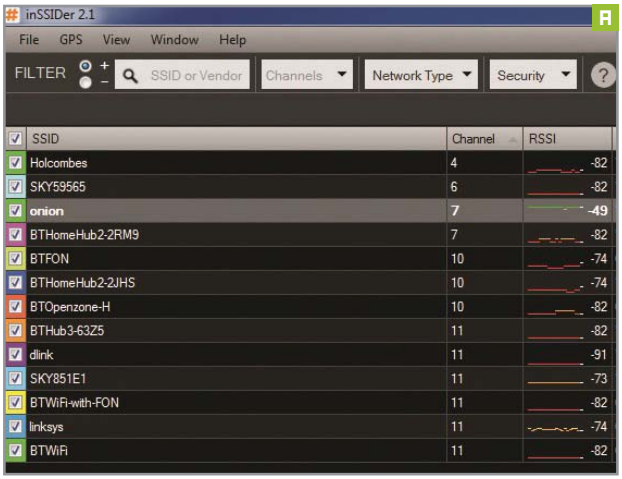

2 CHANGE THE CHANNEL Look at the graph the program presents you (**image B**). This is a map of where the networks in your area fall along the 11 main wireless channels, their height representing their signal strength from your current location. It's up to you to interpret your own graph, since it won't be the same as ours. Find an area with some empty space, or find a quiet channel, and write down its number.

» The channel you've picked out will be your network's new home. If it's nice and clear, you should get a decent boost in speed out of it. But to change the channel on your wireless router, you're going to need to log in to its admin interface, and that's different for every router. We'll show you the steps we took on our Asus WL-700ge—you may need to search the Internet for specifics about your own model.

» Start by firing up a web browser—any will do—and go to the local IP address of the wireless router, in this case 192.168.1.1. This brings up the login screen of the router, where you'll need to enter your login credentials. If you don't know yours, it's likely they're still at their defaults, so dig out the manual or search on the Internet and you should be able to gain access (**image C**).

» From here you need to access the Advanced Setting screen, as the basic wireless settings don't offer the option to change the channel. It's rather simple on this router to change the channel—there's a drop-down box in the middle of the screen that lets us do it (**image D**). Again, your router will have a different settings menu, but the process should be quite similar. While we're poking around, we can also look for a number of other settings that will keep the router speedy.

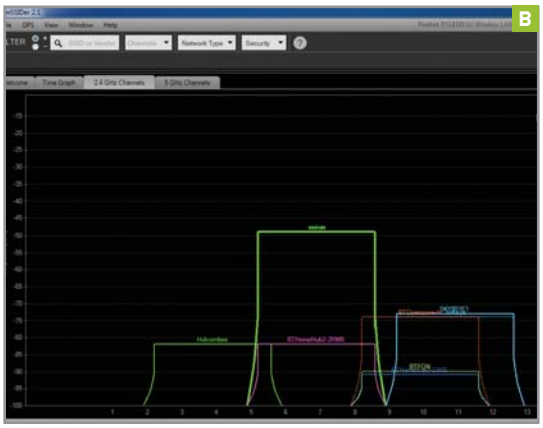

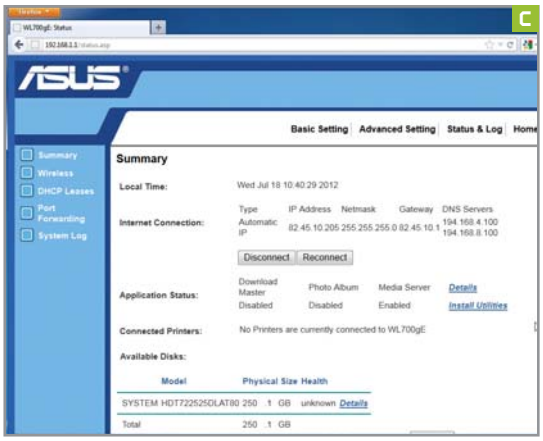

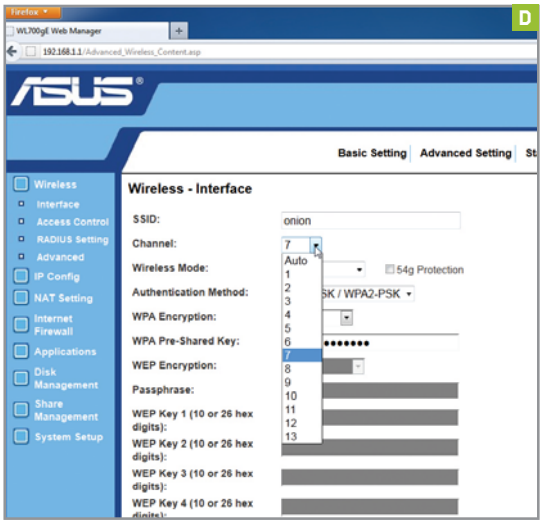

3 MAKE FURTHER CHANGES You're also given the option of setting a wireless mode, in this case choosing to turn off the older 802.11b service (**image E**). If you still have devices that use it, you'll want to leave it switched on, but if you only have modern 802.11g devices in your house, you can safely switch the mode and claw back the tiniest fraction of a speed boost. Then it's time to head to the advanced wireless page.

» Here you can fiddle with all sorts of things, but we'll stick to the bits we know are going to help wireless speed. Turning up the radio power should provide better signal at a distance, so we'll do that. Turning on the "afterburner" mode is only useful if you have wireless cards or dongles in your computers that support it, so we'll leave it off. And frame burst, lower down the page, is a must.

» After you've made the changes, you'll need to apply them and save your settings. Your router will probably reboot at this point (**image F**). Once everything has come back online, run inSSIDer again to check the results of your changes and ensure that your wireless connection has moved from its original channel to its new home. Try a download—you should see the benefits immediately.

» So you've checked out the wireless channels in your area, found the ideal one for your network, and hopefully—depending on your router—switched its channel to some empty space. If it seems to have made things worse, there may be some other interference—a cordless phone, for example—in your new channel, so try a few until you find the best one.

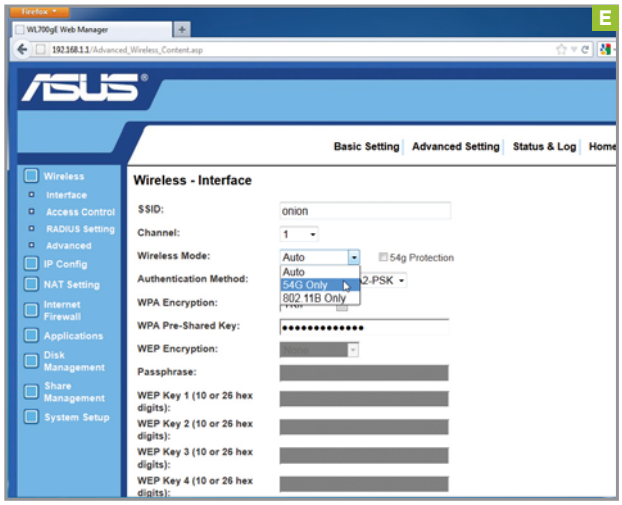

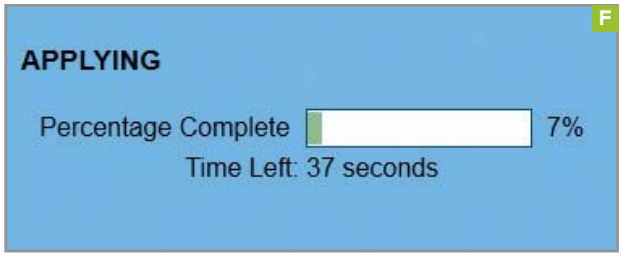

Back Up Your Cloud Storage with GoodSync

YOU'LL NEED THIS

GOODSYNC

This cloud storage backup software is available in free and paid versions at www.goodsync.com.

WITH SO MANY CLOUD storage services available to you, you don't ever truly need to pay for online storage. When your 2GB Dropbox runs out, you can always get 5 free gigs from Amazon. When that runs out, why not open up a SkyDrive account for an additional 7GB? The only problem with this plan is that your files get spread out over different services, which can make it harder to find things, and can also increase your exposure to risk of losing access to files during service outages. In this article, we'll show you how to mitigate both of these problems, by using GoodSync to keep an up-to-date local backup of all the files on multiple cloud storage services. **–ALEX CASTLE**

1 INSTALL GOODSYNC You can find the GoodSync software available for free at www.goodsync.com. The free version does have some limitations—notably on the quantity of files that can be backed up—but it can be used indefinitely and is at least a good way to judge whether the full version is worth $30 to you.

» Download the file and run the installer. There's nothing much to configure during the installation; just see it through to the end and you'll be asked to name your first job. GoodSync organizes all of your syncing tasks into jobs; each job consists of two locations to sync, which can be on your hard drive, USB, or network storage, or even on the cloud (**image G**).

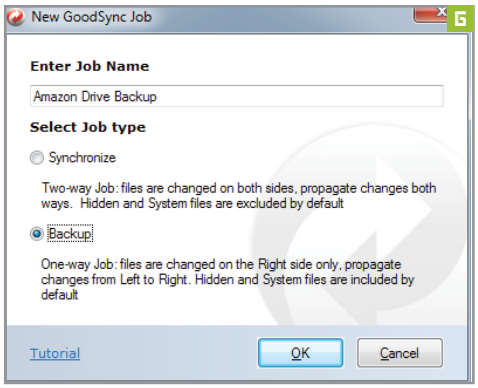

2 PICK YOUR BACKUP LOCATIONS One of the major advantages of GoodSync over a simpler sync application like SyncToy is the ability to connect directly to several of the most popular cloud storage services. If you use cloud storage from Google, Amazon, or Microsoft, you can set up a job that backs them up directly. To do so, just click the Browse button next to the left location field, and in the menu that pops up, select the service you want to connect to. Enter your login credentials and click Go to connect to the service and see a list of folders to sync.

» GoodSync can't connect directly to Dropbox, but it's still easy to back up those files using the software, as long as you have the Dropbox client installed on your PC. Simply click the Browse button, and select My Computer to browse to the location of your My Dropbox folder.

» Next, click the Browse button next to the right-side location field. This time, in the popup window select the location where you want your backup files to be stored. You can click My Computer and choose a local location (**image H**), or select Windows Shares or FTP to select a network location.

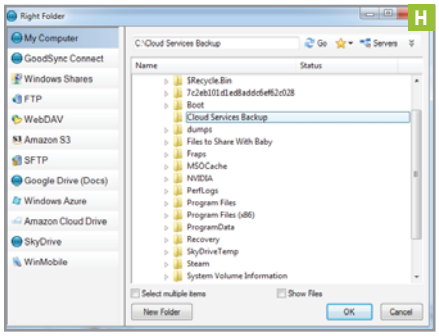

3 SCHEDULE YOUR BACKUP Once you've specified both the left- and right-hand locations, you just have to tell GoodSync how to transfer files between them. To do this, click the small icon in the upper-center of the GoodSync window with a picture of two green arrows. In the window that opens (**image I**) click the Sync Direction to back up from the cloud account to your local location, not vice-versa.

» Finally, click the Auto tab in the backup options window to select how often GoodSync will automatically back up. Unless you need to know that your files are instantly backed up, we recommend conserving system resources by scheduling a backup every two hours or more.

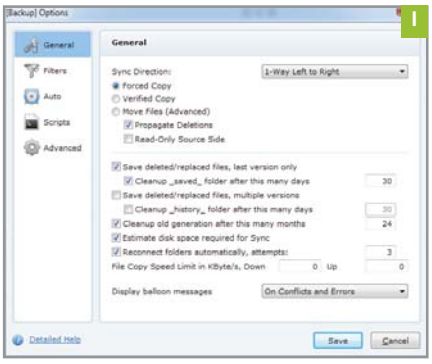

BUILD IT

BY JONATHAN ROBERTS

Bake Your Very Own Raspberry Pi

A step-by-step guide to getting up and running with the world's smallest bare-bones PC—the Raspberry Pi

LENGTH OF TIME: **2 HOURS** | LEVEL OF DIFFICULTY: **ADVANCED**

THE MISSION The Raspberry Pi is a tiny bare-bones PC built by a team in the UK who wanted to bring a low-cost computer to the masses. The idea was born of their goal to increase interest in software development, programming, and general hardware hacking. When they found that a lot of young people didn't have access to PCs these days, they baked a Pi to fix the situation. You order the Pi online (www.raspberrypi.org) and it's shipped to you as a simple circuit board with a few components soldered to it, and that's it. There's no operating system, no power supply, no storage or input devices, no instruction manual, and no cables of any kind. This is what you get for $35, in other words. Don't be scared, though; we're here to help. In this Build It we'll show you what peripherals you need, how to create a bootable SD card and boot the system, and what to do once you've got it up and running. Finally, we'll cover what settings you might want to tweak, and what you can and can't do with the Pi currently.

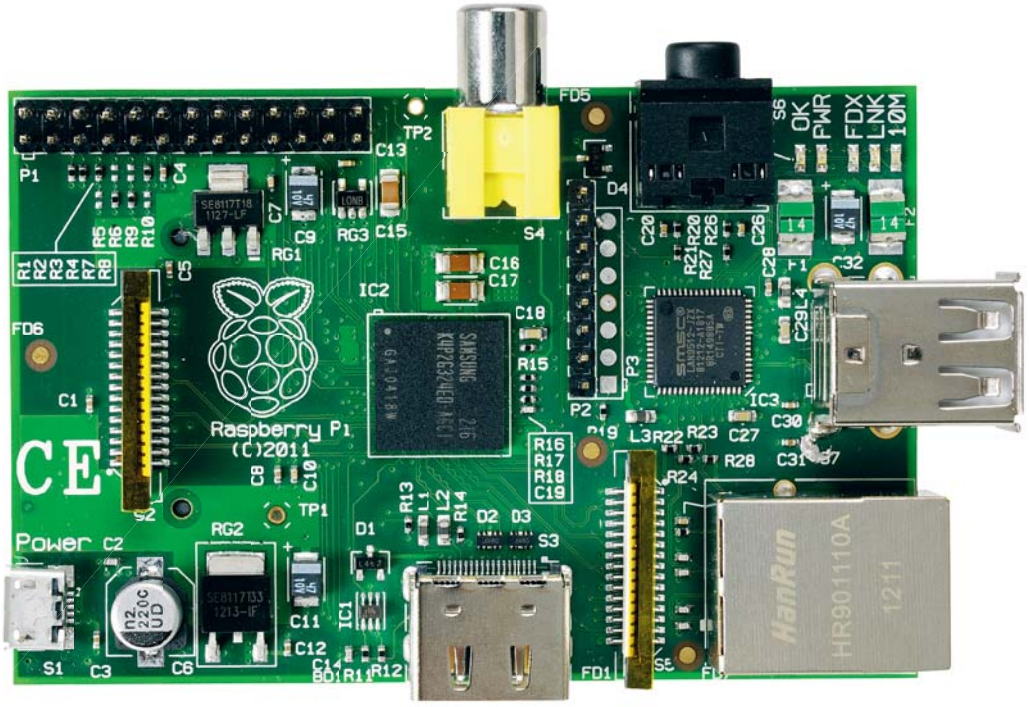

The Raspberry Pi is a credit-card-size Linux PC designed to help people get their feet wet with programming and hardware hacking.

SPECIFICATIONS

CPU/GPU	Broadcom BCM2835 700MHz SoC
RAM	256MB SDRAM
Connectivity	2 USB 2.0 ports, HDMI, composite RCA, SD slot, 10/100 Ethernet, 3.5mm jack
Weight	1.58 oz

WHAT YOU NEED TO GET STARTED

The Raspberry Pi is truly a bare-bones PC, so in order to utilize it in even the most basic ways you're going to need a few pieces of hardware and software first, so let's look at each one individually.

POWER SUPPLY: The power supply has to provide at least 7mA at 5V and have a Mini USB connection. Many smartphone chargers match these specifications, so check yours. We've been using a standard Android phone charger in the office and it works beautifully.

STORAGE: You'll need an SD card, and it must be Class 4 and have at least 4GB of space available. Of course, the more space there is, the more applications and data you'll be able to keep on your Raspberry Pi, so go to town if you want since storage is cheap these days (you can get 32GB for just $20). If you're unsure which to buy, you can find a list of cards known to work at bit.ly/wVO5By. You'll need to be able to read and write to the SD card from a computer other than the Raspberry Pi, so if your PC doesn't have a card reader, you'll need to buy an external reader.

CABLES: The HDMI/composite cable doesn't need to be expensive. You can buy cables online from Monoprice.com for about $5. As for which type of cable you need, HDMI or composite, that depends on what kind of display you're using. If you're hooking up your Raspberry Pi to a very old TV, you'll need to use a composite connection. If your display has a DVI connection, you should buy an HDMI-to-DVI converter, or a dedicated HDMI-to-DVI cable. If it has both composite and HDMI connections, you should opt for the HDMI, as this will give you the best possible image quality.

USB KEYBOARD/MOUSE: Almost any keyboard and mouse combo will do, just be aware that your old PS/2 mouse and keyboard will not work here.

INGREDIENTS

	PART	URL	PRICE
PC	Raspberry Pi	www.raspberrypi.org	$35
PSU	Mini USB Power Supply	www.amazon.com	$5
HDMI Cable	Generic cable from the Internet	www.monoprice.com	$5
Storage	4G SD Card	ww.amazon.com	$5
Media Reader	SD Card Reader	www.amazon.com	$5
Peripherals	USB Keyboard/ Mouse	ww.amazon.com	$30
OS	Debian Linux Distro	www.raspberrypi.org/ downloads	$0
Total			$85

THE OPERATING SYSTEM: Not only do you need extra peripherals before you can get started, the Raspberry Pi is also missing an operating system. To make it work, you need to get a hold of a compatible operating system and write it to the SD card. We'll show you how, below. Note: It's wise to read through all of the steps first before actually embarking on the process.

1 ACQUIRE AN OS

TO GET the operating system onto the SD card, you first need to download it. You can find all the compatible operating systems on the Raspberry Pi website: www.raspberrypi.org/downloads. The developers recommend that you download the Debian version for the time being, which is what we're using in this tutorial. The download is large at 400MB-plus, and there are other, slightly smaller versions, but they don't offer Debian's flexibility.

Once you've downloaded the image, you need to connect the SD card to the computer. If your PC has a built-in card reader, use that; if it doesn't, you'll need to connect a USB card reader.

The download is a zip file, which you should extract to its own folder. Inside the newly created folder you'll find an image file with the .img extension—this file needs to be copied on to the SD card, but it can't be copied like a normal file. The .img file, as the name suggests, is an exact image of a storage device. It contains information about the layout of data on the device and files used to boot the computer, which have to be in exactly the right part of the drive. If you just copy it, the files won't be in the right place, but will be stored on top of the disk's existing structure. Instead, we need to copy it in a way that replaces the disk's structure and data with that of the .img file. We'll do this in the next step.

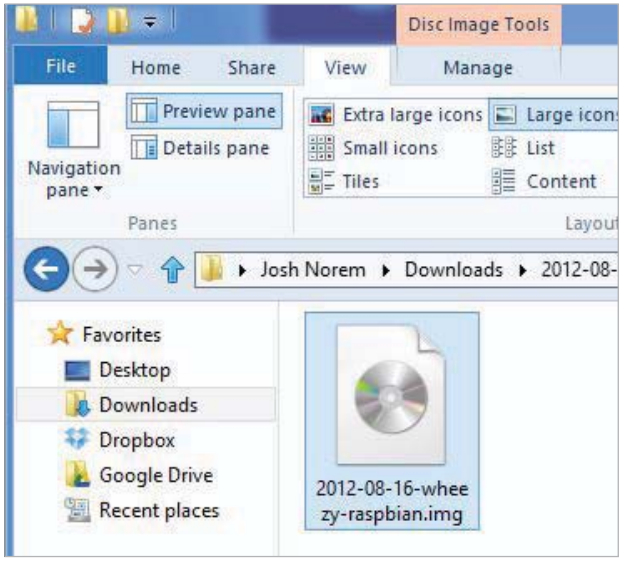

You can run any Linux distro you want on your Pi, but the developers recommend this version of Debian.

2 PREPARE THE SD CARD

WINDOWS USERS will need to use a special program to copy the first byte of the image to the first byte of the disc, the second byte to the second byte, and so on (see sidebar on the right). Linux users can do this using dd, a command-line tool for file-copying.

Using Linux, first put the SD card in your computer, then open a terminal window and type `df -H`. This command will display a list of all storage devices and partitions on your PC, along with their size, allowing you to ascertain the device label your computer has assigned to the SD card.

Look through the list for an entry that has the same size as your memory card, and specifies Media in the Mounted column on the right. You may not find one that's exactly the same size as your card, but there will be something close. Make a note of the information under the Filesystem column; it should be something like `/dev/mmcblk0`.

Next, figure out the path to the downloaded image file. For example, we downloaded it to the downloads folder of our home directory, making its path: /home/maxpc/Downloads/fileX.img.

If you saved it elsewhere on your PC, you'll need to note the full path to that location instead.

Now that you have both the location of the image file and the location of the SD card, head to the command line and make sure you have precise information, as this next command will wipe out anything on the destination device. Ready? OK, enter the following command: `sudo dd if=/home/Maxpc/Downloads/debian6-19-04-2012/debian6-19-04-2012.img of=/dev/mmcblk0`

The part that follows the word "if" is the path to the downloaded image file, and the part after "of" is the SD card location you noted earlier. This will remove all data from the SD card, copy the image file to it, and make it bootable.

3 BOOT UP

WITH THE DANGEROUS part of this project behind us, we can now prepare for our first boot. To do this, simply insert the SD card in the slot on the bottom of the Raspberry Pi, and connect your USB keyboard, mouse, display, and Ethernet cable. Then, with the monitor turned on, connect the power supply and the Raspberry Pi should boot itself.

You'll see a red power light next to the USB connections, and shortly after you should see some flashing green lights. If you don't see these, it means something went wrong with the SD card's formatting. First, unplug the power cable, double-check that the card is pushed into the slot properly, and then reconnect the power. If it still doesn't work correctly, go back and try to configure the SD card again.

If all is well, however, you should see a lot of text scroll by on the screen with a Raspberry Pi logo at the top. The first time you start up, it enters a one-off initialization process. Along the way, it may stop loading when it says "Stopping Portmap Daemon"; if it does stop at this point, you can simply turn the power off and on again, and the RPi will start without any difficulty the next time around.

4 MEET THE PROMPT

EVENTUALLY, the text will stop scrolling across the screen and there will be a line at the bottom that reads `raspberrypi login:` followed by a blinking cursor. This is the login prompt, and it means the machine is waiting for you to enter your username and password. On the Debian image, the username is "pi" and the password "raspberry." Type the username and press Enter, and then type the password and press Enter. It won't show your password as you type it, just in case anyone is looking over your shoulder. After doing that, you'll find the text stops at another line, similar to the last time, and it reads `i@raspberrypi:~$` followed by a blinking cursor. This is the command prompt, which is where you type whatever commands you want the Raspberry Pi to execute for you. If you plan to use your Raspberry Pi to learn about computers and how they work, to run your own server or to do some programming, you'll spend more time with the command prompt, but for now let's get the graphical interface working.

WRITING THE SD CARD IN WINDOWS

If you don't have a Linux PC, you can use the free utility Win32 Disk Imager to move the OS onto your SD card. This tool was originally developed to work with the Netbook Remix of Ubuntu, but it has now become a popular tool for people playing with ARM devices (such as the Raspberry Pi). You can download it from bit.ly/Dh2nE. It's simple to use—just select the .img file in the first field, and select the SD card from the Device field next to that. Then all you need to do is click Write to complete the process.

Make sure you get the correct drive, and consider backing up all your data first because it's possible to lose everything if you make a mistake. Also, this tool will erase all information on the SD card, so be careful when using it.

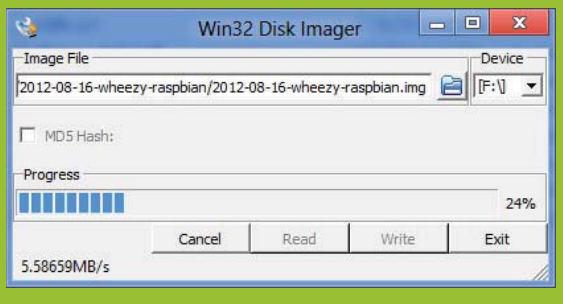

Windows users can use Win32 Disk Imager to write the OS to the SD card.

5 ENTER THE GUI

AT THE COMMAND prompt, type `startx` and press Enter. The screen will flicker for a few moments, then turn black, then a huge Raspberry Pi logo will appear, with a taskbar across the bottom that's similar to Windows. This is LXDE, the Lightweight X11 Desktop Environment. On Linux, users can choose from many graphical interfaces to interact with their computers. LXDE isn't the prettiest or most modern of these, but it's one of the leanest. This makes it perfect for running on hardware without much processing power or memory, like the Raspberry Pi.

Have a look around and click the different buttons to see what they do—you can't break anything at this point, and even if you did, you could just unplug the Pi and reconfigure your SD card and you'd be back where you started.

Although you can't see the command line anymore, it's running in the background. You can access it again by pressing Ctrl + Alt + F1—try it. You should see some obscure-looking text following the startx command you issued. This is the output from the running program, the X Server.

X is the program that Linux uses to control input devices and display information on the screen. It creates a series of virtual consoles, accessed by typing Ctrl + Alt + [F1–F6]. The graphical environment—on the Raspberry Pi at least—will always be accessible on the console represented by F2. Therefore, to get back to LXDE you just press Ctrl + Alt + F2. One thing you might notice as you play and explore is that the desktop doesn't fit your display properly. This is a common problem that some people using HDMI connections have found. It's caused by something called overscan, which is a relic from television sets made from the 1930s through to the 1990s.

6 FIX THE OVERSCAN PROBLEM

TO FIX THE overscan problem, we'll return to the command line. We won't switch virtual consoles as we did before—we'll use a terminal emulator instead. Open the LXDE menu by clicking the icon in the bottom-left of the screen, then click Accessories>LXTerminal. This is another way to access the command line. In the new prompt window, type `sudo leafpad /boot/config.txt`. This will launch a text editor called Leafpad. You don't need to open a file since the command you entered also specifies which file should be opened to edit—config.txt in the /boot directory. In the new window, type `disable_overscan=1` and then save the file. You can then reboot your Raspberry Pi by clicking the logo in the bottom-right of the screen. When it comes back on, the screen should be filled.

While you're looking at it, there are a lot more options you can put in the config file. A lot of them have to do with display modes like resolution, others cover modifying how fast the processor runs—overclocking it, essentially. The hdmi_drive option is useful if you're using an HDMI-to-DVI converter. Set its value to 1 to ensure your screen is at the correct resolution and reboot.

CAN I GAME ON MY PI?

In a word, no—you can't, but you may be able to in the future. The main problem, aside from hardware limitations, is that the Raspberry Pi is very new, so people have only just started tinkering with the hardware, finding and reporting problems, and creating compatible applications that make the most of its capabilities. Also, you may not be surprised to learn that there are some problems with hardware support in the distributions. So far, none of them have X—the software that controls the display—working properly with the integrated graphics system, which means that it's slow to redraw windows.

It's also impossible to get most existing 3D games to work. This is because most 3D games on Linux have been written to interact with the graphics system through something called OpenGL, but Raspberry Pi supports only a subset of this, called OpenGLES. Hopefully, in the future new games will appear and old games will be modified to work with OpenGLES. We expect to see X gain support for integrated graphics, too.

Gaming isn't the Raspberry's Pi's forte, but at least you can play Quake 3 on it.

7 ESTABLISH STORAGE SPACE

AT THIS POINT, you've got your Raspberry Pi running, you know how to start the graphical interface, and you've set the screen resolution. It's a great start, but there's one more tweak we want to describe: how to configure your SD card to make all its storage space available. The reason this is necessary is that when you put the .img file on to the SD card, it automatically divides the drive into three partitions. One of these is used for storing the operating system and all your files (known as the root partition, represented by a /), another is used for providing virtual memory (known as the swap file), and the third is used for storing the files that let the Raspberry Pi start. Because each is seen as another physical device, if there's room on the SD card that's outside of these partitions, the operating system will ignore it.

The Debian image we put on the SD card creates a root partition that's always the same size (around 2GB). As a result, on our 4GB SD card we found ourselves with 2GB of unused space, and only a few hundred megabytes available to store our new applications and files. To fix this problem and reclaim the wasted space, you'll first need a PC with GParted installed. The app is free and supports both Linux and Windows (available at gparted.sourceforge.net).

Before starting, back up your data. Now, with the SD card in your card reader, run GParted. Look in the drop-down menu and choose a device that's the same size as your SD card. If it's the right one, you'll see three partitions listed in the main window and three segments of unallocated space (**image A**). Right-click the linux-swap partition and select Resize/Move. In the window that appears, drag the small box all the way to the right of the box that it's contained within (**image B**). Click the Resize/Move button at the bottom of the screen.

Next, right-click the ext4 partition, select Unmount (**image C**), then right-click and select Resize/Move. Drag its right-hand edge to the end of the box to resize it. Click Resize/Move. This finalizes the moves we want to make, but nothing has been done yet.

If you're sure you've got it right, click the green tick at the top of the screen, then wait patiently while GParted moves and resizes your partitions (**image D**). Take the SD card out, stick it in your Raspberry Pi, and power it up.

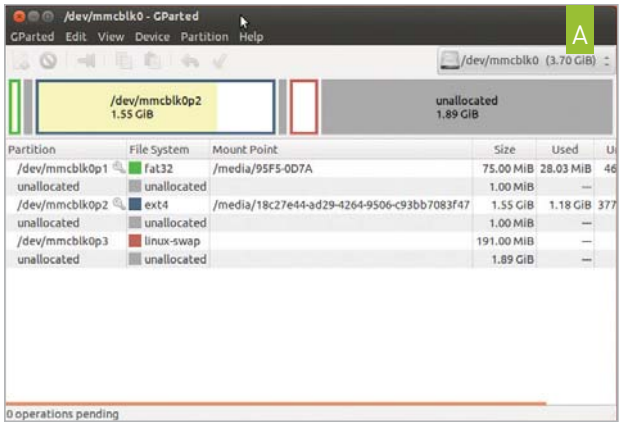

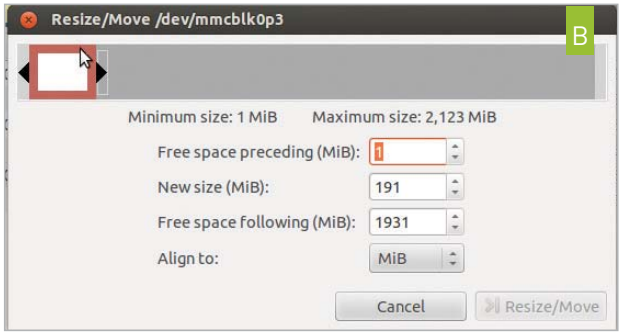

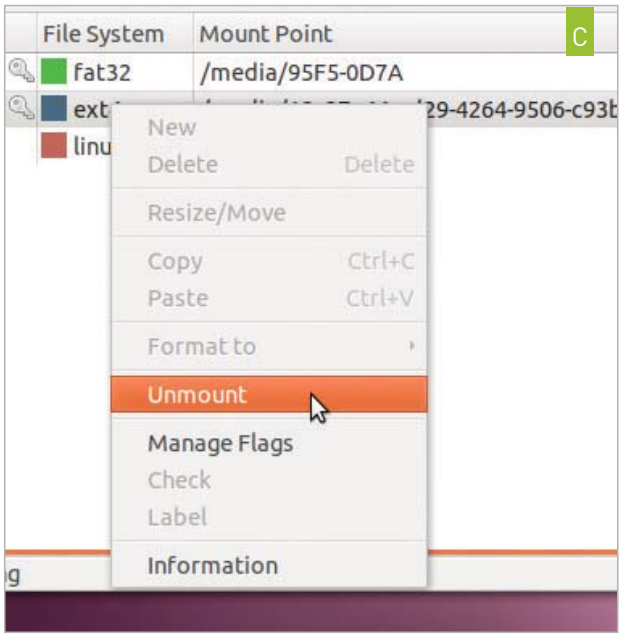

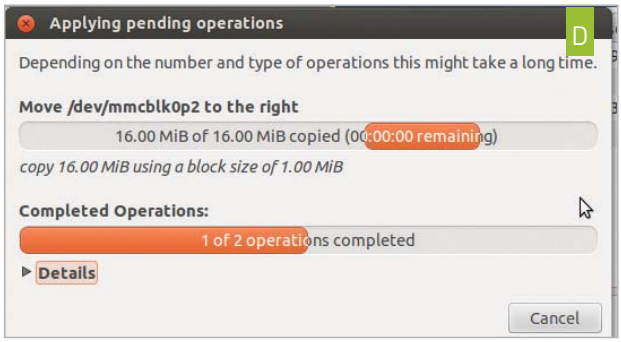

REVIEWS

TESTED. REVIEWED. VERDICTIZED.

INSIDE

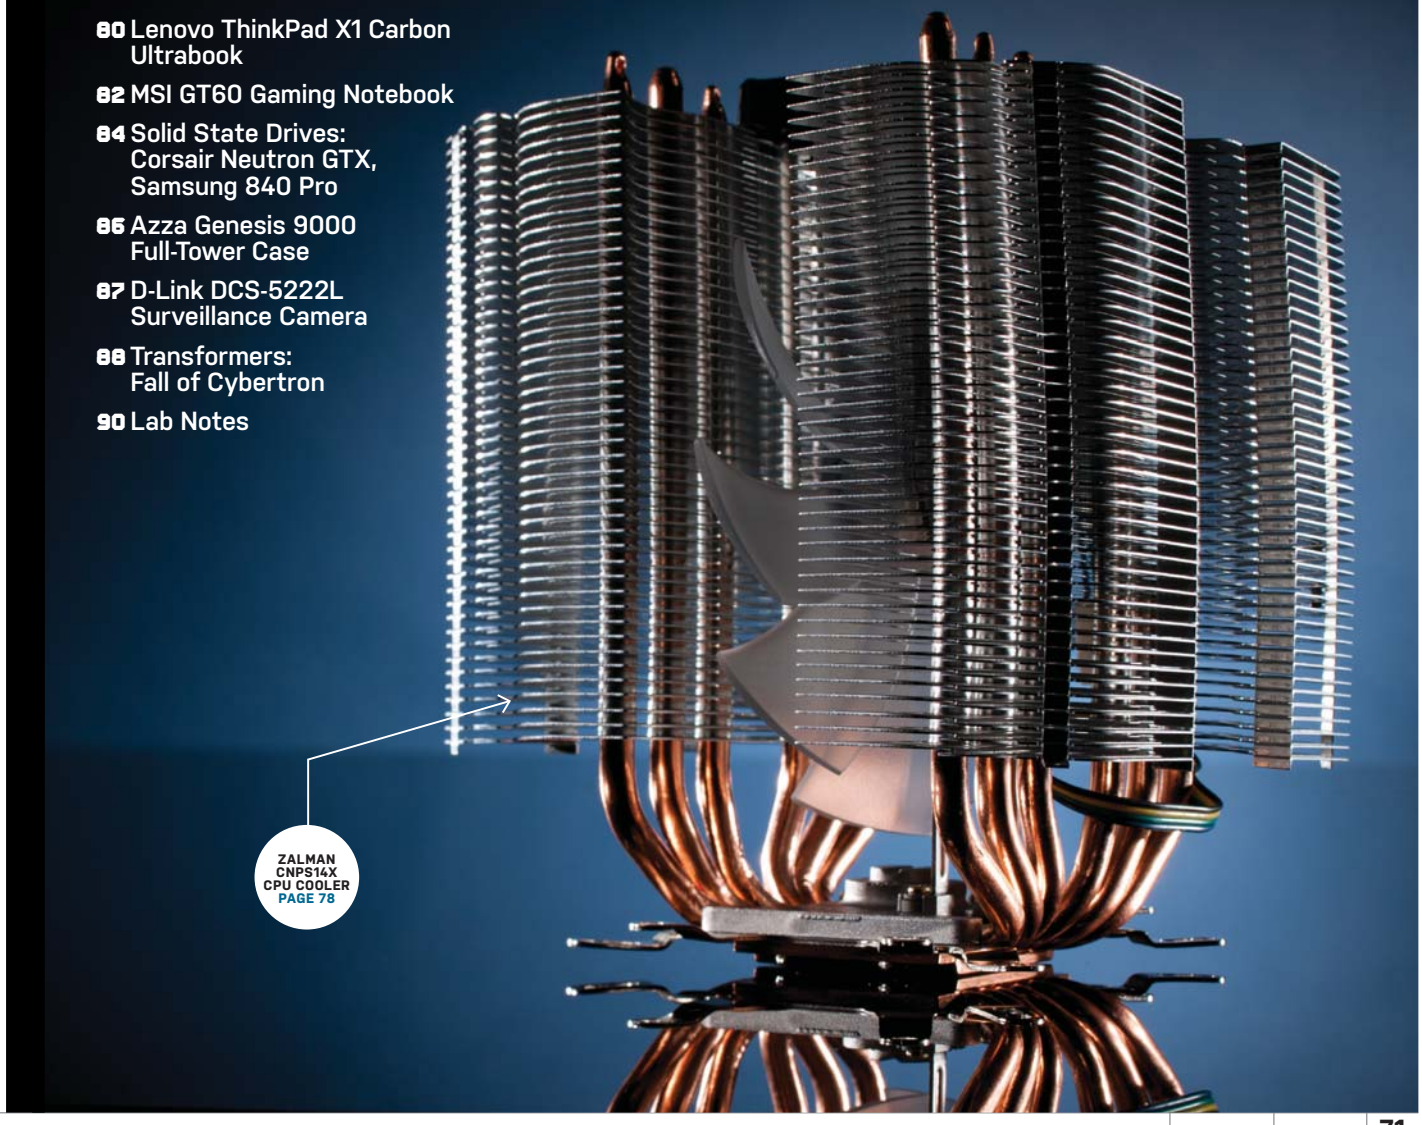

ZALMAN CNPS14X CPU COOLER PAGE 78

ViewSonic Smart Display VSD220

Call it an Android-in-one

IT'S HARD TO review ViewSonic's new Smart Display VSD220 without thinking back to another of the company's unusual products that we reviewed almost a decade ago: the Air Panel V110.

The Air Panel used Microsoft's "Smart Display" technology to essentially let you remotely control your PC over Wi-Fi for browsing and MP3 streaming. Not to rehash ancient history, but Smart Display was just another charred carcass on the long road to a successful consumer tablet computer.

ViewSonic's new Smart Display VSD220, though, is actually nothing like its predecessor and, in fact, shows how much the world has turned. While the Air Panel tried to leverage desktop power, the new Smart Display VSD220 reverses the tide by leveraging tablet technology in a desktop.

The Smart Display VSD220 is essentially a 22-inch, 1920x1080 touch-enabled monitor with Android 4.0.4 in it. That's right, there's no need to hook it up to a PC—you just run Android directly on the monitor. Think of it as an Android-in-one.

The unit features a dual-core, 1GHz TI OMAP 4428 Cortex CPU with a PowerVR SGX 540 GPU. The screen is an LED sidelit panel with optical touch sensors. Unlike the capacitive-touch overlay that your tablet or phone uses, optical is actually affordable in a large screen. A capacitive-touch panel alone at 22-inches would cost about $900.

The unit sports two standard USB ports on the right, a Micro USB port on the left, Ethernet, audio out, Micro HDMI in, and a microSD slot. Android has had USB HID support for some time and adding a USB mouse and keyboard to the VSD220 let us navigate the interface and type mostly without issue. Some of the apps we ran, such as Minecraft, ran only with touch.

In performance, the unit was incredibly underwhelming. We couldn't find the data sheet for the OMAP 4428, but we'd guess that it's two notches below the OMAP 4430 used in the Kindle Fire. As much as the original Fire was a breakthrough tablet, performance on it sucked. In scrolling and panning, you can tell the OMAP4430 and PowerVR SGX GPU are underwater here. The situation is not helped by the OMAP's support for only dual-channel low-power DDR2. Just as with a notebook or desktop PC, the more memory bandwidth you can get, the more performance you'll see in your applications.

Although it's not a direct comparison, we decided to see how the Smart Display VSD220 compared to Google's Nexus 7 tablet in a few benchmarks. The result was ugly for the ViewSonic.

We know that performance isn't the top check-off item on thin-client OSes, but we'd like smooth scrolling at least. ViewSonic engineers surmised that the slow performance could be from the optical digitizer, which doesn't have the resolution of capacitive touch—but the unit felt slow using the mouse to move around, as well.

By now, we're sure you're asking the pesky "why even build such a device?" question. At this point, we don't know. We understand that some computer phobes may find such a low-maintenance device appealing, but for all others, we'd suggest a pass. The unit does have one redeeming quality: Hook an HDMI and USB cable to your PC, and you can run it as a touch-enabled panel for Windows 8.

But for us to even consider this device, we'd want a faster CPU and GPU along with an optimized OS such as Jelly Bean. Then we'll talk. **–GORDON MAH UNG**

VERDICT 5

ViewSonic Smart Display VSD220

+ SMART ALEC Will work with Windows 8 for touch.

– SMART WATER Badly needs more processing and graphics power.

$400, www.viewsonic.com

SPECIFICATIONS

	ViewSonic VSD220	Google Nexus 7
CPU	1GHz dual-core TI OMAP 4428 based on ARM A9	1.3GHz quad-core Nvidia Tegra 3 based on ARM A9
GPU	PowerVR SGX 540	Nvidia GeForce ULP
RAM	1GB LP-DDR2	1GB LP-DDR3
Resolution	1920x1080	1280x720
OS	Android 4.0.4 Ice Cream Sandwich	Android 4.1.1 Jelly Bean

BENCHMARK

	ViewSonic VSD220	Google Nexus 7
GLBenchmark 2.5 Egypt HD (fps)	5.4	**14**
GLBenchmark 2.5 Egypt HD Offscreen (fps)	5.4	**9.4**
GLBenchmark 2.1 Egypt Classic (fps)	14	**55**
GLBenchmark 2.1 Classic Offscreen (fps)	15	**32**
Vellamo 1.0.6 (score)	800	**1,695**
Nenamark 2 (fps)	17.4	**55.8**
Sunspider 0.9.1 (ms)	2,304	**1,713**
LINPACK Single (MFLOPS)	37.3	**45.8**
LINPACK Multi (MFLOPS)	51.1	**133.1**

Besides functioning as a desktop-size Android device, the VSD220 can serve as a stand-alone touchscreen monitor for a full-fledged PC.

Gigabyte's OC Edition offers additional clock speeds and cooling at no extra charge.

The GeForce GTX 660 Debuts

The sweet-spot GPU has finally arrived

The GTX 660 is the first 28nm Kepler board based on a new GPU dubbed GK106, and the final 6-series card to support high-performance features like GPU Boost and SLI. Compared to the GTX 660 Ti, the GTX 660 offers the same 2GB of DDR5 memory, the same 192-bit memory interface, and the same number of ROP units, but loses two SMX units, giving it just 960 CUDA cores compared to 1,344 in the previous card (and 1,536 in the GTX 680). At $230, it's our new favorite GPU in the price-to-performance category. **–JOSH NOREM**

GIGABYTE GTX 660 OC VERSION

Gigabyte's GTX 660 is similar to MSI's board in that it's overclocked and has a cooler with a silly name—Windforce. The board is clocked at the same base and boost clock speeds as the MSI card, too, running at 1,033MHz and 1,098MHz, respectively. The cooler features four copper heat pipes, aluminum fins, and two large 10cm fans breathing down on the whole shebang. Even though the board sports a smallish 7.5-inch PCB, the cooling apparatus is so large that it's 2 inches longer than the PCB and extends the length of the card to 9.5 inches. With a cooler this large you expect it to perform quite well, and it does. It kept the card absolutely silent even when the board was being tortured in the Lab, and allowed it to run at a moderately cool 63 C under full load.

The board's benchmark performance was good enough for it to edge out the MSI card in most tests, but by a margin that's insignificant. Since both cards are clocked the same and spec'd the same, this level of parity is not surprising. On the software front, the included OC Guru II app is easy to use for monitoring the card's vitals and overclocking, but the charts it provides showing change over time are hard to read due to teeny text. There's also a "bundle" in the box but we are putting it in quotes because it's so meager—a single Molex-to-PCIe adapter.

The biggest surprise with the Gigabyte card is its MSRP of $230—exactly the same price as a bone-stock GTX 660, so it's like the extra overclocking and cooling are free. Since it offers roughly the same software, cooling, performance, and bundle as the MSI card (reviewed next), the price advantage is all we need to give the nod to Gigabyte in this contest.

VERDICT 9 KICK ASS!

Gigabyte GTX 660 OC Version

$230, www.gigabyte.com

MSI GTX 660 TWIN FROZR OC EDITION

MSI's GTX 660 is an all-around great card that includes a healthy dollop of overclocking and a side of Frozr to keep it cool. Its base clock speed is a decent 53MHz over stock at 1,033MHz, and when running at full load we saw its boost clock speed rise 130MHz over stock to 1,110MHz, which is also higher than the stock boost-clock spec. The Twin Frozr III cooler sports three copper heat pipes, aluminum fins, and dual 8cm fans housed in a metal-alloy shroud to direct the airflow. Like the other GTX 660 cards, it uses just a single 6-pin power connector, but unlike the others it sports an extra-long 9-inch PCB (Gigabyte's board is just 7.5 inches but the cooler is actually 9 inches long).

MSI's overclocked GTX 660 sports a sexy metallic cooling shroud.

BENCHMARKS

	Gigabyte GTX 660 OC Version	MSI GTX 660 OC Edition	Gigabyte GTX 660Ti OC Version	EVGA GTX 560Ti Classified	XFX Radeon HD 7850
Price	$230	$240	$320	$200 (street)	$230 (street)
3DMark 2011 Perf	7,093	6,935	**9,095**	6,295	6,075
3DMark Vantage Perf	27,858	27,737	**31,468**	25,523	24,584
Unigine Heaven 2.5 (fps)	36	36	**40.2**	31.8	31
Shogun 2, 1080p (fps)	53.76	52	**68.7**	43.1	47
Far Cry 2 / Long (fps)	119	118.4	**131.1**	112.4	103
Dirt 3 (fps)	75.7	76.5	**96.8**	68.7	50
STALKER: CoP DX11 (fps)	42.8	42	**49.8**	38.4	34.7
Just Cause 2 (fps)	60.4	58.6	**67**	55.8	51
Batman: Arkham City (fps)	76	76	**81**	64	60
Metro 2033 (fps)	22	22	22.3	22.6	**23**

Best scores are bolded. Our test bed is a 3.33GHz Core i7-3960X Extreme Edition in an Asus P9X79 motherboard with 16GB of DDR3/1600 and an AX1200 Corsair PSU. The OS is 64-bit Windows Ultimate. All tests are run at 1920x1200 with 4x AA and all settings maxed out, except for the 3DMark tests, and Shogun 2, which is run at 1080p High settings.

SPECIFICATIONS

	Gigabyte GTX 660 OC Version	MSI GeForce GTX 660 Power Edition	GeForce GTX 660Ti	GeForce GTX 670	AMD HD Radeon 7850
Number of Cores	960	960	1,344	1,344	1,024*
Texture Units	80	80	112	112	64
SMX Units	5	5	7	7	16
ROPs	24	24	24	32	32
Base Clock Frequency	1,033MHz	1,033MHz	915MHz	915MHz	860MHz
Boost Clock Frequency	1,098MHz	1,098MHz	980MHz	980MHz	N/A
Memory Clock Frequency	1,502MHz	1,502MHz	1,502MHz	1,502MHz	1,200MHz
L2 Cache Size	384KB	384KB	384KB	512KB	512KB
Frame Buffer Size	2,048MB	2,048MB	2,048MB	2,048MB	2,048MB
Memory Interface	192-bit	192-bit	192-bit	256-bit	256-bit
Memory Data Rate	6Gb/s	6Gb/s	6Gb/s	6Gb/s	4.8Gb/s
Connectors	2x dual-link DVI, 1x HDMI, 1x DisplayPort	2x dual-link DVI, 1x HDMI, 1x DisplayPort	2x dual-link DVI, 1x HDMI, 1x DisplayPort	2x dual-link DVI, 1x HDMI, 1x DisplayPort	2x dual-link DVI, 1x HDMI, 2x DisplayPort
Power Connectors	1x 6-pin	1x 6-pin	2x 6-pin	2x 6-pin	1x 6-pin
Thermal Design Power (TDP)	140W	140W	150W	170W	150W

*Note: Nvidia and AMD graphics compute cores are not directly comparable.

In testing, the MSI board ran neck-and-neck with Gigabyte's similarly clocked offering, losing every test but one by a very slim margin. The cards were also tied in the category of noise/cooling, as they both ran silently under full load at a mild 63 C. The MSI Afterburner software is usable but nothing to email home about and the board's skimpy bundle consists of a single Molex-to-PCIe adapter along with the software CD.

Overall there's a lot to like about MSI's Twin Frozr OC Edition card. It's quiet, cool, and performs very well, and we like the look of its gunmetal shroud and glinting metallic badge. However, the Gigabyte GTX 660 board performs just as well and is just as quiet, and since it costs $10 less, that's enough to give Gigabyte the advantage. If you are an MSI fan and/or a Frozr lover, you can get the 660 OC Edition for the same $230 price via a mail-in rebate, but since we hate rebates we'd rather just go with the Gigabyte and get the savings up front at the virtual cash register.

VERDICT 9

MSI GTX 660 Twin Frozr OC Edition

$240, www.msi.com

The Power Edition is overclocked a smidge, but you can go to town on its core clock and voltages, as well.

MSI GeForce GTX 670 Power Edition

A smokin'-fast GTX 670 card designed for overclocking

WE'RE BIG FANS of the GTX 670 cards, and that's not a joke relating to their generous-size cooling mechanisms, either. Oh no—this is a serious love affair, brought on by the fact that with just a little overclocking, this $400 GPU can be just as fast, if not faster, than the $500 GTX 680 flagship GPU. Not only is achieving that level of performance "sticking it to the man," it's also sound economics.

The reason the GTX 670 can achieve performance parity with the GTX 680 under certain conditions is because the two cards are extremely close in specs. The biggest difference is that the GTX 670 has seven SMX units instead of the GTX 680's eight, resulting in 192 fewer CUDA cores (1,344 compared to 1,536 in the GTX 680) and 16 fewer texture units, making it 112 compared to 128. Otherwise the two cards have the same 2GB of 6GHz GDDR5 memory, the same 256-bit memory interface, same number of ROPs, and the same 1,502MHz memory frequency.

When it comes to clock speeds, MSI's engineers have massaged the clocks by 1,05MHz, which brings its base clock to 1,020MHz (reference boards are clocked at 915MHz). Boost clocks have also been, uh, boosted, by 117MHz, allowing the card to clock all the way up to 1,097MHz when under duress. These are great numbers for a factory overclock, but they are not quite as aggressive as the Asus GeForce GTX 670 DirectCU II TOP (August 2012), which is clocked at 1,058MHz base and 1,137MHz boost. The Asus card is also more expensive, however.

In addition to overclocking, owners of this board can overvolt it as well, which is exceedingly rare these days, so kudos to MSI for allowing it. The included MSI Afterburner software lets you adjust voltage, core and memory clocks, fan speed, and more. The monitoring software is hard to read, but tweaking values is as simple as moving a slider. The dual-fan Twin Frozr cooling apparatus is designed to suck up as much heat as you can throw at it, all while running silently.

In testing, the MSI board ran neck-and-neck with the more expensive board from Asus as well as the much more expensive flagship GTX 680, a result that is both surprising and impressive. In our benchmarks, the MSI and Asus boards shared the top step on the podium on every test, and the MSI board also beat the GTX 680 in six out of 11 tests, but by a margin so narrow we're calling it a three-way tie between the GTX 670 cards from MSI and Asus and the EVGA GTX 680. The MSI board's Twin Frozr cooler hung tough in testing, keeping the proc semi-warm at 75 C but running as quiet as a mime at all times.

Overall we're impressed with this card. It's fast, silent, and you can overclock and tweak it to your heart's desire. Plus it's less expensive than the Asus card, and runs as fast as a GTX 680. Sounds pretty Kick Ass to us. **–JOSH NOREM**

BENCHMARKS

	MSI GTX 670 Power Edition	Asus GTX 670 DirectCU II TOP	Nvidia GTX 670 Reference	EVGA GTX 680	Sapphire Radeon HD 7950 OC
Price	$410	$430	$400	$500	$400
3DMark 2011 Performance	P9,500	**P9,664**	P8,706	P9,555	P7,683
3DMark 2011 Extreme	X3,266	**X3,300**	X2,957	X3,249	X2,562
3DMark Vantage Performance	33,855	**34,700**	32,568	34,339	31,752
Unigine Heaven 2.5 (fps)	29.7	30.7	28.6	**31.2**	26.7
Shogun 2 (fps)	**32.9**	**33.7**	29.9	**35**	24.8
Far Cry 2 / Long (fps)	108.21	**110.5**	103.4	107.3	89.6
Dirt 3 (fps)	70.02	**73.3**	65.5	**73**	53.8
Metro 2033 (fps)	17	16.7	16	16.3	**18.7**
STALKER: CoP SunShade(fps)	**37**	35.4	32.4	34.3	34.1
Just Cause 2 (fps)	54.6	**57.4**	52.7	54.7	45.3
Batman: Arkham City (fps)	**62**	60	57	58	54
Base Clock	1,090MHz	1,058MHz	915MHz	1,006MHz	900MHz
Boost Clock	1,097MHz	1,137MHz	980MHz	1,058MHz	N/A
Memory Clock	1,502MHz	1,502MHz	1,502MHz	1,502MHz	1,250MHz

Best scores bolded. Our GPU test bed consists of a stock-clocked Intel Core i7-3960X on an Asus P9X79 Deluxe board with 16GB DDR3/1600, a 256GB Samsung 830 Series boot SSD, and a 1,050W Thermaltake Toughpower Grand PSU, in a Cosmos II chassis. All tests performed at 2560x1600 with all settings maxed and 4x MSAA except where noted. Power use measured with a Watts Up Pro.

VERDICT

9

KICK ASS!

MSI GeForce GTX 670 Power Edition

+ CHINCHILLA Excellent performance; whisper-quiet; can be overclocked and overvolted.

– DOUBLE CHIN "Bundle" is just two Molex adapters.

$410, www.msi.com

The CNPS14X features six copper heat pipes that feed into two massive heatsinks.

Zalman CNPS14X

A headache to install, but works well

INSTALLING THE ZALMAN CNPS14X CPU cooler is sort of like doing P90X exercises—there's a possibility of extreme discomfort, but if you're tough enough to bear the burden, you will ultimately see good results. The reason the installation is so painful is not so much the size of the cooler—we've installed coolers larger than this before without resorting to profanity—it's the construction of the cooler and fan, which dictates the installation process.

The first culprit is the 14cm fan that spins between a pair of contoured aluminum heatsinks. The fan is attached via screws to the middle of the heatsinks, preventing access to the area around the cooler's baseplate. This is important, as that's where we typically secure the cooler to the CPU socket. Without access to this area, the cooler must be secured by tightening screws that are underneath the heatsink itself, with no access to them from above. Suffice to say, this is no easy task given the cooler's size—it's large enough to require low-profile RAM.

We installed the mounting brackets to the cooler with some difficulty; there was no reference in the manual beyond an unhelpful, thumbnail-size photo. Once we had the brackets secured, we had to remove the motherboard from our LGA2011 system because the CNPS14X's heatsinks protruded so far out that it made tightening the screws near the baseplate of the cooler impossible without additional room to maneuver. Zalman includes a 3-inch spanner designed to tighten the bolts from the outside perimeter of the cooler, but there was simply not enough clearance in our already-built system. We only decided to pull the mobo out after we had already removed the RAM and GPU, as we needed more clearance around the area by the I/O panel and the top of the case.

Once we had the CNPS14X installed, it performed quite well but didn't blow our socks off. Under full load it quietly cooled our 3.3GHz (overclocked to 4.2GHz) Intel Core i7-3960X, but only had one degree Celsius on our current favorite, the Cooler Master Hyper 212 Evo. We also compared the CNPS14X to the similarly priced NZXT Havik 120, and the CNPS14X outperformed it by nearly three degrees, which is impressive. Since Zalman provides clips to add additional fans for a push-pull config, we added two more 14cm fans, but temps only dropped by less than one degree Celsius, so purchasing extra fans doesn't seem like a wise investment.

Overall, we were impressed by the Zalman CNPS14X's effectiveness and quiet operation. Its cooling performance puts it in the top tier of all the CPU coolers we've tested, but unfortunately it was also one of the most difficult coolers to install in recent memory. If you want a massive cooler that's easier to install, we'd go with the NZXT Havik 120, which is roughly the same price, or the Cooler Master Hyper 212 Evo, which offers the same performance for significantly less dough. –JIMMY THANG

BENCHMARKS

	Zalman CNPS14X (Performance mode)	Zalman CNPS14X (Quiet mode)	CM Hyper 212 Evo	NZXT Havik 120
Ambient Air	22.9	22.2	22.8	23.4
Idle Temperature	34.6	34.6	38.6	38
Burn Temperature	**77.1**	77.6	78.6	80.5
Burn - Ambient	**54.2**	55.4	55.8	57.1

All temperatures in degrees Celsius. Best scores bolded. All tests performed using an Intel Core i7-3960X at 4.2GHz, on an Asus Sabertooth X79 motherboard with 16GB DDR3/1600, in a Thermaltake Level 10 GT with stock fans set to Low.

VERDICT **8**

Zalman CNPS14X

ICE T Cools well; quiet; extra fans can be added.

VANILLA ICE Painful install; performance not much better than cheaper Hyper 212 Evo.

$50, www.zalman.com

SPECIFICATIONS

Dimensions H x D x W	5.5 x 5 x 6.3 inches
Weight	1 lb, 9.2 oz
Heat Pipes	6
Stock Fans	1x 14cm 4-pin PWM
Socket Support	Intel LGA775/1155/1156/1366/2011; AMD AM2/AM2+/AM3/AM3+/FM1
Add'l Fan Support	2 (clips included)

Lenovo ThinkPad X1 Carbon

An Ultrabook that makes no Mac pretenses

UNLIKE MANY an Ultrabook, there's no mistaking this one for a MacBook Air, or even an Air wannabe. Staying true to the venerable ThinkPad brand, the X1 Carbon is matte-black through and through, and clad in that distinct rubberized coating that feels nice to the touch, won't easily slip from your grip, and remains blessedly free of fingerprints. It looks every bit the business companion it's intended to be. In fact, the X1 Carbon looks a lot like the ThinkPad X1 we reviewed last year (bit.ly/lEdkj4). But it's grown from 13 inches to 14 inches, and its body has been flattened to Ultrabook standards, measuring just .71 inches at its thickest. Its lap weight, by the way, comes in just under three pounds.

Another notable difference from last year's X1 is the Carbon's carbon-fiber construction (see what they did there?). The material is used in the notebook's shell and its internal "roll cage," making it both lightweight and rugged enough to withstand eight Mil-Spec tests for toughness.

The X1 Carbon does feel sturdy. Despite its thin profile, there's little flex to the keyboard deck when the notebook is held by one corner. The hinges seem solid, and as an added bonus allow the notebook to open a full 180 degrees. We also have to give props to the keyboard. It's one of the most satisfying we've used on an Ultrabook, with nice large keys and satisfying travel. If you like backlighting, the Carbon's got it. The touchpad is also quite nice—smooth and predictable—and TrackPoint is there for folks who like to control the cursor with ThinkPad's signature red nubbin.

The screen is a full 1600x900 pixels, and a matte coating keeps glare and reflections to a minimum. But the vertical viewing angle is fairly narrow, leading to contrast and color degradation off axis.

Compared to our previous Ultrabook zero-point rig—the first-gen Asus Zenbook UX31E—the Ivy Bridge–sporting X1 Carbon scored decisive wins in all the benchmarks, with margins of 15 to 60 percent. But this month we're debuting a new Ultrabook zero-point, an Intel reference design featuring a 1.8GHz i5-3427U. This is the same CPU found in the Carbon X1, but as you can see from the chart, performance wasn't identical. We attribute the Intel rig's wins to a combination of additional memory bandwidth (DDR3/1600 vs. DDR3/1333), which impacts game performance in particular, and a speedier SSD. Both drives have a SATA 6Gb/s controller, but CrystalDiskMark revealed a significant disparity, with Lenovo's SSD achieving read/writes of 372.3- and 173.5Mb/s, respectively, to the Intel's 479.5- and 303.5Mb/s.

Battery life on the X1 Carbon wasn't stellar at a little more than four hours in our video rundown test, but Lenovo's RapidCharge makes rejuicing speedy. We were back at 75 percent battery capacity in 30 minutes and 100 percent in an hour.

The X1 Carbon isn't perfect, and that can make its relatively high price less palatable, but it does offer features you can't get anywhere else, namely ThinkPad quality and a look that's cool in that not-even-trying kind of way.

–KATHERINE STEVENSON

VERDICT 8

Lenovo ThinkPad X1 Carbon

+ CARBON FIBER Understated good looks; solid build quality; nice keyboard/touchpad.

– FIBROMYALGIA Pricey; 128GB SSD; average battery life.

$1,500, www.lenovo.com

SPECIFICATIONS

CPU	1.8GHz Intel Core i5-3427U
RAM	4GB DDR3/1333
Chipset	Intel HM67
Display	14-inch 1600x900 antiglare LCD
Storage	SanDisk 128GB SSD
Connectivity	Mini DisplayPort, USB 3.0, USB 2.0, headphone/mic, 4-in-1 SD card reader, Bluetooth 4.0, 802.11n, 3G WWAN, 720p webcam, USB-to-Ethernet dongle
Lap / Carry	2 lbs, 15.5 oz / 3 lbs, 13.8 oz

BENCHMARKS

	ZERO-POINT	
Premiere Pro CS3 (sec)	840	900 (-6.7%)
Photoshop CS3 (sec)	102	106 (-3.8%)
ProShow Producer (sec)	1,113	1,173 (-5.1%)
MainConcept (sec)	1,904	1,968 (-3.3%)
Quake III (fps)	430.8	353 (-18.1%)
Quake 4 (fps)	73.5	56.1 (-23.7%)
Battery Life (min)	315	251 (-20.3%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point ultraportable is an Intel reference Ultrabook with a 1.8GHz Intel Core i5-3427U, 4GB of DDR3/1600 RAM, integrated graphics, a 240GB SSD, and Windows 7 Home Premium 64-bit.

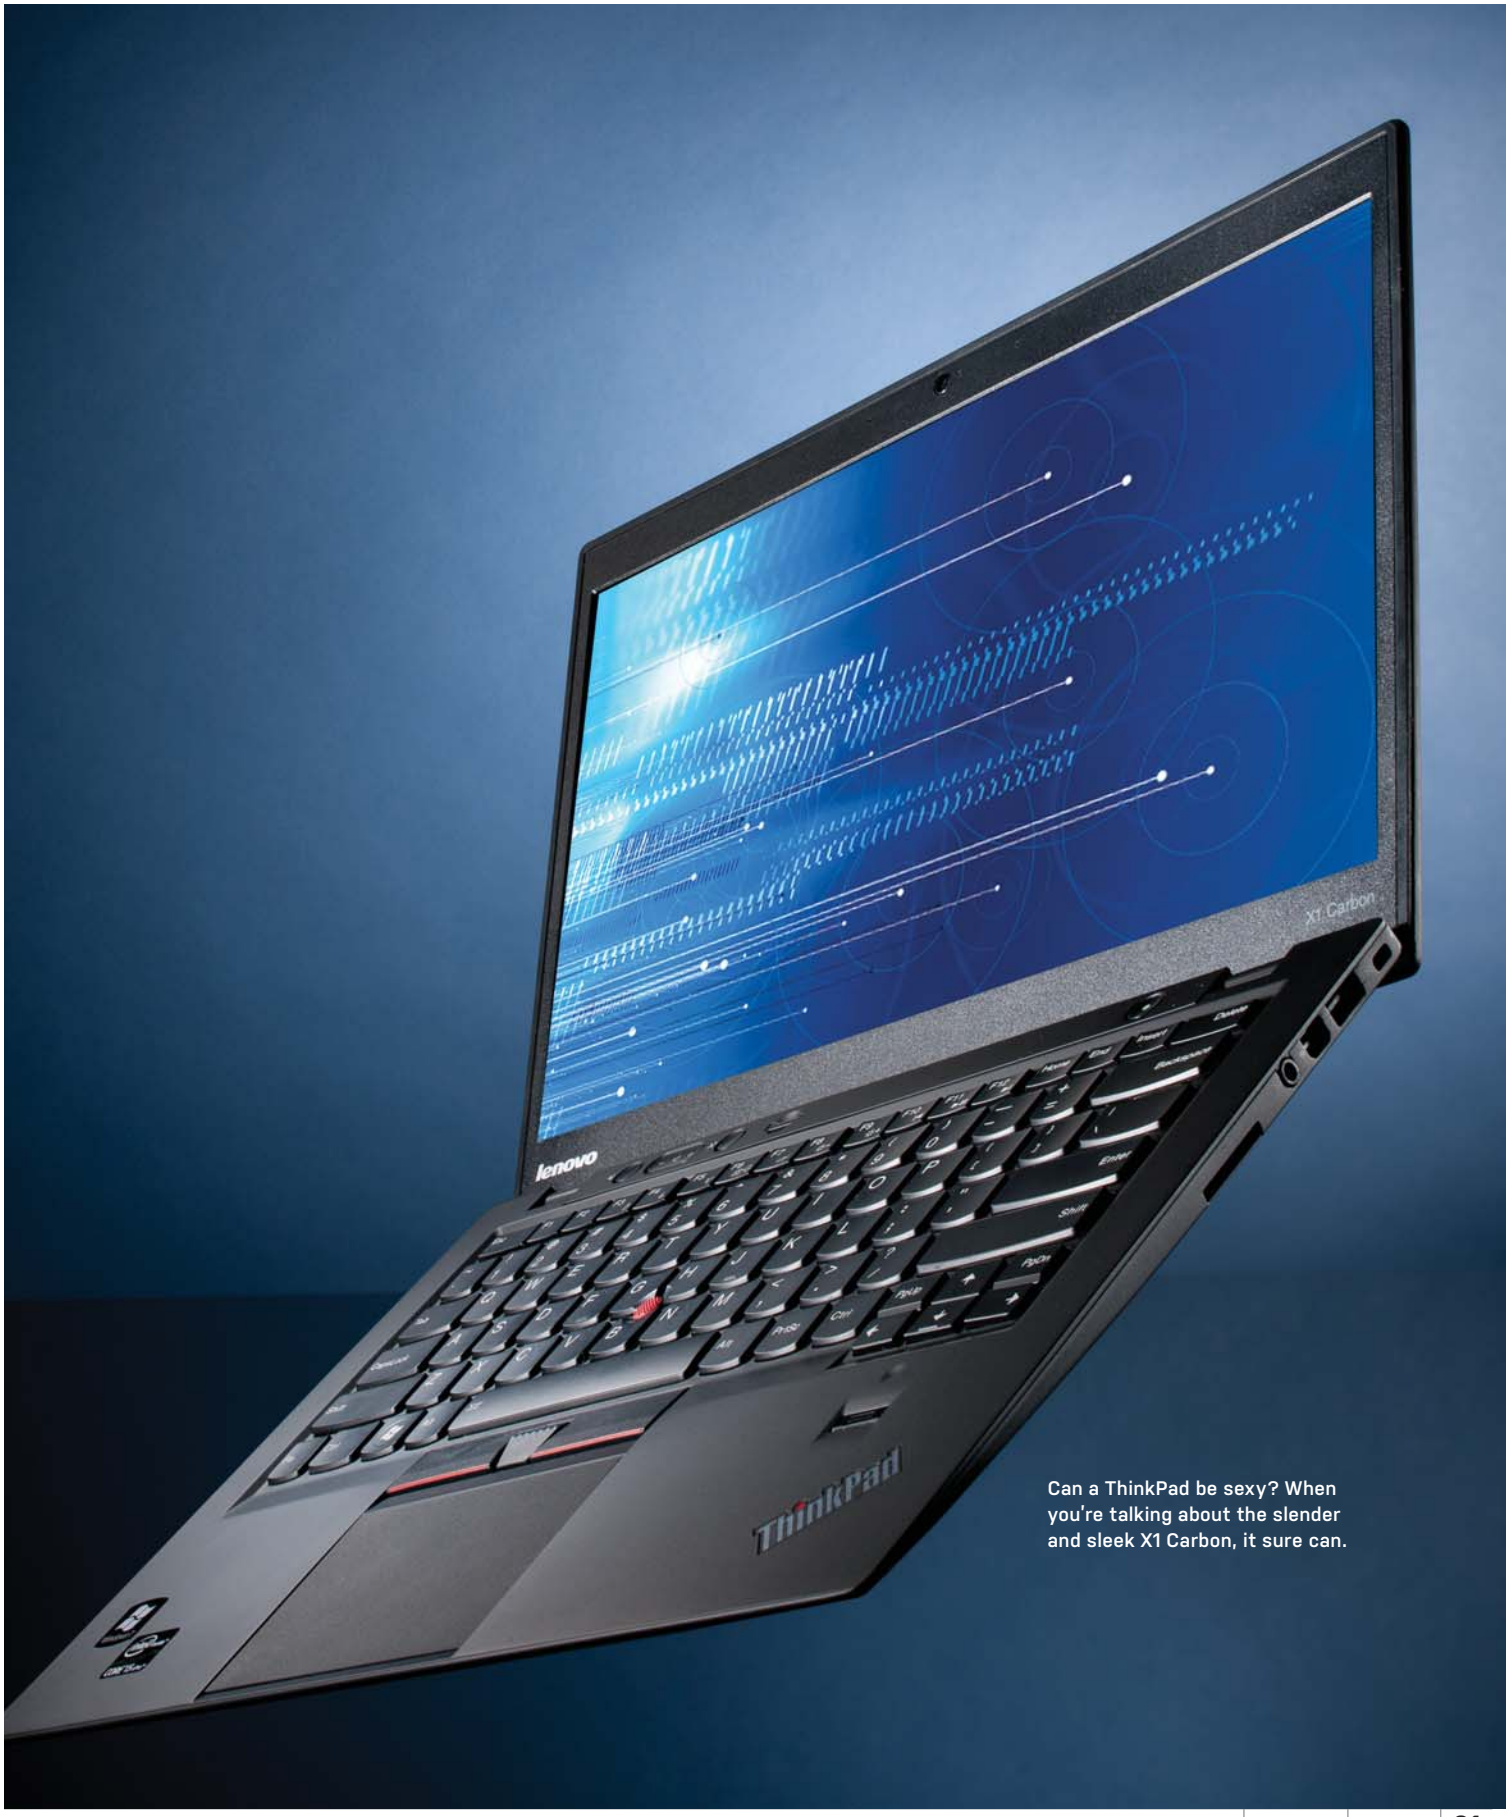

Can a ThinkPad be sexy? When you're talking about the slender and sleek X1 Carbon, it sure can.

MSI GT60

Won't break your back or your bank account

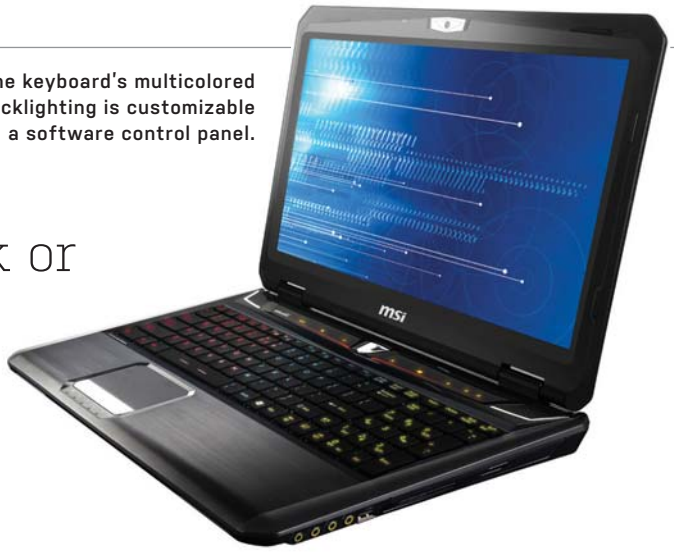

The keyboard's multicolored backlighting is customizable via a software control panel.

WE'VE BEEN so inundated with Ultrabooks these days that we almost forgot how powerful, and hulking, a full-on gaming notebook can be. MSI's GT60 arrived in our Lab to remind us. At 15.6 inches, the GT60 is not the biggest of the big, but it's a beast nonetheless with a 15.5x10.5x2-inch body and a 10-pound carry weight.

Aesthetically, the GT60 is almost entirely black plastic, with some gloss here and there for accent and a black brushed-aluminum panel on its lid. A slightly wedged profile and a multicolored backlit keyboard are probably its most distinguishing features. It's not a stunner, but it's not offensive, either.

Under the hood, the GT60 is rocking an Ivy Bridge 2.3GHz Core i7-3610QM, with a Turbo ceiling of 3.3GHz. It's joined by a GeForce GTX 670M, which is based on Nvidia's 40nm Fermi GPU as opposed to the newer 28nm Kepler. Regardless, the combo decimated our zero-point in the benchmarks by embarrassing margins. It told us what we already knew: that our current notebook benchmarks and zero-point are woefully inadequate for measuring today's performance portables. So we tossed a few of our desktop benchmark's the GT60's way.

In the Stitch.Efx and ProShow tests, the GT60 was only about 20 percent behind our desktop zero-point, which has a Core i7-3930K hexa-core proc. That's because the two apps aren't optimized for more than four cores. In the x264 HD encoding test, the GT60 was at more of a disadvantage to the hexa-core desktop part. In 3DMark 11, the GT60's 670M couldn't hold a candle to the zero-point's GTX 690 dual-GPU desktop part, not surprisingly, turning out a score of X927 vs. the ZP's X5,847. Again, we're not suggesting that a gaming notebook should be compared to a big fat desktop, but we wanted you to have some point of reference, as our zero-point notebook is in the weeds here.

We also ran the STALKER: CoP benchmark on the GT60 and it turned out a respectable 32.8fps using the Ultra setting at 1920x1080, the notebook's native res. Indeed, the GT60 seems up to the rigors of most modern games at the highest settings.

Gamers will appreciate how a touch-sensitive "Turbo" button above the keyboard throttles up the GPU and boosts performance by approximately 4 percent. A touch-sensitive fan button in the same vicinity increases cooling—as well as noise. And Nvidia's Optimus technology automatically switches between GPU and integrated graphics where appropriate.

As is fitting for a gaming notebook, the GT60's Dynaudio Premium speakers and subwoofer, along with a THX software control panel, produce good sound and achieve a satisfyingly loud volume. Less typical of gaming notebooks is the screen's matte finish, but we don't mind, as it eliminates distracting reflections. We also don't mind the better-than-average battery life. Often these desktop replacements top out at around two hours in our video rundown test; the GT60 exceeded three hours.

One thing we aren't crazy about is the keyboard. Despite this being a large notebook, the keys seem small and the overall layout cramped. The shortened right-shift key, in particular, caused us many a typing error. We're also dismayed by the absence of any SSD storage. Given the caliber of the other parts, a fast boot drive, at the very least, seems warranted. Of course, that would add to the GT60's price tag.

At $1,500, the GT60 remains affordable and offers a sound gaming system/desktop replacement, if not an outstanding one.

–KATHERINE STEVENSON

SPECIFICATIONS

CPU	2.3GHz Intel Core i7-3610QM
RAM	12GB DDR3/1600
Chipset	Intel HM77
GPU	Nvidia GeForce GTX 670M
Display	15.6-inch, 1920x1080 LCD
Storage	Two 500GB hard drives (7,200rpm)
Optical Drive	8x DVD burner
Connectivity	Ethernet, VGA, HDMI, eSATA, 7-in-1 media reader, 3x USB 3.0, USB 2.0, audio in, audio out, headphone, mic, 2MP webcam, , Bluetooth 4.0, 802.11a/b/g/n
Lap / Carry	7 lbs, 13.5 oz / 10 lbs, 1 oz

VERDICT 8

MSI GT60

BACKGAMMON Strong CPU and GPU performance; good battery life; nice screen.

CHECKERS Cramped keyboard; no SSD.

$1,500, www.msi.com

BENCHMARKS

	ZERO-POINT	
Premiere Pro CS3 (sec)	899	420 (+114%)
Photoshop CS3 (sec)	131	79.3
ProShow Producer (sec)	876	542
MainConcept (sec)	1,782	1,028
Far Cry 2 (fps)	48.5	68.3
Call of Duty 4 (fps)	62.2	157.7 (+153.6%)
Battery Life (min)	96	187

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point notebook is an Asus G73Jw-A1 with a 1.73GHz Intel Core i7-740QM, 8GB DDR3/1066, two 500GB Seagate 7,200rpm hard drives, a GeForce GTX 460M, and Windows 7 Home Premium 64-bit. Far Cry 2 tested at 1680x1050 with 4x AA; Call of Duty tested at 1680x1050 with 4x AA and 4x anisotropic filtering.

Corsair's Neutron GTX packs a punch via a new controller from A Link Media.

Flash Fight

Corsair and Samsung debut new SSDs and controllers in a battle for SSD-ominance

The pace of development in the SSD world is staggeringly awesome, as each generation of SSD controllers has delivered substantial increases in performance and reliability, while at the same time we've seen flash prices drop like a stone. It's a great time to be storing and accessing data, for sure, but we've also seen the market dominated by a trio of SSD controllers from SandForce, Marvell, and Indilinx, with different vendors applying their own tweaks to the drives' firmware to differentiate them. Though these controllers are all pretty sweet, we were beyond stoked to see two brand-new drives from Samsung and Corsair arrive this month, both with all-new SSD controllers. Will either of them put a dent in the SandForce/Marvell juggernaut? Read on to find out! **–JOSH NOREM**

CORSAIR NEUTRON GTX 256GB

Corsair's Neutron GTX SSD is the first solid-state drive to arrive in the Lab sporting a brand-new eight-channel controller from a company named Link A Media Devices (LAMD for short). The controller promises top-notch I/O performance, especially in a multitasking environment, making this a drive specifically targeted at hardcore enthusiasts—which is you, since you're reading this.

The Neutron GTX uses Toshiba 24nm MLC toggle-mode NAND flash memory and includes some enterprise-level technology designed to enhance reliability. It also supports the Windows Trim command and employs a routine of wear-leveling and garbage collection to help maintain performance over time. The LAMD controller uses a multicore ARM microcontroller, and hooks up via a SATA 6Gb/s interface.

The enclosure itself is a slim 7mm unit, so it will fit in any notebook or Ultrabook, and if you want to stick it in your PC, Corsair provides a 3.5-inch bay adapter. It also features a metal chassis, giving it a rigid feel while keeping it amazingly light at the same time.

During testing, the GTX posted impressive benchmark numbers across the board, making a remarkable debut appearance on our benchmark chart. In sequential read tests it averaged 435MB/s, which is decent but not earth-shattering. Its sequential write speeds, however, make it the second-fastest drive we've ever tested, losing only to the Samsung drive, and just by a hair.

When the workload increased to 32 commands in a queue, the GTX kept up relatively well but couldn't keep pace with the SandForce-equipped Patriot and Intel drives. In our simulated real-world test, PCMark Vantage, the GTX achieved the second-fastest score ever, behind the new Samsung. Overall the only area where it lagged in any significant way was in heavily queued 32-command workloads.

The GTX is an impressive drive, and we can't wait to see if other manufacturers jump on the LAMD train. Even though the Samsung 840 Pro steals its thunder a bit, Corsair's new SSD is still one of the fastest we have ever tested.

VERDICT **8**

Corsair Neutron GTX 256GB

$250, www.corsair.com

We loved Samsung's 830 Series SSDs, and we love its new 840 Series even more.

BENCHMARKS

	Corsair Neutron GTX	Samsung 840	Patriot Pyro SE	Intel Series 520	OCZ Octane
Controller	Link A Media	Samsung	SF-2281	SF-2281	Indilinx Everest
Capacity	240GB	256GB	240GB	240GB	512GB
Price	$250	$270	$250	$240	$750
CrystalDiskMark					
Sustained Read (MB/s)	434	**516.9**	482	470.6	445.4
Sustained Write (MB/s)	469	**497.9**	300.3	299	315.5
AS SSD					
4KB Read (IOPS)	5,675	**7,225**	4,986	5,655	5,546
4KB Write (IOPS)	11,423	12,579	**14,179**	14,123	10,417
ATTO					
64KB File Read (MB/s)	494	**517**	443.24	422.81	408.57
64KB File Write (MB/s)	393	**545**	487.9	490.29	287.02
Iometer					
4KB Random Write (32QD)	46,694	49,238	**91,171.26**	87,713.73	22,073.97
Premiere Pro Encode Write (sec)	**421**	**421**	424	423	425
PCMark Vantage x64 HDD	64,060	**68,160**	61,686	49,622	57,030

Best scores are bolded. Our current test bed is a 3.1GHz Core i3-2100 processor on an Asus P8P67 Pro (B3 chipset) running Windows 7 Professional 64-bit. All tests used onboard 6Gb/s SATA ports with latest Intel drivers.

SAMSUNG 840 PRO 256GB

Samsung brings considerable resources to bear as the world's largest manufacturer of flash memory, and it has quite a stellar track record in this market, as well. The 470 Series in early 2011 earned a verdict of 8 from us, and the 830 Series, which came out at the end of 2011, received not only a 9/Kick Ass but a spot on our vaunted Best of the Best list as the fastest SSD money can buy. That's a tough record to beat, but the company is trying to do just that with its all-new 840 Series SSDs, available in Pro and non-Pro versions. We focus on the Pro model in this review.

Unlike Corsair and other vendors that must purchase drive controllers from a third party and then tweak them to fit the flash memory that was also purchased from a third party, Samsung builds its own controllers from scratch. It also builds its own NAND flash. And its own DRAM. So the entire package is 100 percent made by Samsung. Don't get us wrong, we're not criticizing how Corsair, OCZ, and others assemble SSD drives from third-party vendors, but Samsung is uniquely capable of putting the whole drive together by itself, which it claims provides better performance. After testing the 840 Series drives, it looks like Samsung might be onto something here.

Like the Corsair Neutron GTX, the 840 Pro comes in a 7mm chassis and uses toggle-mode NAND flash, though it's built on a smaller 21nm process than Corsair's 24nm flash from Toshiba. Samsung uses a multicore controller as well, allowing for increased efficiency as the drive fills up. The drives are actually quite close in terms of build and specs, so their close performance isn't much of a surprise.

Even though the two drives were quite close in benchmark performance, the Samsung drive had a small edge in every single test we ran. In sequential read and write tests the 840 showed us the fastest scores we've ever seen, with its closest competitor being both the Neutron GTX and the SandForce Patriot Pyro SE. In our 32-command-queue tests, the Samsung also trumped the Neutron GTX but lost big-time to the SandForce-based drives, as they are simply head and shoulders above everyone right now. These were the only tests in which the Samsung failed to set a new record.

In the final tally, the Samsung 840 Pro took top honors in seven out of nine benchmarks, so it looks like it'll be taking its predecessor's place on our Best of the Best list. Samsung also claims this drive is energy-efficient and "robust," which is great news, but in the end, just icing on a very, very fast cake.

Samsung 840 Pro 256GB

$270, www.samsung.com

Azza Genesis 9000

Our love for this feature-packed case runs deep

HOW DO WE love thee, Azza Genesis 9000 full-tower case? Let us count the ways. Is it your gorgeous, pearl-white (or charcoal) chassis? The thin, blue or red lit-up lines that adorn your body and provide us with fond memories of the oh-so-delightful *Tron: Legacy*? The ample space within your interior that allows us to install a motherboard, a kitchen sink, and two power supplies at once? The list continues.

The steel-and-plastic chassis weighs a bit more than you might expect. Standing just over two feet in height—with a reasonable depth that's just shy of two feet—this is not a case you should (or would want to) keep hidden under a desk. Its design is lovely and lit-up without dipping a toe into tacky, as evidenced by the slick detail of the case's top ventilation fins sandwiched between elegant strip lighting.

When you pop off the left side panel, you'll be confused to find yourself staring at the back of your motherboard tray. That's right: Azza's flipped everything around. Not only do you install components from the right side, you also slap them into this spacious chassis upside-down. The motherboard's CPU and cooler hover over the case's bottom two 14cm fans, directly to the right of the case's standard power supply bay (if you need room for two, you just have to take out the rightmost 14cm fan). Two larger 23cm fans expel warm air out of the case's top, but they can also be removed to make room for a mammoth 48cm liquid-cooling radiator.

The aforementioned fans—and one additional 12cm fan—come pre-wired to the case's built-in on-off switch. This button, glowing-blue outline and all, mirrors the case's power button on the top-front of the chassis. While we appreciate the ability to transform the case from shining star to silent beast, we would have preferred separate toggles for the lights and fans—and some way to minimize the din of the system's cooling instead of eliminating it altogether.

Installing hard drives and 5.25-inch components into the case's five and four bays, respectively, is a snap. Built-in locking mechanisms easily hold most of your parts in place. While three hard drives require you to first screw them into provided trays, it's a mild bummer that's quickly offset by the case's two built-in hotswap drive bays.

Cables for the case's two top-panel USB 3.0 ports, two USB 2.0 ports, and one eSATA port come tucked away on the side of the Genesis 9000. There's ample room for cable and tubing management beyond that, and we especially enjoy the hole cut into the back of the motherboard tray for easy CPU-cooler management—if you don't feel like sliding out the entire tray, that is.

To say that there are more features on the Azza Genesis 9000 than we have room to write about isn't an exaggeration: That includes brackets that help hold your videocards in place, rubberized tubing holes, and how Azza perfectly wraps a power cord extension around the case's bottom (so you can still plug in your system from behind with its front-mounted PSU setup). So, we'll keep our conclusion brief: This chassis is a must-have, period. **–DAVID MURPHY**

VERDICT 9 KICK ASS!

Azza Genesis 9000

+ SEGA GENESIS Beautiful design; plenty of space; excellent cable management; inverted for enhanced cooling.

– SEGA CD No fan-speed adjustment other than off; a bit loud; installing hard drives outside of hotswap bays requires screws.

$180 (retail), www.azzatek.com

No, we didn't invert the image. You actually install your components into this case upside-down: perfect for PC building enthusiasts on the International Space Station.

There's no need to poke holes in the firewall with the DCS-5222L.

D-Link DCS-5222L

Build your own 1,000-eyed beast

D-LINK'S STRATEGY to get you into its MyDlink IP cameras is quite insidious. First, you buy one of the company's low-cost cams to look at the baby in the crib. Then, you buy one to point out the front door. Pretty soon, you have a 1,000-eyed beast that would make Orwell envious.

Well, D-Linker, prepare to add the new DCS-5222L to your many-eyed monster. The DCS-5222L is the latest camera in the MyDlink fleet and brings with it pan, tilt, and zoom capabilities as well as on-board storage for video and still images, visible infrared night-vision imaging, and a passive infrared motion detector.

Pan, tilt, zoom (PTZ) isn't new to D-Link's webcams, but the DCS-5222L is the first with "HD" resolution. In this case, that translates to 1280x800 pixels.

MyDlink-enabled cameras are different from many webcams in that you don't have to poke holes in your firewall or any of that rigmarole. The camera requires that you set up an account with MyDlink, which lets you stream live video and sound to a browser that supports Java.

Setup isn't as horrible as with older IP cameras, but also not always snag-free. We had to update the firmware on the camera before we could get it up and running. If ease-of-setup on the Dropcam that we reviewed in September is a 10, the DCS-5222L is probably a 7.

Image quality of the DCS-5222L is OK, but also not on par with the Dropcam. In fact, it's difficult to even get HD out of the D-Link at all remotely. We could only access higher resolutions by going into the camera's setup screen. Like the two previous D-Link cameras we've tested, the DCS-942L and DCS-932L, the camera tends to switch to night mode too quickly, leaving you with a black-and-white image rather than color. D-Link provides a free Android and iOS app for viewing and controlling the camera—but also only at low resolutions.

The camera includes 802.11n, Ethernet, an external speaker jack (so you can talk through the camera), and a microSD slot for recording images and video. Since the camera doesn't actively stream all video all the time like the Dropcam, capturing video to the microSD card is key. You can also set up a local machine using free D-Link software to record all video, but most folks will likely access video remotely.

We must also note that the DCS-5222L doesn't carry the monthly service fee that seems to be a trend among IP camera providers these days.

Control of the camera is acceptable. There's about a three-second lag between telling the camera to move and when it does.

Overall, the DCS-5222L is a decently priced and decently featured IP camera. The installation is workable and the image quality and access are actually pretty reliable from either phone or web browser. The DCS-5222L's best place is probably alongside others in a whole network of MyDlink cameras. **–GORDON MAH UNG**

VERDICT 8

D-Link DCS-5222L

+ ARGUS Doesn't require monthly surcharge to run.

– MEDUSA Clunky web interface; switches to night mode too easily.

$250, www.mydlink.com

Nothing says "Autobot BFF" like a Transformer the size of a freakin' building. Not the guy to hack off in this game.

Transformers: Fall of Cybertron

Oh, we've got trouble, right here in Autobot city

THERE'S SOME magical quality about the Transformers brand, a wonderful beauty in the idea that giant, walking Rock 'Em Sock 'Em robots—who absolutely hate each other, we note—can transform into the coolest of cars, the heaviest of machinery, the biggest of guns... or even larger, walking Rock 'Em, Sock 'Em robots.

But if you, like us, have spent the last few years watching Michael Bay transform the franchise into a toilet, you're probably a bit skeptical about anything Transformers-related that hits the shelves or screens nowadays. Worry not. You need only play Transformers: Fall of Cybertron to rekindle your love affair with all things morphing, punching, and shooting.

Unfortunately, while Transformers: Fall of Cybertron might be an excellent Transformers game, it's not an excellent game in general. It's a fairly nondescript shoot-'em-up that fits the genre's trappings to a giant, Transformers-size T, and that's just the single-player campaign. Multiplayer has all the liveliness of a quarter-filled Energon cube—even given the ever-present bit about being able to seamlessly transform from robot to vehicle at a moment's notice.

Developer High Moon Studios takes the unique approach of framing most of the game's single-player levels from the perspective of a single Transformer. And, no, you're not stuck just fighting as the good guys—thankfully. One can only take so much Optimus Prime preaching.

The game switches from the friendly Autobots to the mean Decepticons about halfway through, until the very last level, which is a kind-of schizophrenic-like romp between each faction. The game's big finale forces you to make a definitive choice that you can probably see coming a mile away. No spoilers, but you don't have to be Perceptor to figure out how this one wraps up—except for the ending, the ever-bouncing story is one of the game's more compelling elements for Transformers fans.

While we love the game's lovely look and feel (minus the pre-rendered cutscenes, in which the quality varies greatly), the raw mechanics—basically, one Transformer with one special power that you always use throughout the level—forces a single gimmick on players that starts to grow tiresome in each of the game's 13 hour-long levels.

In some cases, your character feels absurdly overpowered, like when you gain access to nearly unlimited artillery barrages in chapters two and three. Fun, but not satisfying. When you're intentionally transformed into a robot god, however, it's a blast: Controlling Bruticus is a delightful reward for the borderline-tedium of steamrolling your way through Cybertron.

Combat in Transformers: Fall of Cybertron feels very Gears of War. In fact, the entire plot of the first third of the game—discovering a lake full of Energon at some processing plant, only to have to fight off waves of Insecticons while doing so—screams Imulsion and Lambents.

You get a ton of options to customize the look of your multiplayer Transformers character, but we wish you could pick from an assortment of vehicles to transform into.

In the Transformers world, it's never too close for missiles—but you can still transform and switch to guns if you really want to.

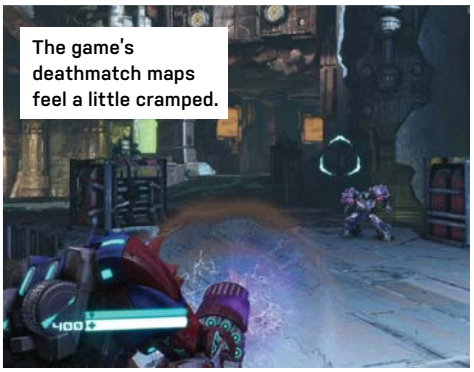

The game's deathmatch maps feel a little cramped.

In all seriousness, much of the game feels like you're hopping from combat zone to combat zone and just waiting out hordes of enemy robots—while waiting for your Halo-like shield to recharge. The game doesn't actually have a cover system, or even a way to duck. Instead, you just maneuver your Transformer behind an object and hit a key to swap your gun between hands, then shift the over-the-shoulder view from side-to-side. The game's major strategy eventually becomes, "How easily can I cheap-shot that robot by angling my gun just barely over the top of a box?"

When the game tries to get fancy—like the big "stealth" section during Cliffjumper's chapter—it doesn't feel very polished. Deus Ex: Human Revolution, this bit is not. I frequently found it easier to use the Autobot's stealth capabilities as a quick means for an absurdly obvious one-shot kill, like cloaking right in front of a Decepticon, moving slightly to its rear, and bashing in its robotic noggin' via one-button kill.

The very premise of Transformers: Fall of Cybertron might center on artificial intelligence, but this game's enemies don't appear to be all that aware of their surroundings (especially when you take out one of their peers all of a few feet away; don't giant robots make noise when they hit the ground?).

For whatever reason, these giant, walking robots don't have any place to store much ammo or weaponry—even Gears of War's meager humans could support four items of destruction and plenty of ammo. You only get to hold two weapons at a time as a typical Transformer, which you swap out at one of the many Teletran 1 kiosks littering each level. It's here where you go about the usual buying new weapons, buying perks, and upgrading stuff part of most shooters. The game's currency, Energon, comes from your dispatched enemies or crates and other items you destroy within missions.

No, it doesn't get much sillier than watching a giant, walking tank punch stacks of crates for cash—except maybe when you have to purchase access to a gun that you're *already holding*.

While its core "transforming" mechanic is certainly fun, the game might have been better as less of a forced romp through a small, specific set of terrain, and more of a choose-your-own-path, objective-driven shooter. Or, dare we say it, a Lego-style game: You get 15 of your favorite Transformers to choose from to complete a chapter, with the levels designed to accommodate those who like stealth, those who like maneuverability, or those who like blowing holes in walls and robots alike. It would make Transformers: Fall of Cybertron worth more than a single play-through.

The game's multiplayer mode does little to hook you into its long-term, character-leveling experience. Deathmatch levels feel cramped, game modes are archetypal and stale, and even its "unique" Escalation mode is just a small and simplified version of Gears of War's Horde Mode—we'd rather play Super Monday Night Combat. Yes, you can customize your own transformer, but critically missing is the option to select from a ton of vehicles or forms you'd want it to turn into. We don't care about paint jobs and armor. We want MPCBot to transform into a giant T-Rex.

While Transformers: Fall of Cybertron isn't going to win any awards for its gameplay, its standardized action elements are lifted by all the fun and unique Transformers tie-ins: from the great story, to the epic music and character cameos, to Peter Cullen himself voicing everyone's favorite semi truck. Don't play this game because you want a realistic action shooter devoid of all those stereotypical "finding 85 hidden things per level" bits. Play this game because you love Transformers and you want to experience a new story-with-shooting that's presented enjoyably, not excellently.

–DAVID MURPHY

6

Transformers: Fall of Cybertron

+ MEGATRON (GUN-MODE) Strong graphics and soundtrack (minus some pre-rendered scenes); plenty of well-designed characters and cameos; compelling storyline.

– MEGATRON (HUGO WEAVING) Stale shooting gameplay; a bit easy on standard difficulties; one-character, one-level mechanic gets a little tiresome; multiplayer lacks unique Transformers experience.

$60, www.transformersgame.com, ESRB: T

LAB NOTES

JIMMY THANG **ONLINE MANIGING EDITOR**

Steam's Big Picture Mode Is a Big Deal

Is it also a console killer?

VALVE HAS teased an HDTV-friendly version of Steam for a while now, but with Steam Big Picture Mode finally out in open beta, I'd argue its awesomeness could change PC gaming, and even the entire industry as we know it.

For the uninitiated, Big Picture Mode is not simply a re-skin of the desktop Steam we know and love; rather, it looks like an entirely new OS that's built from the ground up for HDTVs. While it works fine with a keyboard-and-mouse setup, Valve has clearly designed Big Picture Mode for Xbox 360 controllers, as its various buttons seem tailor-made for gliding through the interface's menus. Considering Big Picture Mode features its own built-in web browser, this also means you can easily use your PC to watch YouTube, Hulu, and Netflix movies from the comfort of your couch. Because it's such a solid PC solution, it also begs the question, "Is this what the rumored 'Steam Box' will use?" See my guide on setting up Big Picture Mode at bit.ly/Ovz6xp.

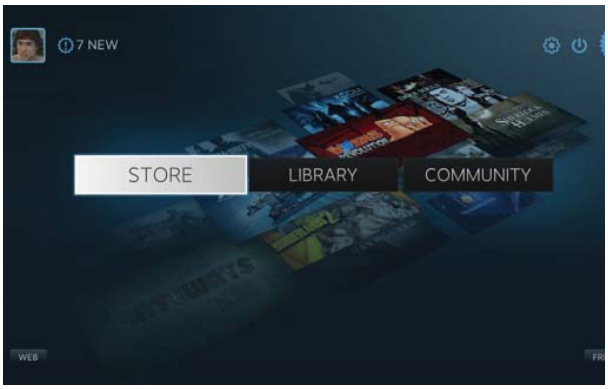

Josh Norem
Senior Editor

This month I split my time between testing the new GeForce GTX 660 GPUs and new SSDs from Corsair and Samsung. Sadly, it doesn't look like there will be any new GPUs on the horizon for a while, so I'll be hibernating in the Lab this winter evaluating SSD benchmarks until I can see a GPU's shadow again.

Katherine Stevenson
Editor-in-Chief

The Hardware Hall of Fame poll that we posted on Facebook (see page 93), got me thinking about gear that changed my life, and I have to give props to the Dell 2405FPW. Moving from a 19-inch screen to this 24-inch beauty made me feel like a baller. Of course, now I need at least two large screens to be satisfied.

Gordon Mah Ung
Deputy Editor

I'm gathering survival items for an upcoming zombie-game feature and I've decided my coworkers aren't going to make the cut. If they're afraid to eat expired emergency rations, how can I count on them to hack the head off a zombie or do the right thing if, say, Josh gets bit?

Richard Koscher
Art Director

The more hype there is around Apple products, the more I am getting turned off of the brand. As mentioned before, I've been using Apple products since 1995, but I can't identify with those people in line anymore. What a bunch of lame *&^%#. The first order of business when my contract is up is to switch to a Windows Phone. And I mean it!

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> GPU Sound Check
> Win 8 Hate
> Misleading Ads

GPU Sound Check?

I would appreciate it if you could add a reasonable and consistent rating of the min/max noise levels for all videocards you review. You rarely talk about this factor in your reviews. Personally, I prefer a quiet videocard to one that sounds like a jet taking off. I realize there are those who will sacrifice anything and everything for the sake of speed. However, there are also people like me who value silence as well as performance. Without such a rating we are left in the dark.

—Robert Tanner

EDITOR JOSH NOREM RESPONDS: Though our name is *Maximum PC* we are actually *not* in the camp of people you are referring to, Robert. We will not forgive or ignore a super-loud cooling mechanism unless there are very special circumstances, and even then we will only give a product a teeny, tiny bit of latitude. Case in point: the GTX 690 from Nvidia. It's a bit noisy, but it is also a gargantuan dual-GPU card, with only a single fan placed in the center of the cooling apparatus, so when the card is stressed, the fan gets as loud as any card in recent memory. But this is a dual-GPU card, and it's as rare as a unicorn, so we cut it a little bit of slack. The noise is not unbearable, but noticeable. All the other cards we've reviewed lately, from the high end to the midrange, have all been remarkably quiet, and we do mention a card's noise level in each review. We do not list a specific decibel level, and we won't be able to do so in the future, as we simply don't have the resources to do such testing properly—soundproof room, high-end microphones, etc. However, we will always mention a card's noise level in a review, and if it's noisy, we'll certainly make note.

Calling Bunk on Chromebird

I've been a subscriber since the magazine was known as *boot*, and it keeps on getting better. It's because of my longtime subscription that I noticed an error in Nathan Edwards's article in the September edition (Build It). The 2008 Dream Machine that he modified wasn't chrome-plated, it was nickel-plated. Maybe he should rechristen the machine the "Nicklebird."

—Ian J. Plamondon

EDITOR-IN-CHIEF KATHERINE STEVENSON RESPONDS: Good catch, Ian. Despite the fact that a few of us editors were here when that Dream Machine was first built, that technicality escaped us. Thanks for keeping us honest!

Touch This, Windows 8

Gordon Mah Ung's comments about Windows 8 in the October 2012 issue (Ed Note) underscore precisely what is wrong with this abomination of an operating system. He feels that the Interface Formerly Known As Metro is "a surprising joy to use" on a touchscreen.

There's no doubt that Windows 8 is Microsoft's touchscreen operating system, but his positive comments amount to painting lipstick on a pig.

What about the rest of us who use *real* computers to actually get something done—those of us who represent Microsoft's largest *current* customer base? In a ruthless and reckless effort to grab a piece of the iPad and Android tablet markets, Microsoft is throwing us under the bus. The current desktop paradigm has evolved gradually and is now ubiquitous *because it works*. A tablet/hybrid device isn't the most useful tool for design and simulation work, video and audio editing, complex spreadsheet work, equation editing, technical writing, publishing, architectural drawing, photography work, or even serious gaming.

Sure, I can toss my beautiful huge monitors and buy touchscreens, but why would I do that, especially when a touch interface is not what I want or need? Even if I had applications where touch would be a practical option, I still have no desire to lean forward across my keyboard to finger, poke, diddle, swipe, and smudge my display.

To me, the Windows 8 interface is an abortion on the non-touch-enabled desktop/laptop platform. Foisting it on *everyone,* regardless of need or hardware platform, is insanity. If I'm going to be forced to relearn an operating system from scratch, then like Microsoft, I too will likely opt for a more radical option.

—S. Hiverz

DEPUTY EDITOR GORDON MAH UNG RESPONDS: You won't get much disagreement

↘ submit your questions to: comments@maximumpc.com

> I HAVE NO DESIRE TO FINGER, POKE, DIDDLE, SWIPE, AND SMUDGE MY DISPLAY

from me: I did my own screaming at my monitor when trying to figure out the super-secret handshake for using IFKAM when using a mouse and keyboard. In fact, the whole staff comically stood around a beautiful nontouch 30-inch panel painfully trying to figure out how to navigate in IFKAM and it didn't go over well. I can't foresee someone tossing two 30-inch panels to get touch. And yes, forcing the interface on mouse/keyboard users is going to royally piss people off.

However, I stand by my outlook that IFKAM is pretty compelling for touch users. So, keep your monster content-creation rig with its 30-inch panels running Windows 7—even I'm keeping mine. But on a new all-in-one or convertible tablet with touch panel, I look forward to IFKAM.

An Ad in Edit Clothing

I'm curious why there was an ad structured to seem like regular content in the magazine on page 20 of the October 2012 issue. It seemed to be part of the Doctor section. It was for a bogus application called iTVmediaPlayer.

This application is just a scam. It aggregates sites one can already get online (some of which are not always legal—e.g., JustinTV.com) and plays it through a window with their own ads attached. There is no benefit from using this program.

You seem to have been misled to include this ad on the Doctor page. You should not have an ad for this software anywhere in you otherwise excellent magazine.

—Carleton Tanner

EDITOR-IN-CHIEF KATHERINE STEVENSON RESPONDS: Because of the strict separation between advertising and editorial that we observe at *Maximum PC*, we only discovered this ad after the magazine was printed, and we were just as offended as you. In the future any content of this nature will clearly be labeled "Advertisement." I apologize to anyone who was misled.

[NOW ONLINE]

OPTIMIZE CHROME

Let's face it, the light-and-fast Google Chrome browser is the only way to web surf—no question. But whether you're new to the browser or an old hand, we've got some tricks to improve your mileage. Check out our Google Chrome Optimization Guide for ways to tweak settings you didn't even know were there. bit.ly/QqpRNz

Facebook Polls

Reader Picks for the Hardware Hall of Fame

What components do our Facebook fans feel are deserving of special honors? Here are the most popular mentions.

TOP PICK IN ALL CATEGORIES

Nvidia GeForce 8800 GT

HONORABLE MENTIONS

GRAPHICS
3dfx Voodoo 2 and 3, ATI 9800 Pro, Riva TNT

CPU
AMD K6-2, AMD 64 FX, AMD Thunderbird 1.4GHz, Intel Pentium III, Intel Core 2 Quad Q6600

STORAGE
Quantum Bigfoot Hard Drive, WD Raptor 74GD, IBM 75GXP Deskstar

SOUND
Creative Labs Sound Blaster, Aureal Vortex A3D

What Hardware Are You Upgrading Next?

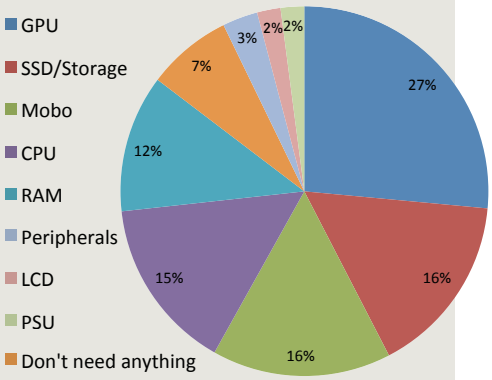

Like our page at
www.facebook.com/maximumpc

THE BUILDS

BASELINE

INGREDIENTS

PART		URL
Case	Fractal Design Define R4	www.fractal-design.com
PSU	Corsair HX650	www.corsair.com
Mobo	Gigabyte Z77X-UP4 TH	www.gigabyte.us
CPU	Intel Core i5-3570K	www.intel.com
Cooler	Cooler Master Hyper 212 Evo	www.coolermaster.com
GPU	Gigabyte GTX 660 Power Edition	www.gigabyte.us
RAM	8GB Patriot Gamer DDR3/1600	www.patriotmemory.com
Optical Drive	Samsung SH-222BB	www.samsung.com
Solid-State Drive	128GB Samsung 830 Series	www.samsung.com
Hard Drive	3TB Seagate Barracuda	www.seagate.com
OS	Windows 7 Home Premium 64-bit	www.microsoft.com

Approximate Price: $1,200

THE BIG NEWS this month is that the price of our Baseline rig has dropped a bit, going from $1,350 last month to a low price of just $1,200. The drop is due to two factors—prices for PC parts keep falling, and we swapped out only one component, trading the Gigabyte GTX 660 Ti for the less expensive GTX 660. Since this rig is built for 1080p gaming, the GTX 660 fills the bill quite well and costs $80 less than the 660 Ti. We thought about inserting a Radeon HD 7870, given recent price drops that bring it in line with the GTX 660, but the GTX 660 is less expensive and faster still, so in it goes. All other parts remain unchanged.

PERFORMANCE

INGREDIENTS

PART		URL
Case	NZXT Phantom 410	www.nzxt.com
PSU	Corsair HX750	www.corsair.com
Mobo	Asus Sabertooth X79	www.asus.com
CPU	Intel i7-3820 @4.7GHz (overclocked)	www.intel.com
Cooler	NZXT Havik 120	wwww.nzxt.com
GPU	Asus GTX 670 DirectCU II TOP	www.asus.com
RAM	16GB Corsair Vengeance DDR3/1600	www.corsair.com
Optical Drive	LG WH12LS39 BD-R burner	www.lg.com
Solid-State Drive	256GB Samsung 830 Series	www.samsung.com
Hard Drive	Seagate Barracuda 3TB	www.seagate.com
OS	Windows 7 Professional 64-bit	www.microsoft.com

Approximate Price: $1,633

THE PERFORMANCE build undergoes extremely light modifications this month, as there hasn't been a lot of new hardware introduced, and to be honest, this rig is already pretty bitchin'. We're sticking with our Sandy Bridge-E Core i7-3820, 16GB of Corsair RAM, and our Asus GeForce GTX 670 GPU, which all makes for a potent package and represents the perfect middle ground between Baseline and Ultra. We have upgraded our SSD to the 256GB 830 model from Samsung, however, as its dollar-to-gigabyte ratio has dropped, and let's be honest—a 128GB boot drive is barely cutting it these days.

ULTRA

SINCE WE dished on the parts required for a budget build last month, this month we've brought back the Ultra configuration—a towering hulk of a PC designed to shred benchmarks and your line of credit. The Ultra configuration is about two stages below Dream Machine in that its budget is still "real world," yet it is built using the best components available in every category.

The six-core Intel Core i7-3930K is running the show with a Corsair H100 water cooler keeping it frosty, even overclocked to 4.8GHz. A modest 16GB of RAM from Corsair has found a cozy home in the Asus P9X79 motherboard, which is great for overclocking and will accommodate another GTX 690 if we ever need to play Crysis 3 on three displays at once. On the storage front, we're sticking with 6TB of 7,200rpm storage for data and backup, and have upgraded our boot drive to the new 256GB Samsung 840 Series SSD, which set seven out of nine benchmark records in the Lab this month. We also switched our LG Blu-ray drive in favor of a Lite-On model, but the writing is on the wall for physical media so we'll probably remove this category soon. Finally, we're sticking with Windows 7 Pro for now even though Windows 8 has arrived. We'll definitely upgrade at some point, but probably not until SP1 arrives.

INGREDIENTS

PART		URL
CASE	Cooler Master Cosmos II	www.coolermaster.com
PSU	Thermaltake Toughpower Grand 1,050W	www.thermaltakeusa.com
Mobo	Asus P9X79 Deluxe	www.asus.com
CPU	Intel i7-3930K	www.intel.com
Cooler	Corsair H100	www.corsair.com
GPU	Asus GTX 690	www.asus.com
RAM	16GB Corsair Vengeance	www.corsair.com
Optical Drive	Lite-On BD-R burner	www.liteonit.com
Solid State Drive	Samsung 840 Series 256GB SSD	www.samsung.com
Hard Drive	Seagate Barracuda 3TB x2	www.seagate.com
OS	Windows 7 Professional 64-bit	www.microsoft.com

Approximate Price: $3,421

For our complete Best of the Best list of recommended components, visit www.maximumpc.com/best-of-the-best.

SUGGESTED PAIRINGS

Kick-ass peripherals for your new rig

PREMIUM MONITOR
Dell UltraSharp U3011
$1,200, www.dell.com

KEYBOARD
Microsoft Natural Keyboard
$40, www.microsoft.com

ROUTER
Asus RT-N66U Dark Knight
$160, www.asus.com

DVD BURNER
Samsung SH-222BB
$20, www.samsung.com

AIR COOLER
Phanteks PH-TC14PE
$90, www.phanteks.com

MAXIMUM PC (ISSN 1522-4279) is published 13 times a year, monthly plus Holiday issue following December issue, by Future US, Inc., 4000 Shoreline Court, Suite 400, South San Francisco, CA 94080. Phone: (650) 872-1642. Fax: (650) 872-2207. Website: www.futureus.com. Periodicals postage paid in San Bruno, CA and at additional mailing offices. Newsstand distribution is handled by Time Warner Retail. Basic subscription rates: one year (12 issues) US: $14.95; Canada: US$19.95; Foreign: US$29.95. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R128220688). PMA #40612608. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to Maximum PC, PO Box 5852, Harlan, IA 51593-1352. Standard Mail enclosure in the following editions: None. Ride-Along enclosure in the following editions: None. Returns: Pitney Bowes, PO Box 25542, London, ON N6C 6B2, Canada. Future US, Inc. also publishes @Gamer, Crochet Today!, MacLife, Nintendo Power, The Official Xbox Magazine, PlayStation: The Official Magazine, PC Gamer, and Windows: The Official Magazine.