KILLER APPS
On This Month's Disc

SIX CORES, GOOD FOUR CORES, BAD.
AMD's Phenom II X6 1090T Revealed! *p. 36*

DISPLAY MYTHS BUSTED
Don't believe the specs! *p. 44*

MAXIMUM PC

MINIMUM BS • JULY 2010

High-Performance UPGRADES (at Minimum Cost!)

18 Parts that prove budget components can still kick ass. *p. 22*

- Videocard
- SSD Storage
- CPUs
- CPU Cooler
- Blu-ray Drive
- 802.11n Router
- Power Supply
- 23-inch Display
- ***AND MORE!***

These three Special Edition Maximum PC coins can be yours! *Details p. 23*

ALIENWARE'S M11X
The fastest gaming ultra-portable we've ever tested! *p. 76*

EBOOK TIPS & TRICKS
Master your library across multiple devices—today! *p. 63*

0 70686 38753 6
$8.99US
07

www.maximumpc.com

Future MEDIA WITH PASSION

8 PRODUCTS = $3.7 MILLION! Ridiculously Expensive PC Hardware *p.14*

Unleash your PC's Potential...

Try MAXIMUM PC

2 FREE Trial Issues

Each issue of Maximum PC features:

- **Brutally honest product reviews**
- **Hard-hitting editorials**
- **Tips to blast your machine's performance**
- **Insightful and innovative How-To's**
- **A CD loaded with new software, utility and game demos**

Reserve your 2 FREE Trial Issues today! There's no obligation.

To order, head to:
www.maximumpc.com/archive

JULY

FEATURES

DEPARTMENTS

Quickstart

R & D

In the Lab

LETTERS

32 ▶

MAXIMUM PC

EDITORIAL
EDITOR IN CHIEF **George Jones**
DEPUTY EDITOR **Katherine Stevenson**
SENIOR EDITOR **Gordon Mah Ung**
REVIEWS EDITOR **Michael Brown**
SENIOR ASSOCIATE EDITOR **Nathan Edwards**
ONLINE MANAGING EDITOR **Alex Castle**
ONLINE ASSISTANT EDITOR **Alan Fackler**
CONTRIBUTING WRITERS **Loyd Case, Pulkit Chandna, Nathan Grayson, Tom Halfhill, Dave James, Paul Lilly, Thomas McDonald, Quinn Norton, Zack Stern**
COPY EDITOR **Mary Ricci**
PODCAST PRODUCER **Andy Bauman**
EDITOR EMERITUS **Andrew Sanchez**
CONTACT **comments@maximumpc.com**

ART
ART DIRECTOR **Natalie Jeday**
CONTRIBUTING ART DIRECTOR **Boni Uzilevsky**
PHOTO EDITOR **Mark Madeo**
ASSOCIATE PHOTOGRAPHER **Samantha Berg**
CONTRIBUTING PHOTOGRAPHER **Patrick Kawahara**

BUSINESS
VP / PUBLISHING DIRECTOR **Kate Byrne**
kbyrne@futureus.com
NATIONAL SALES DIRECTOR **Jane Evans**
jevans@futureus.com
REGIONAL SALES DIRECTOR **David White**
dwhite@futureus.com
REGIONAL SALES MANAGER **Anthony Losanno**
alosanno@futureus.com
SALES MANAGER **Aida Rodriguez**
arodriguez@futureus.com
INTEGRATED SALES DIRECTOR **Joe Pomparelli**
jpomparelli@futureus.com
SENIOR MARKETING MANAGER **Andrea Recio-Ang**
arecio-ang@futureus.com
MARKETING COORDINATOR **Allyson Kardel**
akardel@futureus.com
ADVERTISING COORDINATOR **Jose Urrutia**
jurrutia@futureus.com

PRODUCTION
PRODUCTION DIRECTOR **Michael Hollister**
SENIOR PRODUCTION COORDINATOR **Dan Mallory**
PRINT ORDER COORDINATOR **Jennifer Lim**

CONSUMER MARKETING
DIRECTOR CONSUMER MARKETING **Rich McCarthy**
CIRCULATION DIRECTOR **Crystal Hudson**
NEWSSTAND DIRECTOR **Bill Shewey**
CONSUMER MARKETING OPERATIONS DIRECTOR **Lisa Radler**
RENEWAL AND BILLING MANAGER **Mike Hill**
BUSINESS MANAGER **Elliot Kiger**
SR. ONLINE CONSUMER MARKETING DIRECTOR **Jennifer Trinker**
CUSTOMER SERVICE MANAGER **Mike Frassica**

FUTURE US, INC
4000 Shoreline Court, Suite 400, South San Francisco, CA 94080
www.futureus-inc.com
Tel 650.872.1642

PRESIDENT **John Marcom**
VP / CFO **John Sutton**
VP / INTERNET DEVELOPMENT **Tyson Daugherty**
VP / PUBLISHING DIRECTOR MUSIC **Anthony Danzi**
VP / PUBLISHING DIRECTOR TECH AND LIVING **Kate Byrne**
VP / SALES AND MARKETING **Rachelle Considine**
EDITORIAL DIRECTOR TECH AND LIVING **Jon Phillips**
EDITORIAL DIRECTOR GAMES GROUP **Stephen Pierce**
EDITORIAL DIRECTOR MUSIC **Brad Tolinski**
DIRECTOR OF HUMAN RESOURCES **Nancy Durlester DuBois**

Future MEDIA WITH PASSION

Future US, Inc. is part of Future plc. Future produces carefully targeted special-interest magazines, websites and events for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read or visit. Today we publish more than 150 magazines, 65 websites and a growing number of events in the US, UK, France and Italy. Over 100 international editions of our magazines are also published in 30 other countries across the world.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR).

FUTURE plc
30 Monmouth St., Bath, Avon, BA1 2BW, England
www.futureplc.com
Tel +44 1225 442244

NON-EXECUTIVE CHAIRMAN **Roger Parry**
CHIEF EXECUTIVE **Stevie Spring**
GROUP FINANCE DIRECTOR **John Bowman**
Tel +44 1225 442244
www.futureplc.com

REPRINTS For reprints, contact Marshall Boomer, Reprint Operations Specialist, 717.399.1900 ext. 123 or email: marshall.boomer@theygsgroup.com

SUBSCRIPTION QUERIES Tel 800.274.3421
www.maximumpc.com/customerservice
Maximum PC ISSN: 1522-4279

Unfrozen Caveman Editor

For any tech-head, Maximum PC headquarters are hallowed grounds. A sacred place. I worked here as Editor-in-Chief back in 2003 and 2004, and I couldn't be happier about being back. (By the way: There's no truth to the rumor that previous EIC Will Smith defeated me in a traditional Vulcan Death Match and sent me packing.)

Not surprisingly, Maximum PC is a different place than it was in 2004. I mean that both literally and figuratively, and it's making me feel a little bit like Unfrozen Caveman Editor-in-Chief. For starters, our office space is much improved. Ample natural and indirect light. A huge laboratory. Large test benches and tons of storage space.

But more importantly, *Maximum PC* is now online in a big way. Six years ago, being online meant you had a website. Now it means websites, message boards, Facebook, Twitter, and more. These extra dimensions might intimidate your average everyday journalist, but for us, they're just extra opportunities to talk directly to and with you about PC-based technology. Knowing full well how smart our readers are, I'm super-excited about that.

As I write this editorial, I'm looking back at my last issue of *Maximum PC*. February 2005. In it we covered a slew of Windows XP tips, an 850-watt power supply, a 40GB USB 2.0 external drive, a handful of MP3 players, and a kick-ass Velocity Micro system with a 2.6GHz AMD Athlon 64 FX-55 processor, 1GB of RAM, two 74GB WD Raptors in RAID 0, and a 250GB storage drive. (That Athlon CPU was a fave of mine, by the way.)

Wow, have times changed. I'm inclined to view the futurist notion of pending "technological singularity" with a raised eyebrow, but after contemplating the rapid technological evolution of the last six years, I'm beginning to wonder if the futurists have it right. Maybe we are just hitting the beginning of some crazy hockey-stick-shaped growth cycle. Smartphones. Touch-based computing. Cloud-based computing. Multicore processors. Solid State storage. Social networking. These aren't just evolutionary hops. They're paradigm shifts—legitimate transformations in the way computers work and the way we use them.

Singularity or not, I guarantee this: The next decade is going to be extremely interesting. I'm happy to be back at a magazine that not only "gets" these things, but embraces them.

P.S. In honor of our upgrades-oriented cover story, email me at george@maximumpc.com (or if you prefer to tweet, hit me @lizzunchbox). Tell me about the best or worst PC upgrading moment in your life. I'm going to pick the 10 best tales—you'll each get a handful of coins.

THE BIG CHEESE'S PICKS

3 Million Mobos from Nan-Ping
page 54

Bowers & Wilkins MM-1 Computer Speakers
page 80

The Power of Six: AMD's Thuban
page 36

LETTERS POLICY Please send comments, questions, and Pocky to george@maximumpc.com. Include your full name, city of residence, and phone number with your correspondence. Unfortunately, George is unable to respond personally to all queries.

THE NEWS

The FCC's New Vision for PC TV

Farewell, CableCARD; we hardly knew ye —MICHAEL BROWN

Those who believe that the interests of consumers have been routinely subjugated to the will of big business—Hollywood, cable companies, ISPs, phone companies, and so on—for most of the last decade will applaud the Federal Communication Commission's latest move. The FCC is soliciting ideas for a new technology standard that will finally integrate PCs, gaming consoles, and televisions with digital cable and satellite television services. Yes, they're talking about a replacement for the much-maligned CableCARD. The agency has dubbed this theoretical successor AllVid.

You don't need to read much beyond the first two paragraphs of the FCC's Notice of Inquiry to get excited about what's envisioned. "…the Commission seeks comment on specific steps we can take to unleash competition in the retail market for smart, set-top video devices that are compatible with all multichannel video programming distributor (MVPD) services."

By specifying "all MVPD services," the FCC is describing not just cable TV systems, but also satellite TV and any similar services that might pop up in the future. It gets better. "In particular, we wish to explore the potential for allowing any electronics manufacturer to offer smart video devices at retail that can be used with the services of any MVPD and *without the need to coordinate or negotiate with MVPDs*" (emphasis ours). The commission wants to return to the days when consumers "…could connect their 'cable ready' video cassette recorders and television sets directly to a cable operator's system without the need for any other equipment."

But the folks at Ceton Corporation, who are about to ship a brand-new quad-tuner CableCARD product (the InfiniTV 4) into retail channels, are saying hey, not so fast. It would be easy to dismiss Ceton's complaint as self-interested—after all, the company has invested considerable R&D resources to develop this new card—but Ceton has a solid argument against throwing the baby out with the bath water. While the FCC has pegged the many flaws inherent in the CableCARD spec, there's no denying that it's the only technology you can buy today that enables a PC to connect to a digital cable TV service.

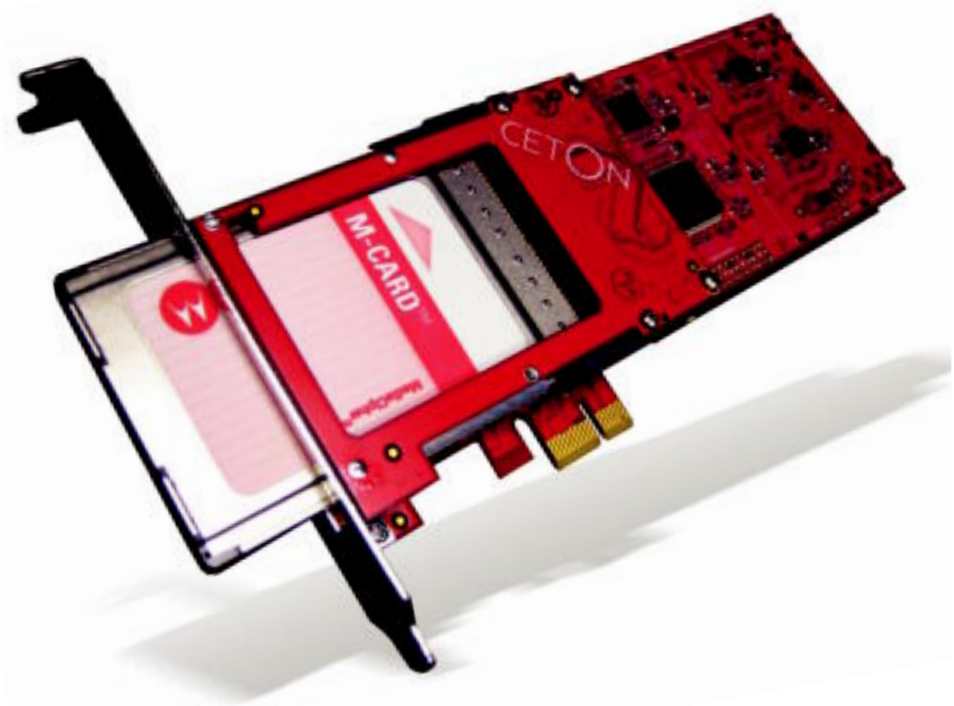

The FCC's AllVid proposal could dampen the market's enthusiasm for Ceton Corporation's brand-new InfiniTV 4 quad CableCARD tuner.

"We think it's premature to call the death of CableCARD," says Ceton's chief marketing officer, Ed Graczyk. "CableCARD answers most of the FCC's requirements. It lets us deliver on a single PC board the ability to receive and decode four [digital cable TV] streams." Graczyk admits that CableCARD has its shortcomings—including its inability to support program guides or deliver video-on-demand services, as well as spotty support among service providers. "In-the-field support is very touch and go," Graczyk says, "depending on which cable operator you have and which market you happen to be in. And the entire pairing process is so antiquated—the way you need to activate the card; there's no reason that couldn't be automated."

But it's hard to argue with Graczyk's point that a broadband-connected PC with access to services such as Netflix, Hulu, YouTube, and to a CableCARD is vastly superior to any dumb set-top box. And it's available today, whereas AllVid is just a collection of goals. "We are proponents of the FCC's ideas, but many of them can be realized with CableCARD today," Graczyk says. "The technology has been around for several years on paper, but it's only really been on the market for a couple of years."

TV on the Go

Broadcasters unite to create a national mobile TV service

Imagine getting over-the-air television programming from your smartphone, netbook, or car—no Internet connection required. That's what a newly formed coalition of national and local broadcasters, known collectively as Pearl Mobile TV, are working on. Currently, the group consists of three national television networks—Fox, NBC/Telemundo, and Ion Television—along with nine of the largest owners of local stations, which plan to pool their existing spectrum, content, and marketing resources in support of the effort.

The proposed service, which was made possible by the national adoption of the DTV broadcast standard, has received the blessing of FCC Chairman Julius Genachowski, as it assists in the Commission's goal of unclogging the nation's wireless broadband network. The broadcast group expects to have enough bandwidth to deliver mobile video to 150 million U.S. consumers, and predicts as many as 150 mobile DTV stations by the end of the year. –KS

For mobile DTV to be realized, device makers will need to build TV tuners into their products.

Google Gets a Scolding

Google has taken flak for its seeming lack of regard for people's privacy. First there was fuss over the company's intrusive Street View technology; then Buzz, Google's social-networking extension for Gmail, raised hackles for being an "opt-out" service.

Now, the data-protection authorities from 10 countries have come together to lambast the company. In a letter addressed to Google CEO Eric Schmidt, the Privacy Commissioner of Canada, Jennifer Stoddart, slammed the Internet behemoth for "introducing services without due regard for the privacy of its users."

The letter, dated April 19, is also signed by Stoddart's counterparts in France, Germany, Israel, Italy, Ireland, Netherlands, New Zealand, Spain, and the United Kingdom. Stoddart calls on Google to ensure that its services honor fundamental privacy principles and asks the company to outline ways in which it plans to ensure such conformity. –PC

MS Pulls Plug on Courier

Rumors that the company would ship the sexy dual-screen tablet early next year have been put to rest

After tech blog Gizmodo broke word of the Courier's existence last September, it gave gadget hounds—particularly those with an aversion to Apple—a tablet to dream about. Now, not only has Microsoft nixed plans for Courier, but word that HP is abandoning the Slate dashes hopes of *any* Windows tablet. –KS

Apple Catches a Hummingbird

Apple keeps getting more serious about custom-chip design. In its latest move, the company has stealthily acquired a small processor-design company called Intrinsity, based in Austin, Texas. Apple hasn't publicly announced the deal, the price paid, or its plans.

But it's not as mysterious as it looks. Lately, Intrinsity has specialized in making processor cores that are fully compatible with those from ARM, the British company whose processors are found in almost every cell phone. Apple's iPods, iPhones, and iPads also use the ARM architecture. (Of course, Macs use the Intel x86.)

Intrinsity's processor cores are generally faster and consume less power than ARM's, relative to their performance. Last year, Intrinsity introduced a core nicknamed Hummingbird. It's 100 percent compatible with the ARM Cortex-A8, one of ARM's most powerful processors. When manufactured in a 45-nanometer low-power CMOS process, Hummingbird can reach 1GHz while using less than 750 milliwatts.

That clock frequency may not seem fast when compared with PC processors, but it's the combination of gigahertz speed and sub-watt power consumption that's important for mobile devices, like smartphones and tablets.

Note that Intrinsity designs processor *cores*, not processor chips. A core is merely a building block for a custom system-on-a-chip (SoC). An SoC typically surrounds the CPU core with peripheral logic (such as USB), a graphics engine, and memory. SoCs are vital for small devices that, unlike PCs, have no room for a large motherboard with many separate chips.

Hummingbird is rumored to be the ARM-compatible processor in the iPad's custom-designed A4 chip. (Apple isn't talking.) Even if the rumor is wrong, Hummingbird will likely appear in a future iPhone or iPad.

I believe Intrinsity has developed an even faster processor that's compatible with ARM's dual-core Cortex-A9. I call it Hummingbird-II. I think this dual-core clone will exceed 2GHz, and I predict it will appear in a future iPhone or iPad.

By designing its own custom chips, Apple gains more control over its platforms and can continue differentiating them from competing products. Differentiation is the key to maintaining the higher prices and profit margins that make Apple so successful.

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

Powerline Networking Standard Advances

The IEEE moved one step closer to establishing a global standard for broadband over power-line networks as 81 percent of IEEE P1901 Working Group members voted to approve Draft 3.0 of the IEEE standard in April. The votes were accompanied by 857 comments, which the working group will discuss at their meetings in June, August, and October, when a final vote is expected to take place.

The IEEE P1901 standard is based largely on HomePlug Powerline AV technology, and the HomePlug Powerline Alliance will serve as the certification body for the IEEE P1901 standard in much the way that the Wi-Fi Alliance certifies IEEE 802.11 wireless networking devices. Current HomePlug AV products should be interoperable with future IEEE P1901 devices. **—MB**

If all goes according to plan, HomePlug AV devices such as this Trendnet TPL-303E2K will be interoperable with future products based on the final IEEE P1901 Broadband Over Powerline standard.

Adobe, Apple Trade Jabs

Following Apple's change to the terms of its iPhone 4.0 software developer kit license, which effectively blocks Adobe from porting Flash support to the iPhone and iPad, the two companies have been waging a public battle.

"As developers for the iPhone have learned, if you want to develop for the iPhone you have to be prepared for Apple to reject or restrict your development at any time, and for seemingly any reason," said Mike Chambers, Adobe's principal product manager for the Flash platform.

It didn't take long for Steve Jobs himself to respond in an open letter on Apple's website. Jobs lambasted Flash on a variety of fronts, including reliability, security, and performance, and concluded his tirade by saying, "Flash is no longer necessary to watch video or consume any kind of web content," and that, "Adobe should focus more on creating great HTML5 tools for the future, and less on criticizing Apple for leaving the past behind." Ouch. **—PL**

PC Sales Surge

Who says the PC's days are over? Sales of personal computers appear to be on a major rebound, with market researchers IDC and Gartner both reporting an increase in sales figures in the U.S. and globally.

According to Gartner, PC sales were up 27.4 percent for the first quarter of 2010 compared to 2009. Its competitor IDC, meanwhile, reported sales of PCs up 24.2 percent over the previous year.

The increase in sales is attributed to large companies replacing old machines, as well as a rise in laptop sales and even higher-performance CPUs. In the U.S., HP continued to be the number-one PC maker, followed by Dell, Acer, Toshiba, and then Apple. **—GU**

GAME THEORY

Two Stories, One Meaning

Look, I've been on the Internet since the Internet was on. In those weird early days, I saw enough of alt.binaries.ohmygodicantbelievetheyredoingthat to illumine the darker corners of humanity.

Thus, I was not at all surprised that some ostensibly human-like creatures get their kicks watching videos of attractive women crushing small animals to death under stiletto heels. It's a fallen world, and any cursory glance at history or headlines shows us that men are capable of unspeakable depravity.

I *was* surprised, however, when the U.S. Supreme Court ruled that these "crush" and dogfight videos were, in fact, protected speech. In their April decision on *United States v. Stevens*, the court voted 8-1 to overturn a ban on the sale and distribution of these videos. Chief Justice Roberts wrote that, although the acts depicted are in fact illegal, the ban on images of these creates a "criminal prohibition of alarming breadth," and violates First Amendment protections.

Why is any of this of any concern to gamers?

Simple: The Roberts court is about to get its first swing at the gaming piñata. In a case bearing the sublime name of *Schwarzenegger v. Video Software Dealers Association*, California's laws banning the sale of violent games to anyone under the age of 18 have already failed to pass muster with various courts. Most recently, the notoriously liberal 9th U.S. Circuit Court of Appeals ruled that the law was in violation of the First and 14th Amendments, clearing the path to a Supreme Court challenge.

The *Stevens* decision, although not an exact analog, is a pretty clear indication of where the Roberts court stands on freedom of speech, with obvious applications to video and computer games. After all, if eight out of nine justices can rule that video recordings of illegal acts of animal cruelty are protected speech, then the court is packed with more First Amendment absolutists than the staff of *Reason* magazine.

Of course, what the VSDA and the development community *should* be asking themselves is, "How did our content become so depraved that we're relying on animal torture videos as precedent?"

Thomas L. McDonald has been covering games for 20 years. He is an editor at large for *Games* magazine.

Intel Unlocks Cheap Chips

Company finally offers overclocking-friendly features in midrange CPUs

Intel breaks with tradition by offering two unlocked LGA1156 CPUs, and makes a serious price cut to one of them. The new 2.93GHz Core i7-875K will cost $342—down from the $562 Intel had been charging for its 2.93GHz Core i7-870. Oddly, Intel will charge $216 for the 3.2GHz Core i5-655K. That's up from the $176 it charges for the equivalent locked model. These processors should be available by the time you read this. **—GU**

Hulu for a Fee

Details of a paid subscription model emerge

If a recent report in the *LA Times* is true, Hulu is on the verge of testing its long-rumored subscription service. According to "people with knowledge of the plans," the website will offer a Hulu Plus service for $9.95 a month that will give viewers access to a comprehensive selection of the television content that has so far been available for free. Free content will still be available, say the sources, but in limited supply, restricted to the five most recent episodes of popular network shows. Content on Hulu, whether paid or free, will continue to carry commercials, and it's expected that the number of commercial spots will increase. **—KS**

LOC Archives Twitter

In April, the Library of Congress acquired a copy of the entire public Twitter archive. Why would the LOC want to store billions of possibly inane tweets forever? Wrote Matt Raymond, an LOC blogger, "Individually, tweets might seem insignificant, but viewed in the aggregate, they can be a resource for future generations to understand life in the 21st century." **—NE**

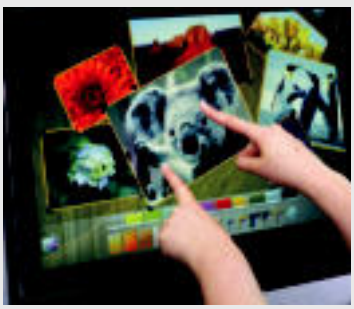

You've Got the Touch

For everyone who's been wondering how Windows 7's multitouch functionality might change their lives, Microsoft is now offering Touch Pack for Windows 7 for free (http://bit.ly/cfBDMJ). The pack lets you install and experience three multitouch casual games and three Microsoft Surface applications, including Surface Globe and Surface Collage, provided you have a Win7 PC paired with a multitouch-capable monitor. **—KS**

BYTE RIGHTS

From British Backrooms to Your Room

Back in April, the British Parliament passed the Digital Economy Act (nicknamed DEAct), a law affecting the whole of the British Internet. It gives everyone in Britain lit fiber, unlimited *Doctor Who*, and a Corgi. Just kidding. It takes away their human rights.

DEAct doesn't really kick in for a while, but it will require national monitoring of ISPs, and also require ISPs to monitor their users. People accused of file sharing face account capping and suspension, unless they have the time and money to fight it, because punishment comes before trial in DEAct. The presumption of innocence is enshrined in hundreds of years of Anglo-American law, and pretty much every human rights declaration, but the Labour government just threw that out for any Brits who use the Internet. It's a bit like the British government turning off your electricity because someone accused you of adultery. Parliament did get rid of DEAct's Great Firewall of Britain provision, which allowed national blocking of websites that might permit copyright infringement, but plans to bring it back in the new session.

How did this legislative abomination get past the reasonably sane people of Britain? It was done during the washup—a lame-duck period in the British political cycle when few representatives are even in town, and laws get passed by backroom deals with powerful interests. DEAct, with the paw prints of record labels all over it, was no exception. Like the secretive ACTA, DEAct was the undemocratic product of secret meetings between government and lobbyists.

What measures like ACTA and DEAct really show us is that rights holders realize that if they had to go through the regular democratic process, We The People would run them out of town on a rail. I don't like sounding paranoid, but they are acting like we are the enemy, to be fought and controlled, all while being bilked.

Why should we Yanks care? Because now that the idea has a foot in the door, it's sure to cross the pond soon. Let's pay attention when it does.

Quinn Norton writes about copyright for Wired News and other publications. Her work has ranged from legal journalism to the inner life of pirate organizations.

THE LIST

8 Signs a Geek Has Too Much Money

1 Leatherman Del Rey

Custom-made to order, the Leatherman Del Rey 25th Anniversary multitool is crafted by artist Adrian Pallarols and boasts handles made of 18K gold from the Andes.

$40,000

www.leatherman.com

2 CHROME-PLATED PC

Our dual-Core 2 Extreme Dream Machine 2008 isn't anything to boast about these days, but its custom chrome-plating remains the height of decadence.

$5,000–10,000

www.computer-choppers.com

3 Sharp LB-1085 LCD

You're looking at 108 inches of 1080p goodness. It comes with three HDMI ports, two component ports, and a single DVI input, so you can peel your eyelids back with multiple HD devices. It weighs 430 pounds, but if you're buying this, you're not lugging your own hardware.

$150,000

http://bit.ly/aJZV5V

4 TRANSMISSION AUDIO ULTIMATE SPEAKER SYSTEMS

Imagine if you will, six seven-foot-high panels loaded with supertweeter ribbons and subwoofers bundled together with six 500-watt dual-mono amplifiers and a preamp. If your friends object to the price, you can jack up the volume and blow them out the nearest window.

$2,000,000/set

www.transmissionaudio.com

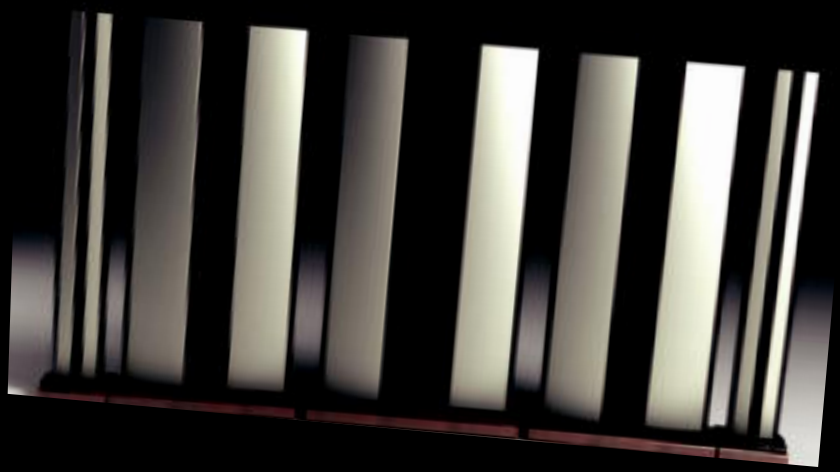

5 Xten Pininfarina Office Chair

Made by the renowned team that dresses Ferrari interiors, the Pininfarina Xten is maximized for comfort and durability, offering a Dynatec fabric seat and the patented Armtronic system, which adjusts to fit people of all sizes—but only the uppermost economic strata.

$1,500,000

(yes, you read that right!)

www.xtenchair.com

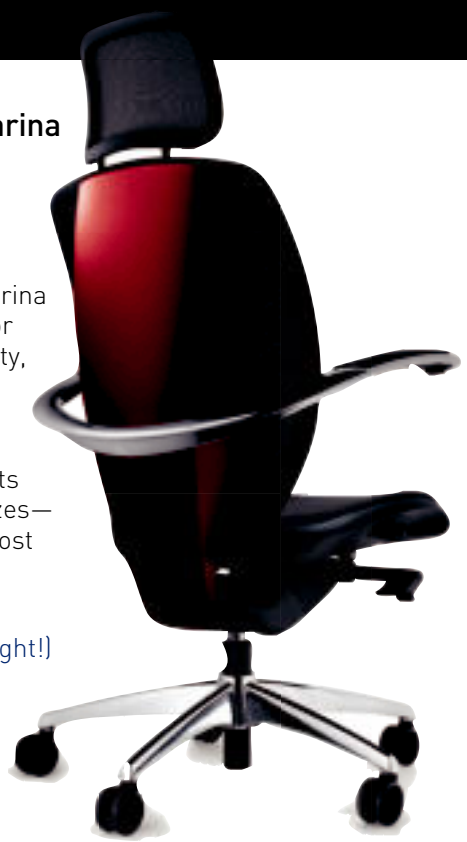

6 SONY XEL-1 OLED TV

This 11-inch display is small, but it gives you the bragging rights of owning a TV based on the long-vaunted OLED technology. Organic light-emitting diodes are reported to improve contrast ratios and offer superior color quality, including absolute black levels. Couple that with a panel that's 3mm thin, and you've got some big technology in a tiny package.

$2,500,

http:bit.ly/7G9uT6

7 ART LEBEDEV OPTIMUS MAXIMUS KEYBOARD

Every one of this board's 113 keys is embedded with a micro OLED display capable of up to 10fps video, allowing the user to reprogram each of the keys to perform different tasks for different apps.

$2,400

http://bit.ly/313k8e

8 Intelligent Design Emotion Mouse

When no luxury is too small, an artfully designed wireless laser mouse hand-crafted from grade 1 titanium and high-quality plastic resin makes sense.

$1,200

www.intelligent-design.nl

This month the Doctor tackles...

Setting Your Processor Affinity
PCI Express
Shrinking RAID Partitions

Are Default Drivers Enough?

Every time I try to join a multiplayer server in Battlefield 2, my computer freezes with no sound for nearly a minute, then resumes again. I can't seem to diagnose the problem. I have installed the newest patch (1.5). I have reformatted my hard drive and reinstalled Windows 7 Home Premium 64-bit. All the drivers were automatically installed by Windows, so I only had to install my soundcard drivers. My computer is connected to a cable modem via a Linksys wireless router. I have also manually updated my PunkBuster files. Any idea what's causing the pausing?

—rtschchand

There's your problem right there, as they say. Your issue actually highlights one of the things Windows 7 does *too* well. Because Win7's native drivers are so good (especially compared to Vista's), nearly everything works fine immediately after install, and you might think the days of manually hunting down and installing drivers are gone. If only.

Although the Windows 7 native drivers are quite solid, it's always better to manually install the latest version of your motherboard chipset and videocard drivers from the manufacturers' respective websites. You should even check on your network card drivers. If that doesn't fix it, make sure that if you're running a firewall or antivirus program, it's not blocking BF2—or worse, popping up notifications under the game window where you can't see them.

What's the Deal with PCI Express?

Are x1 PCI Express cards and x16 PCI Express slots compatible? Can I insert my x1 PCI Express soundcard into a x16 PCI Express slot? And maybe this question isn't relevant, but why do motherboard companies always put their x1 slots between two x16 slots? Even if you install one videocard it will block access to the x1 slot. Of course, you can install your videocard in a secondary x16 PCI Express slot, but how do you know it's not running at x8 speed? Moreover, if you want to install only one videocard, the motherboard manufacturer instructs you to install it in the topmost x16 slot. But is that required?

—Vitaliy

Short answer: Yes, a x1 PCI Express card will fit and is fully compatible with a x16 PCI Express slot. Why do most board companies sandwich x1 slots between x16 slots? They do that because if you use a double-width GPU, you only lose access to the x1. If you had a x16 slot next to it, you would lose access to that one. How do you know what speed the slot is running at? Most boards will have the speed printed somewhere on the PCB, but the easiest way is to consult your motherboard manual. We suspect that most motherboard makers recommend that you always install a card in the top x16 slot because those slots are almost always x16. Since the lower x16 slots on many boards actually run at x8 or even x4 data rates, installing in any slot other than the top could be problematic. If you're curious about what speed your GPU is running at, download the free GPU-Z tool from www.techpowerup.com. In the Bus Interface tab of the utility, you can see what mode your GPU is running under.

High idle temps? Replace your stock cooler with a higher-performing one.

Is My CPU Running Too Hot?

I have a custom-built gaming computer made by Magic Micro housed in a Thermaltake Soprano case. It has 4GB of RAM, a 3GHz Core 2 Duo (model unknown), a 500-watt Antec power supply, two 1TB Seagate hard drives, and a Sapphire Radeon HD 4870. Things are fun and fast here. The desktop sits right under the air-conditioning vent in my library. Using SpeedFan, I recorded the max temps to be 158 degrees in the Core 0-1 areas and 124 degrees at the CPU. Am I in trouble? Normally it runs around 98 degrees combined, but it was a hot night. I turned the AC on and it all cooled down quickly back to 98 F. Is my desktop OK or do I need to do something?

—Chris Buemi

Your desktop's probably fine, Chris. 158 F (70 C) is within normal operating temperatures for your processor (likely

a Core 2 Duo E8400, by the way), though definitely on the high end. The ambient air temperature in your house with the AC off no doubt contributed to the high temperatures.

Though the temps you're seeing aren't too worrying, they're higher than they should be. Make sure nothing is blocking your intake and exhaust fans—vacuum to remove dust, pet hair, small rodents, etc.—and make sure your CPU cooler is firmly mounted. If it is, and Magic Micro stuck with the stock cooler, you might want to consider an upgrade—as always, we recommend the $30 Cooler Master Hyper 212+ for an inexpensive, high-performance CPU cooler.

Shrinking Partitions in RAID

I read your answer regarding partition resizing with 64-bit Vista ("64-Bit Partition Resizer," December 2009). My 1TB of storage consists of two 500GB hard drives in RAID 0. I would like to shrink my partition to allow a dual-boot with 64-bit Windows 7. Should Vista's Disk Management utility be able to handle this (the menu option I get is Shrink Volume)?

—Lawrence

If you're using a hardware RAID setup (either onboard or through an add-in card), Vista will see your RAID as a single disk, therefore you ought to be able to use Disk Management to shrink the partition—in theory. However, we tested this method on a RAID 0 of two 40GB Intel X-25V SSDs on a 64-bit Windows 7 machine,

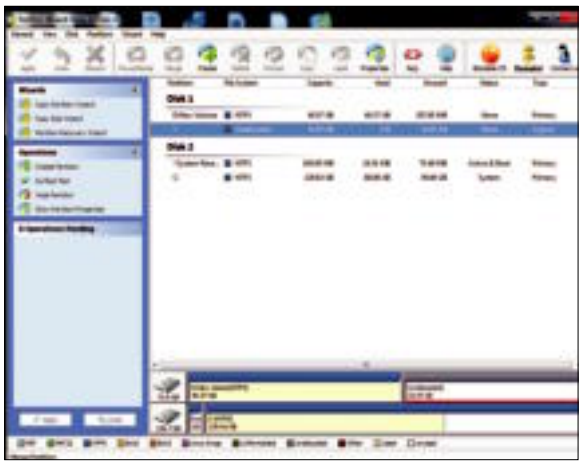

Partition Wizard 5 can resize partitions where the Windows tool fails—even with RAID. And the free version has 64-bit support, a rare feat.

and although the Shrink Volume let us set a new size for the partition, the actual shrinking procedure failed. Fortunately, the free home version of Partition Wizard 5 (www.partitionwizard.com) succeeded where Windows' built-in solution didn't, allowing us to shrink the partition on the disk and create a new one in the freshly unallocated space. Given that you're attempting to shrink a system partition on a RAID, though, you should definitely defrag the array and back up everything beforehand; some RAID controllers can be weird about shrinking partitions, especially on striped drives.

One final note: We question the utility of dual-booting Vista and Win 7. There's almost no software that runs on Vista but not Win7, and Win7 is all-around better, more stable, and more secure. Dual-booting is useful if you're running two totally different operating systems (say a Linux partition and a Windows partition), but we're not sure why you'd need to boot into Vista at all, once you have Win7.

Playback Sucks

I recently purchased a Sony Handycam HDR-CX150. It records in full 1080p HD and saves the files in .m2ts format. When I first transfer these files from the cam to the computer and play the .m2ts files on my Sony software, the video is clear, crisp, and looks the way I think full HD video should look. But .m2ts won't play on any media player on my computer except VLC Player, and then the quality is horrible—it tears and won't play right. So I am forced to convert it to a full HD H.264 file or something of the same quality. I have tried every converter program from HandBrake to AVS Video Converter. I don't care about cost or hard drive space; I just want to know if there is any way to convert my full HD .m2ts files to a format that will not lose any, or barely any, video quality. I don't care if the output file is larger than the input file! I have 12TB of storage between my computer and server. I'm using a 2.66GHz Core i7-920

SUBMIT YOUR QUESTION Are flames shooting out of the back of your rig? First, grab a fire extinguisher and douse the flames. Once the pyrotechnic display has fizzled, email the doctor at doctor@maximumpc.com for advice on how to solve your technological woes.

overclocked to 3.7GHz, a Radeon HD 4890, and 6GB of DDR3; I don't think my hardware is the problem.

—Christopher Basquit

Christopher, the opinion of the *Maximum PC* staff is that rather than try to re-encode your .m2ts files and potentially lose out on that sweet HD goodness, you should invest in a media player that handles .m2ts files with ease, like the latest versions of InterVideo WinDVD or CyberLink PowerDVD Ultra. That way you won't lose any of the quality of the original.

Incidentally, .m2ts is a container format just like .mkv; the underlying video codec can vary. And the reason VLC seizes up when you try to play HD .m2ts files is that the player lacks GPU acceleration—although the recent 1.1 preview build supports it.

Yet Another RAM Timing Question

Can you please explain to me this idea of the front-side bus and memory timings being the same speed? This doesn't make any sense to me. I have an EVGA 750i FTW board, a 2.4GHz Core 2 Quad Q6600 running at 333MHz FSB, or 1,333MHz effective, for 3GHz, two EVGA GeForce GTX 260 cards, and 8GB of OCZ Gold 1,066MHz. I had the memory set (all "unlinked") to 1,066MHz at 5-7-7-18 28, 2T @ 2.15 V. Several forums on OC'ing the Q6600 recommended an 800MHz speed, so now I have the memory set to 800, 5-6-6-16 26, 2T @ 2.1 V. I don't see any real difference in performance. Should 1066 not be faster than 800? What is the story, Doc?

—Spencer Taylor

Unfortunately, Spencer, there are very few applications that are currently able to use the memory bandwidth available on today's PCs. That's mainly the result of Intel's very efficient and very large on-die caches. These large caches were specifically designed to ameliorate bandwidth bottlenecks.

In fact, if you were to run your system in single-channel mode instead of dual-channel (and effectively gut your bandwidth), you'd be unlikely to see a difference in the vast majority of games and apps you use. Only synthetic memory tests such as SiSoft Sandra or Everest would show you that your RAM at higher clock speeds is actually giving you increased bandwidth. You should also keep in mind that it looks like the PC2-8500 OCZ Gold Edition has rated timings of 5-6-6-18 (CAS-TRCD-TRP-TRAS) at 2.1 volts when run at 1,066MHz. You were actually running relaxed timings initially. By running at 800MHz, you can probably run even tighter timings than you tried. However, the Doctor must again note that you'll probably see only very slight performance improvements, if you see any at all.

RAM Preventing Wakeup?

Doc, I have a 2-year-old Dell XPS 720. It has a 2.4GHz Intel Core 2 Quad Q6600, a GeForce 8800 GT, and runs on Windows XP. It shipped with two 1GB DIMMs. A few months ago, I decided to upgrade to four 1GB DIMMs. I bought the RAM from Crucial. The machine boots and shows the correct amount of RAM and runs great. My problem is coming back from a sleep or hibernate state. The PC's hard drive just spins and nothing happens. If I take out one DIMM, it will boot right up. I can shut it off, replace the fourth stick of RAM and it comes right up. It doesn't matter which stick I remove or what slots they are in. I've updated the video drivers and BIOS but it didn't make any difference. Any ideas?

—Jim Leiferman

While running greater than 4GB of RAM with Windows XP is known to cause hibernation issues, the Doc's uncertain about the cause of your problem. The sad fact is, you could write an entire book troubleshooting hibernation and sleep issues in Windows XP. When you add RAM, your hiberfil.sys file will grow in size to match the amount of RAM you have. Your page file may also increase in size. If you are running particularly tight on hard drive space, this could be the source of your issue. It's a little hard to follow what you did, but you didn't actually add RAM to the machine while it was hibernated, did you? Adding RAM to Windows XP while in hibernation would likely induce strange results with the OS. If you're just ready to give up on trying to figure this out, the compromise might be to run just 3GB of RAM instead of 4GB. Since 32-bit Windows XP cannot generally access more than 3.5GB of RAM, you won't be giving up too much by losing that one DIMM. The lost memory bandwidth is also unlikely to impact performance very much, either.

Persnickety Second DVD Drive

I have a problem with my second DVD drive: Sometimes it works, sometimes it doesn't. My system is a home-built Biostar TA790GXB A2+ motherboard with an AMD Athlon II X2 250 running at 3GHz, a GeForce GTS 250 videocard, 8GB of RAM, two 500GB Hitachi hard drives, and Windows 7 32-bit, with all the required updates installed. There is also an LG CH08LS10 SATA Blu-ray ray drive (assigned drive F by Windows)

Now to the problem: drive E. Sometimes it runs, sometimes it doesn't, and there is no rhyme or reason to it. Sometimes it starts fine, recognizing DVDs or CDs and burning DVDs and CDs as necessary, and runs well all day. Sometimes it starts and won't even recognize any CD or DVD (to read or to burn) I put in it—and won't all day.

Originally, I had installed an LG SATA DVD drive. When it acted up as described above, I replaced it with a BenQ DW 1640 from my old computer. It's an IDE drive. Same thing happened. No problem with the Blu-ray drive.

So, there you have it. Two different drives, using different controllers on the motherboard but malfunctioning the same way. Any ideas? Does Win7 think people shouldn't have more than one, or more than one type of, burner? Is there some jumper I should be moving or software setting I'm missing? Did I just find a Win7 bug?

—Rick Thompson

The Doctor is stumped as well, Rick. We don't know of Windows 7 having any problems supporting multiple optical drives. You've clearly done all the obvious troubleshooting by trying different drives, cables, and controllers, so faulty hardware doesn't seem to be the issue. You didn't mention whether you're running any drive-cloning or virtualization software, but if you are, we'd recommend you uninstall it. Otherwise, you might try updating your BIOS, or as a last resort, reinstalling the OS to see if that remedies the problem.

DOES WIN7 THINK PEOPLE SHOULDN'T HAVE MORE THAN ONE, OR MORE THAN ONE TYPE OF, BURNER?

HIGH-PERFORMANCE UPGRADES...
ON THE CHEAP

Upgrading your PC can be a head-spinning process. Our Lab experts help you sort through the chaos with 18 products that won't break the bank

BY THE MAXIMUM PC STAFF

The art of the PC upgrade is simultaneously an expression and a test of one's diagnostic skills, computing savvy, and fiscal sensibilities. Identify the bottleneck. Research the parts that will fix the bottleneck. Remove the bottleneck.

As always, price and performance are the pivot points. After all, you can't just toss $1,000 at your system to level it up. Well, you can, but in most cases you'd be a fool for doing so.

When the *Maximum PC* staff convened in conference room Spock to plan this story, we decided to establish some ground rules. First, we challenged ourselves to stick to our theme of a successful budget upgrade. This meant avoiding the tendency to fall back on the most expensive, best-of-breed components in each category.

Instead we forced ourselves to take a more nuanced approach. In each category, we expended considerable energy determining which product(s) owned the sweet spot—top-left on the 2x2 grid if you're graph-happy—of the price-performance ratio. Staying consistent with our real-world theme, we used real-world pricing from sites like NewEgg and Amazon. Because we're talking about upgrading an existing machine, you'll find no case or mobo recommendations here.

Without further adieu, we happily present the results of our research. Inside you'll find a bevy of product recommendations that prove you don't have to break the bank to achieve substantial gains in performance.

THESE COINS CAN BE YOURS!
Tell us about your best or worst PC upgrading moment. We're going to pick the top 50 replies and coin you. More details at www.maximumpc.com/upgrade_contest.

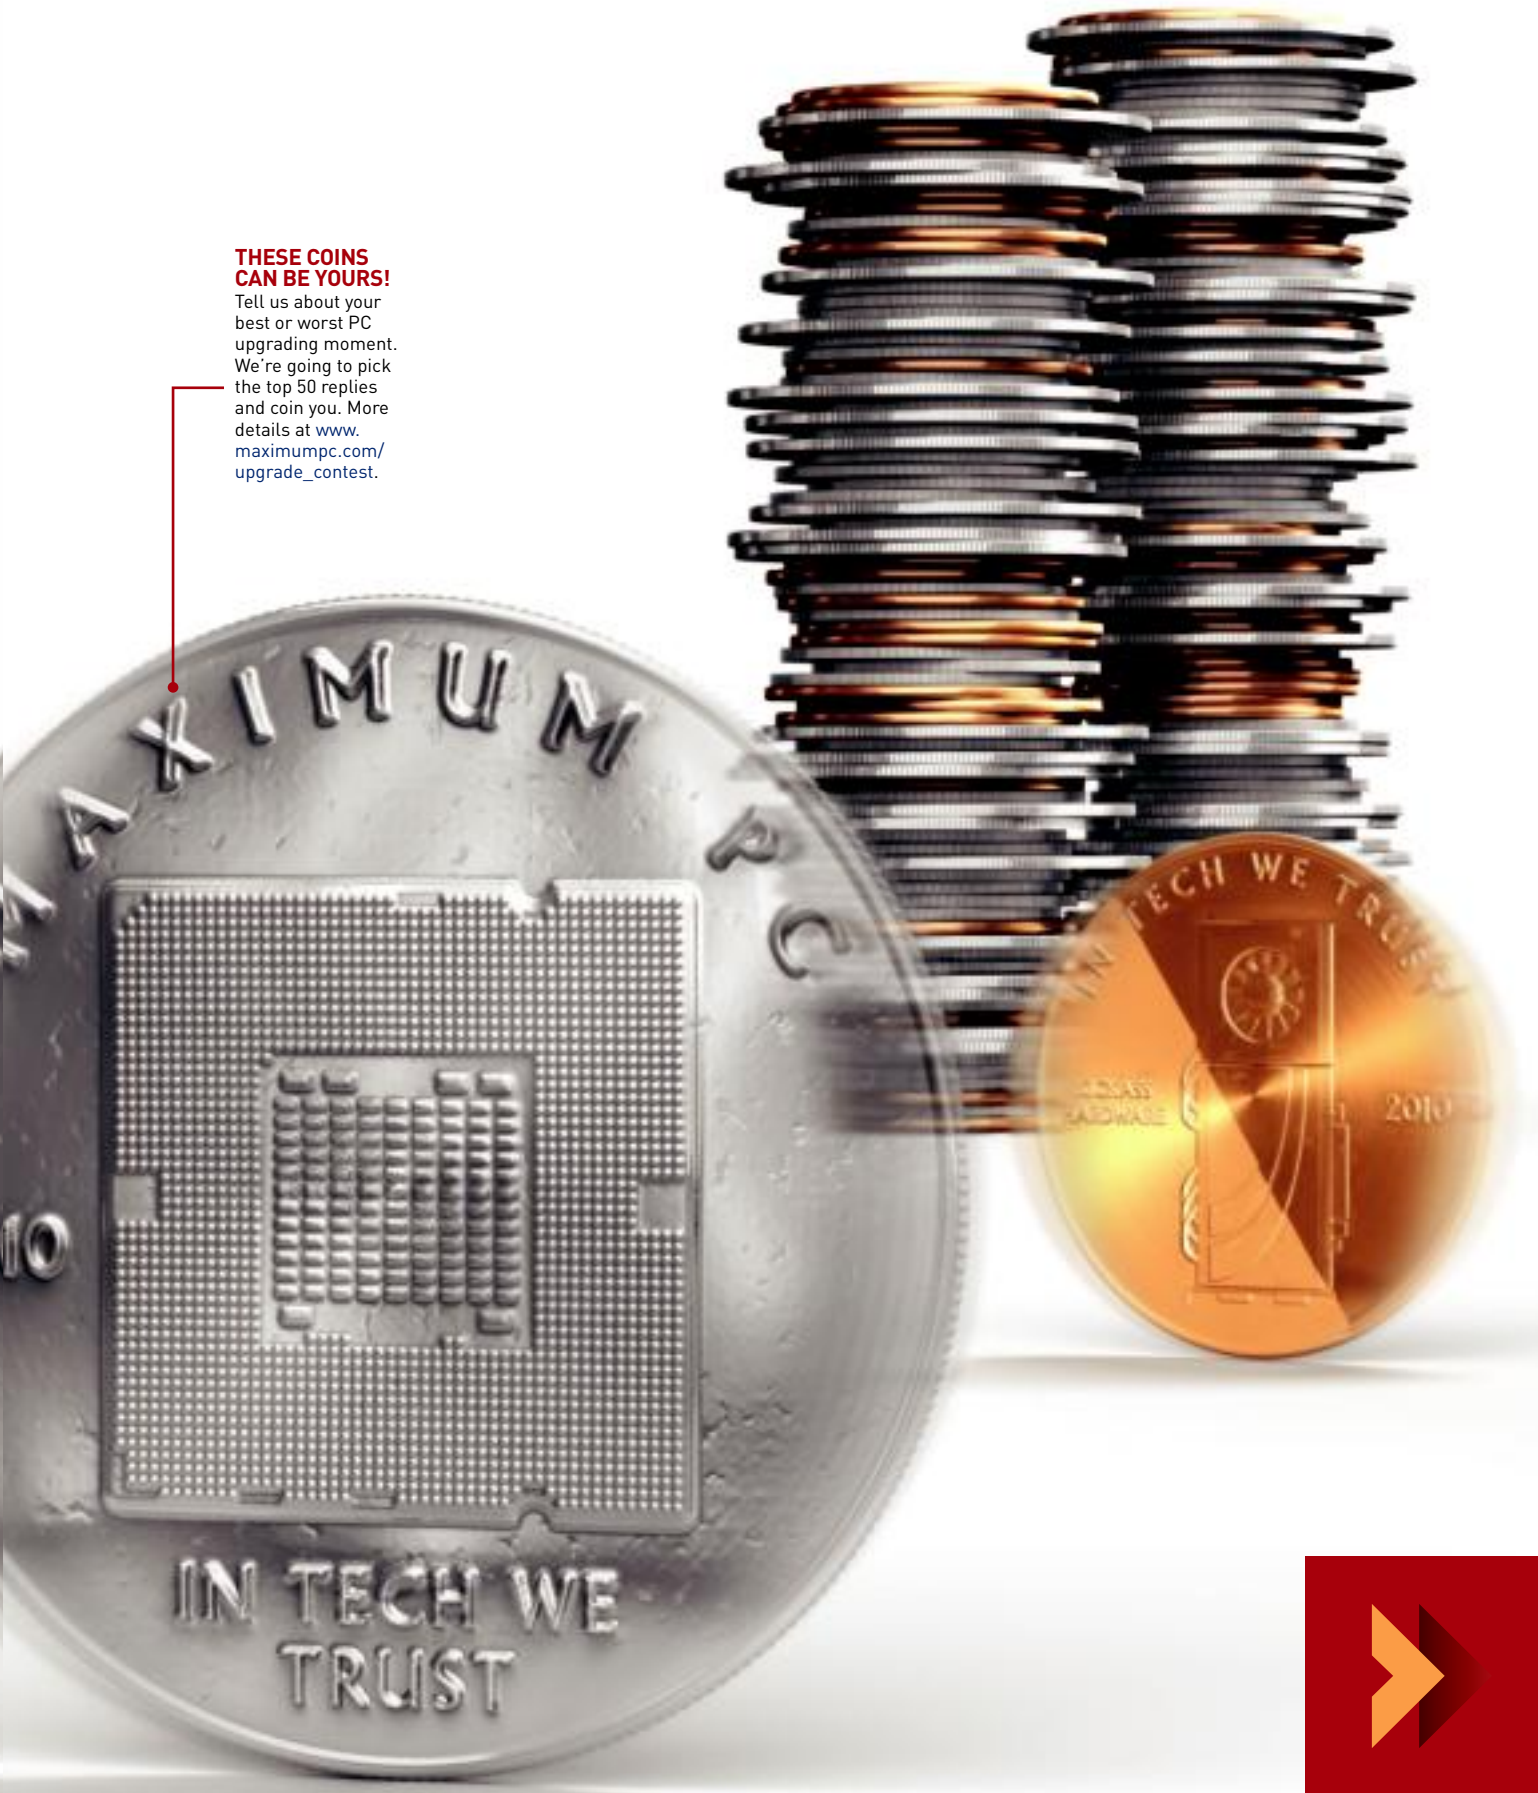

SOLID STATE DRIVE

40GB Intel X-25V

It's easy to argue that a budget SSD doesn't actually exist. That said, a $125 solid state drive can qualify as a budget upgrade in some contexts—and only some of those contexts involve recreational drugs.

Intel's X-25V solid state drive (the V stands for Value) doesn't have the fastest sustained write speeds (think 50MB/s, not 200MB/s), but its sustained read speeds top 150MB/s and its random-access writes are triple any of its peers'. This makes it perfect as an OS drive, which relies more on reads and writes than on sustained writes.

If you don't mind keeping data on an external drive or SD card, a 40GB Intel X-25V can also offer a substantial speed boost to the 5,400rpm drive on your netbook or older laptop. And if you're moving to Windows 7, the X-25V supports TRIM, which will prevent performance degradation. $125 is a lot for a hard drive, but for an SSD, it's downright reasonable given the performance bump you'll experience.

- ✔ SSD FOR $125
- ✔ TRIM SUPPORT PREVENTS DEGRADATION

Give your aging laptop a kick in the pants by replacing its hard drive with an SSD.

MECHANICAL DRIVE

Seagate Barracuda 7200.12

In the old days, the prospective hard drive buyer had to choose between high performance and high capacity. Heck, if you're planning on upgrading, you probably don't have either.

Fortunately, while solid state drives have thoroughly usurped the highest end of the performance spectrum, mechanical drives still rule the capacity roost, and they're only getting faster. To wit: the 1TB Seagate Barracuda 7200.12, which costs just $80 and offers sustained read and write speeds of over 100MB/s.

While it can't match the speeds or random-access times of WD's VelociRaptor drives or SSDs, the 1TB Barracuda is capacious enough for all your apps and data—unless you're in the habit of ripping Blu-ray discs, of course. So, if your OS drive is getting long in the tooth (or just running out of room), moving to a 1TB Barracuda 7200.12 will buy you some breathing room and a substantial speed boost.

- ✔ 1TB FOR $80 DEFINES BUDGET UPGRADE
- ✔ PERFECT SINGLE-DRIVE SOLUTION

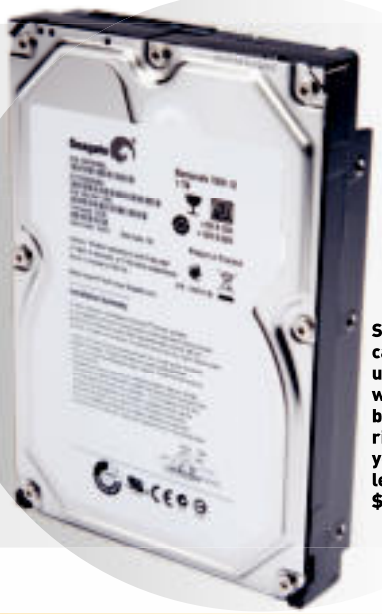

Speed and capacity unheard of when you bought your rig, now yours for less than $100.

OPTICAL DRIVE

Samsung SH-B083L Combo Drive

In our opinion, a BD burner is still too pricey an upgrade for its limited usefulness, but a BD ROM combo drive, like Samsung's SH-B083L, makes sense.

If you're currently performing DVD chores with a 16x burner, an upgrade to a higher burn-speed rating is beyond cheap (shoot, our current Best of the Best 22x Samsung SH-S223 is $20), but not all that satisfying in terms of performance gains. With DVD media stuck at 16x, higher-rated drives only exceed that limit when burning to discs of a particular brand. And even then, you're looking at a time savings of maybe a minute. Big whoop.

Instead, consider the benefits of upgrading to a BD-ROM combo drive. You can get Samsung's SH-B083L for $100. It gives you the ability to enjoy HD Blu-ray movies on your newly upgraded display, while still offering respectable 16x DVD+/-R write speeds. In our tests, the SH-B083L's performance was on par with the more expensive Plextor PX-B320SA (which received a 9/Kick Ass verdict in December 2009) in everything but DVD ripping, where the Samsung took 15:17 to copy a dual-layer disc vs. 10:47.

- ✔ AFFORDABLE, SPEEDY BLU-RAY PERFORMANCE

VIDEOCARD

ATI Radeon HD 5850

When it comes to videocards, you can count on today's $300 product being superior to the top-shelf product from two generations back. That's certainly the case with cards based on the ATI Radeon HD 5850 GPU, which not only deliver superb performance, but do so without requiring a massive power supply.

What might it be replacing? If your gaming rig is three years old and you invested in a high-end videocard, it would have been based on Nvidia's 8800 GTX, and the card alone would have set you back $600. Besides costing a fortune, that card required a massive heatsink and fan and sucked power from two 6-pin power cables in addition to what it drew from the PCI Express bus (165 watts in total). That GPU boasted amazing performance at the time, and it heralded the arrival of DirectX 10. Today, the card is performance-limited with next-gen DX10 games and it doesn't support DX11 at all.

- ✔ HELLO, DIRECTX 11 GAMES!
- ✔ PERFECT REPLACEMENT FOR THE 8800 GTX

A Radeon HD 5850 card will deliver excellent performance and should remain viable for years to come—as long as you don't upgrade to a 30-inch display. At 1920x1200 resolution with antialiasing disabled, these cards can run Crysis at 30fps. Boost AA to 4x and you'll lose just four frames per second in a game that used to bring even the highest-end GPUs to their knees. You'll fare even better with other titles: Far Cry 2, for example, can easily hit more than 60fps at 1920x1200.

Upgrading to the HD 5850 is a simple decision in other ways, too: It's 9.5 inches long, so it will fit in any case that housed an 8800 GTX, and you won't need a new power supply. Lower price, excellent frame rates, and decreased power consumption—what's not to like?

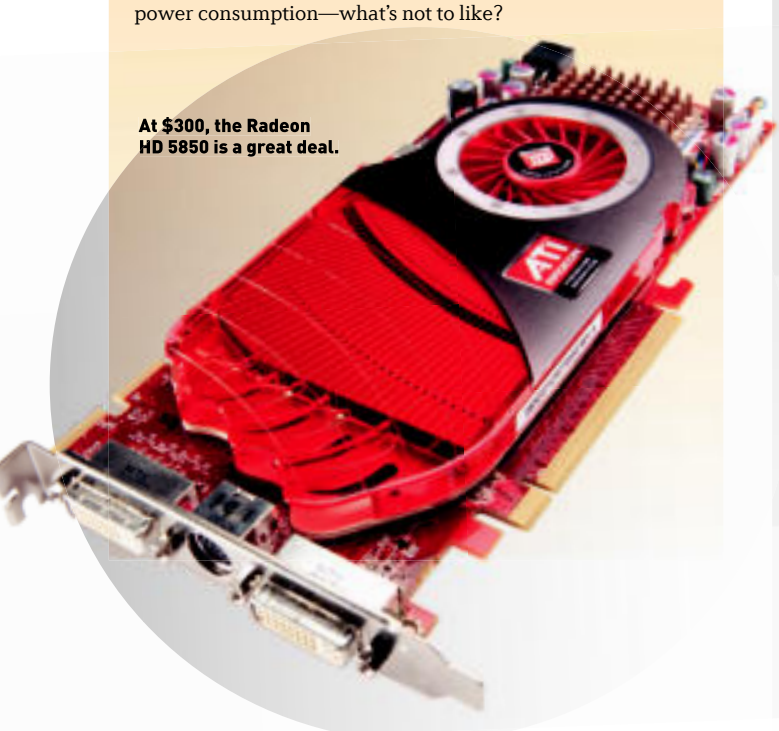

At $300, the Radeon HD 5850 is a great deal.

DISPLAY

This IPS display excelled in our Lab tests.

Viewsonic VP 2365

Twisted Nematic LCD panels blow. After running through our DisplayMate, Blu-ray, and gaming gauntlet of Lab tests, the TN displays we've reviewed retreated with their DVI cables tucked between their legs. So what's a budget upgrader to do?

If you want our advice—and you do—pick up ViewSonic's VP2365wb. It's a 23-inch IPS panel offering 8-bit color depth. It's equipped with a four-port USB hub and a height-adjustable stand that tilts, rotates, *and* pivots. And you can find it selling online for about 300 bucks.

You *will* encounter trade-offs: Although it's marketed as a "professional" monitor, its max resolution is a consumer-ish 1920x1080. It's dimmer than its pricier competitors, and it doesn't have an HDMI input. But in Lab tests, we had no problem playing games or movies, and it's a better photo-editing monitor than any TN display we've tested.

- ✔ IN-PLANE SWITCHING DISPLAY OFFERS SUPERIOR IMAGE AND VIEWING QUALITY

WI-FI ROUTER

Belkin Play Router

Belkin has been hit or miss on the router front over the past few years, but its Play router is a definite hit. Here's a concurrent dual-band 802.11n router (it runs 2.4GHz and 5GHz radios simultaneously) with a virtual guest network, a USB port that can share either a storage device or a printer over the network, and very respectable throughput and range that sells for less than $100.

The router is self-healing, too. It automatically detects and attempts to resolve network problems, and it will automatically reinitialize itself on a weekly basis (you choose the day and time—or turn off the feature if you don't like it). If that doesn't deliver enough value for you, Belkin also throws several applications into the mix. Memory Safe is a utility that runs on your client PCs and automatically backs up whichever directories you designate to an external drive attached to the router. Music Mover is an UPnP- and DLNA-compliant media server. And Daily DJ analyzes your music library and automatically creates playlists based on one of three user-designated moods: High Energy, Steady Groove, or Kick Back. We haven't used this last feature long enough to have a solid opinion about it, but it wouldn't detract from this router's value even it if was unusable.

In fact, there's just one feature we find wanting on the Play router: It has a four-port 100Mb/s switch, versus a gigabit switch.

Belkin's Play router is loaded with high-end features, including two radios and the ability to host a virtual guest network.

- ✔ SELF-HEALING ROUTER
- ✔ BUILT-IN UPNP/DLNA MEDIA SERVER

RAM

The Real Question is: 4GB or 6GB?

For budget buyers, it makes no sense to get caught up in the bandwidth wars that memory makers are waging today. The truth is, unless you use applications with particularly high bandwidth requirements, DDR3/1333 will work fine.

Ultimately, the amount of RAM you should run in a modern PC really depends on your CPU. If you are running an AMD system with dual-channel DDR3, the minimum is 4GB. Likewise, if you are rolling a dual-channel Intel system, then 4GB should be in your sights. Intel systems with tri-channel memory should run a minimum of 6GB of RAM. Anything above 6GB is gravy.

Invariably, first-time upgraders want to know whose memory to buy. Since RAM is generally a commodity component, our guideline is to stick with known brands: Corsair, Crucial, Kingston, OCZ, Patriot. No yellow-box memory, please.

The good news is that RAM prices seem to have stabilized somewhat. We found 4GB of brand-name DDR3 for $100 on the street, with 6GB of brand-name DDR3/1333 in the $160 range.

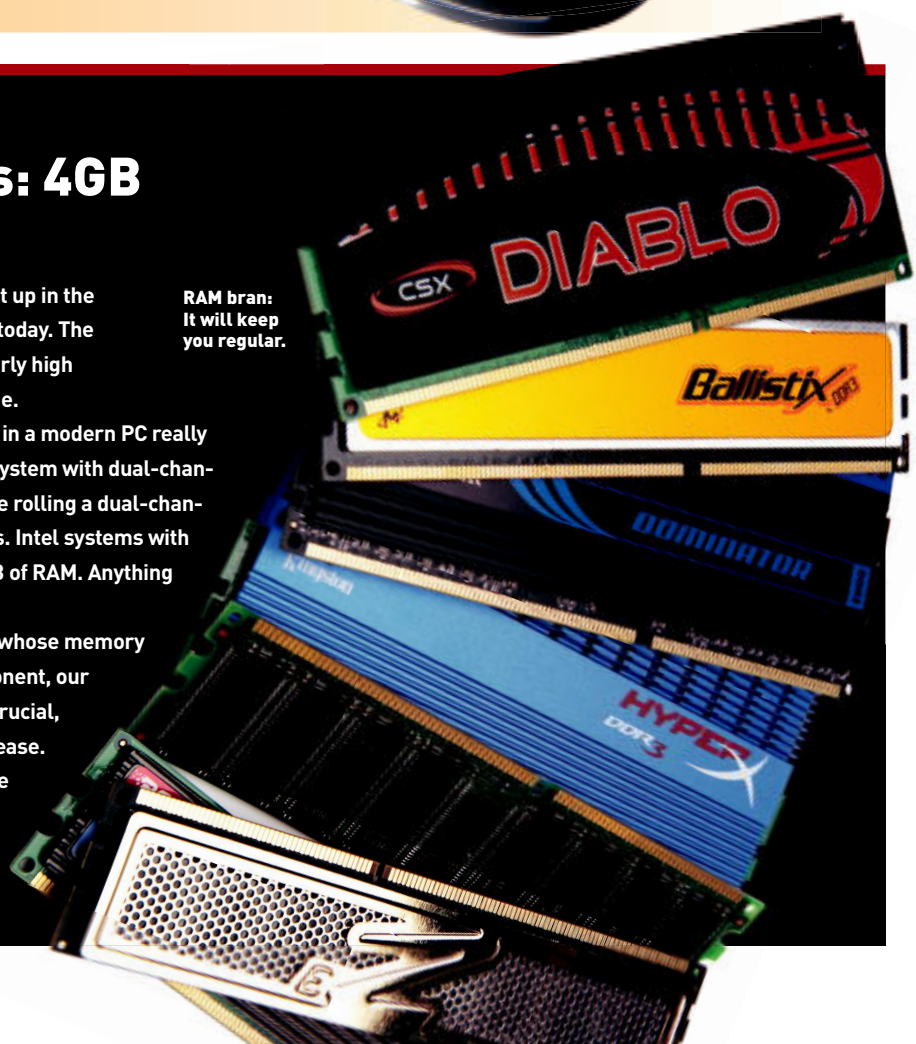

RAM bran: It will keep you regular.

MOUSE AND KEYBOARD

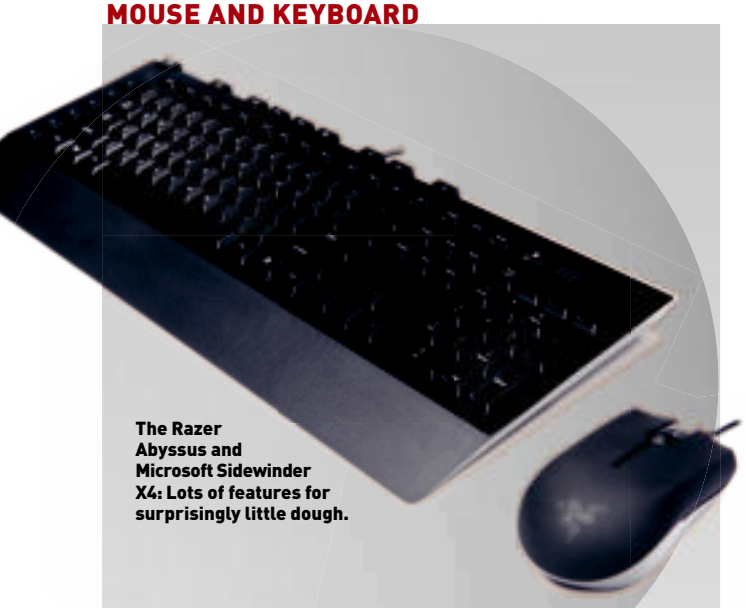

The Razer Abyssus and Microsoft Sidewinder X4: Lots of features for surprisingly little dough.

Razer Abyssus Gaming Mouse and Microsoft SideWinder X4 Keyboard

Basic mice and keyboards are commodity-priced goods, available for as little as 10 dollars. They get the job done just fine. But if you're planning to do any gaming at all, you owe it to yourself to upgrade to a gaming-grade mouse and keyboard combo. This upgrade—one of the cheapest you can make—may very well make the biggest difference.

For a budget gaming mouse, we recommend the Razer Abyssus ($30). Although it sits at the bottom of Razer's sizeable gaming lineup and lacks a few of the features we appreciate in a mouse, such as thumb buttons, the Abyssus will feel like a noticeable step up from any non-gaming mouse. With ultra-tactile buttons, a 1,000Hz polling rate, and a very-respectable 3,500dpi optical sensor, the Abyssus should be more than responsive enough for all but the most hardcore gamers.

Our keyboard recommendation is Microsoft's Sidewinder X4 ($50), which eschews some of the over-the-top bells and whistles of its more expensive X6 sibling, but retains all the features we really care about in a gaming keyboard. These include anti-ghosting (which allows many simultaneous key presses), programmable macro keys, and multiple profiles that switch when you load a game. Physically, the keyboard's a real beauty, and the extra-springy keys are a joy to use for extended gaming or typing.

- ✔ 1,000HZ POLLING RATE = RESPONSIVE GAMING
- ✔ ANTI-GHOSTING KEYBOARD

POWER SUPPLY UNIT

Corsair 750TX

Picture a raft full of PC components. It'll take seven days for the rescue boat to arrive, but only five days of food and water remains. Who gets pushed off the raft first? The GPU? The CPU? No way. They're first-class passengers. The case? The lowly keyboard? Don't kid yourself. The power supply will be the first to go. Do you know why? Because no one respects the power supply.

And sadly, that's the strategy everyone takes when they build a budget PC. We mean everyone. Hell, we've even occasionally given the PSU short shrift when push came to shove.

Fortunately, Corsair's 750TX is one component that might force something else to swim with the sharks. (Yeah, take that, mouse!) With a five-year warranty, a high power-efficiency rating, and a single 62-amp rail, this PSU will keep any budget PC running, even on those sweltering summer days when your components are broiling at 120 degrees. With a street price of $99 and SLI certification for dual GeForce GTX 470 cards, the 750TX strikes a good balance between budget and midrange. Sure, it lacks modular cables, but that just means you can't misplace the cables.

- ✔ SLI CERTIFIED FOR DUAL GEFORCE GTX 470 CARDS

No one respects the power supply. Except us. And you (hopefully).

SPEAKERS

Logitech Z523

- ✔ SURPRISINGLY RICH SOUND
- ✔ REAR-FACING DRIVERS = LARGE SWEET SPOT

Few things suck harder than cheap speakers—well, except maybe cheap TN displays. So we have to wonder how Logitech manages to sell the 2.1-channel Z523 speaker system for less than a hundred bucks. Heck, we've seen them selling online for as little as $70!

Now, we'll be the first to admit that these puppies can't compete with the likes of B&W's glorious MM-1 system, which we review on page 80, but if you're seriously considering those bad boys, you aren't reading a story about budget upgrades. The Z523 isn't appropriate for critical listening, but it can fill a small room and it has an exceedingly large sweet spot, thanks to the presence of a second driver mounted on the back of the two satellite speakers. These rear-facing drivers bounce audio waves off the wall behind them so that the sound arrives at your ears a microsecond after they've heard the front speakers.

The 40-watt amp in the subwoofer sends 9.5 watts to the two-inch full-range dome drivers in the satellites and 21 watts into a four-inch down-firing subwoofer, which is augmented by a six-inch side-facing pressure driver. There's an input for a digital media player, and separate volume controls for stereo and for bass, so you can crank the lows for gaming. Lastly, there's a headphone jack for those times when you'd prefer to rock out in private.

It's hard to believe Logitech's Z523 speakers cost just $70.

HEADPHONES

Creative Fatal1ty HS-1000 USB Headphones

- ✔ X-FI DRIVERS PRODUCE HIGH FIDELITY
- ✔ REMOVABLE MIC REDUCES GEEK FACTOR

At around $60, the Fatal1ty HS-1000 headset is hardly the cheapest on the market, but it contains several features we consider a must.

First, we like our gaming headsets to have cans big enough to surround the ear, for maximum noise isolation and comfort. The HS-1000 fits the bill here, with large, oval cans and ample foam padding. Second, we need a decent, adjustable mic—bonus points if it's removable, for when we're not playing games. The Fatal1ty gets a gold star here, as well. Third, the set should have some sort of in-line volume/mic control, to make it easy to fine-tune your settings mid-game. Creative's set has this as well.

Last and most importantly, for us to call any gaming headset an "upgrade," it has to actually sound good. That means bass deep enough to let you really feel each exploding frag grenade, and highs that let you hear the bullets whizzing past your head. This category is what makes the HS-1000 really shine as a budget headset, thanks in large part to Creative's excellent X-Fi drivers, which emulate the functions of a full-fledged X-Fi soundcard, including EAX effects and simulated surround sound.

You could go cheaper, and you could certainly go more expensive, but we think that the Creative Fatal1ty HS-1000 hits the budget sweet-spot: strong performance at an amazing price.

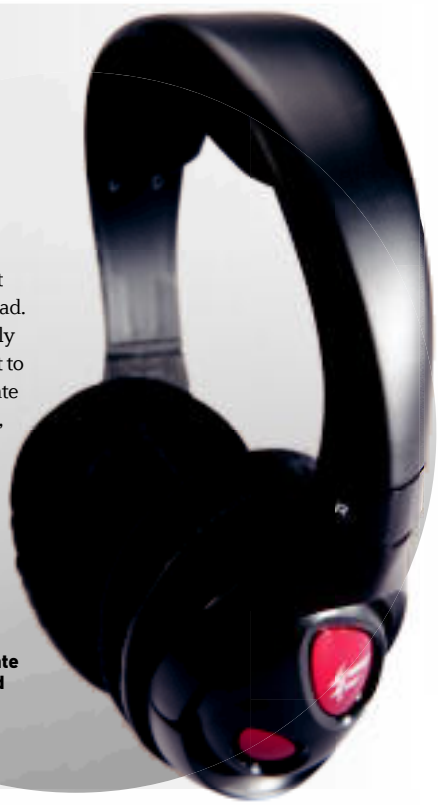

Creative's X-Fi drivers emulate the functions of a full-fledged X-Fi soundcard.

CPU Upgrades

Does it make sense to upgrade your CPU today? What's the sweet spot for price and performance? And how about Intel's LGA775? Answers below!

When it's time to pep up an old PC, the CPU is usually the first candidate that springs to mind. By leveraging microarchitecture changes, larger cache sizes, and additional cores from a new chip, you can turn that tired old dog into a prancing pony.

At least, that's always been the promise of a CPU upgrade. While we're certainly champions of the value of a fast processor, we're also the first to admit that the CPU is not always the most severe bottleneck holding you back. So before we weigh in on the intricacies of which CPU you should upgrade to and how to form a logical upgrade plan, here are a few reasons why you should think

THE GPU-CPU DIVIDE

You're familiar with the not-so-quiet war between the GPU and CPU crowds these days, right? While both factions seem quite happy to float big, fat stinking lies about the other on occasion, we generally agree that if your PC suffers from low frame rates in games, investing in a bigger GPU will usually be more impactful in delivering higher frame rates and a better gaming experience. This is especially true for those of us who play games at resolutions of 1920x1200 or higher.

This doesn't mean the CPU is worthless in gaming. You probably won't be happy with the performance you get by pairing a Radeon HD 5970 with a

ıt if
6GHz
), your
ent
'orce
r-
tead
aveat:
rting
more
eing
ıad-
s such
tal
been
ıexa-
s. Our
is that
h a

'-860 tar- ıance

high-clocked dual-core in the 3GHz or higher range, although the new baseline you should shoot for is a tri- or quad-core processor in the 2.5GHz or greater range for gaming. More on that later.

LOW RAM? FULL HARD DRIVE? BE WARNED!

Other scenarios to consider before upgrading your CPU are instances when you have abnormally low amounts of system RAM or a particularly full hard drive. If you're running 2GB or less memory on a modern OS, strongly consider moving to 4GB. And sure, that hard drive may have been fast when it was an empty 1TB vessel, but at 80 percent capacity, it will read and write much slower because the heads have to grab data from the inner portions of the platter. In this case, buying a secondary hard drive and migrating data files to the new drive will improve overall performance more than a new chip.

HOW TO PICK 'EM

OK, let's get down to it. When weighing your own chip upgrade, when is the right time to pull the trigger?

Within the same family of chips, adding clock speed will normally give you corresponding performance benefits. For example, upgrading from a 2.4GHz Core 2 Duo E6600 to a 3GHz Core 2 Duo E6850 will give you roughly a 25 percent boost in application performance. With this in mind, you should ask yourself if it's worth buying a $180 CPU for a 25 percent lift.

If you run multithreaded apps such as encoders, RAW file converters, or 3D modeling, you'll see the biggest performance gain by adding extra CPU cores. Upgrading from a dual-core 2.93GHz Core i3-530 to a 2.93GHz Core i7-870, for example, can yield a 100 percent increase in threaded app performance.

Since we're working under the parameters of an upgrade story, we'll skip major overhauls, such as moving from a Core 2 Duo E6700 on the LGA775 to a Core i7-930. Such an undertaking requires a new motherboard, new RAM, a new heatsink—and well, just about new everything. We're going to stick with upgrades that work within a certain platform.

LGA1366 PICKS

Intel's premium socket has never really been your budget buddy. Motherboards that support it demand a price premium over other Intel platforms and also require buying three sticks of RAM instead of two. Still, Intel's original 2.66GHz Core i7-920 has always been a rocking deal, and its replacement, the $300 2.8GHz Core i7-930, is our first choice as a step up.

If you're looking for a higher-end upgrade, we recommend passing on Intel's $500 3.2GHz Core i7-960. Assuming that the Internet rumors are correct, it makes more sense to wait for the 3.2GHz Core i7-970 successor, which will allow you to make the jump to a hexa-core processor. The i7-970 is expected to release as early as the end of this summer. The price is unknown, but we expect it to cost around $500–$550.

Other than this, there's not much more maneuvering room on Intel's LGA1366 platform for budget-minded shoppers. Let's move on.

LGA1156 PICKS

The LGA1156 is Intel's real budget platform. On the top end, the 2.93GHz Core i7-870 on LGA1156 is so fast that, until the hexa-core Core i7-980x was introduced, it was hard for us to justify even considering LGA1366 processors at all.

If you originally built a machine using the dual-core 2.93GHz Core i3-530 and you're finding that this $113 chip doesn't have the pep for your video encoding or content-creation needs, the sweet spot in the LGA1156 category for budget buyers is the 2.8GHz Core i7-860. At $284, this gives you a Hyper-Threaded quad-core processor. In terms of performance, upgrading to the Core i7-860 cut the time it took us to encode video from an iPhone in HandBrake by 51 percent, and RAW conversion for still photos was cut by as much as 84 percent.

If $284 is too rich for your blood, our next pick is Intel's Core i5-750. With this $200 CPU, you lose Hyper-Threading, but retain Turbo Boost. Again, compared to the Core i3-530, your HandBrake encodes would be cut by 40 percent and RAW conversions by 40 percent. Generally speaking, gaming performance with the Core i5-750 is also much improved.

So what to do if you already have a Core i5-750 or Core i7-860? The next step up is the stellar (albeit pricier)

CPU COOLER

Cooler Master Hyper 212 Plus

The CM Hyper 212 Plus quite epitomizes the "budget upgrade" concept. It requires minimal investment ($30 and half an hour) but can yield big returns for nearly any system. Skeptical? So were we. The Hyper 212 Plus is a CPU cooler with direct-contact heat pipes, which give it excellent performance for its size. It's one of the best air coolers we've ever tested: On our test bed, it cooled a stock-clocked Q6700 at 100 percent CPU utilization down to just 43 C—an 18 C difference from Intel's stock heatsink.

But why bother upgrading your CPU cooler at all? For the clocks, of course. Overclocking your CPU is the cheapest way to squeeze more performance out of your rig, but overclocked chips put out more heat. By getting a better cooler, you can sustain higher overclocks. Given that the Hyper 212 Plus can install on virtually any socket, performs better than any other air cooler we've tested, and only costs $30, it's one of the best things you can do for your PC.

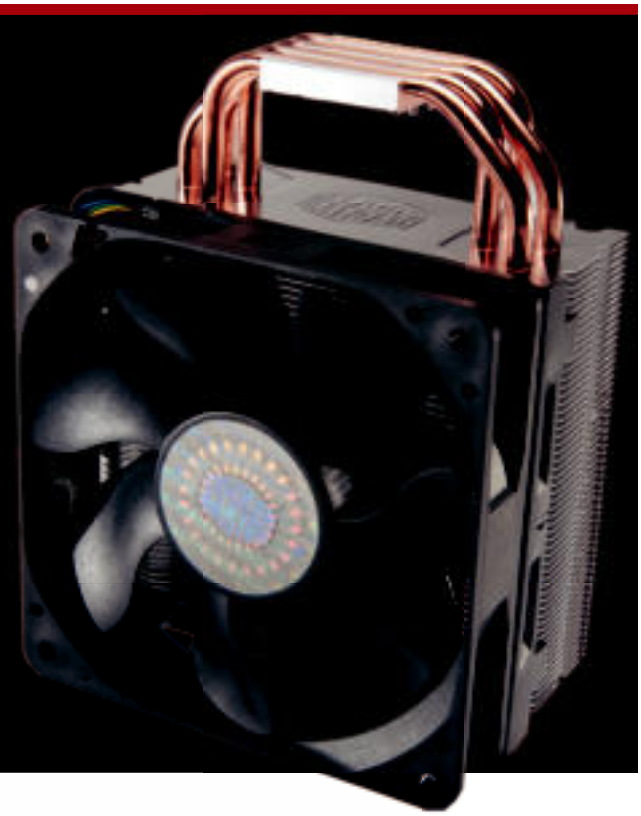

2.93GHz Core i7-870 chip. This CPU is slower than Intel's $1,000 Core i7-975 Extreme Edition but not by as much as you'd expect from a chip costing half the price. One thing to keep in mind: Internet chatter indicates a potential replacement for the Core i7-870 in coming months. Dubbed the Core i7-880, this new chip is said to clock in at 3.06GHz with Turbo Boost pushing it as high as 3.73GHz. Expected to sell at the same $550 price point as the Core i7-870, it's probably worth waiting a few months for.

OUR PICKS

INTEL

- 2.8GHz Core i7-930
- 2.8GHz Core i7-860
- 2.8GHz Core 2 Quad Q9500

AMD

- 2.8GHz Athlon II 630
- 3.2GHz Phenom II X61090T

AM2+/AM3 PICKS

The AM2+/AM3 platform is a glorious playground for upgrading. Because this platform allows anything from power-conserving dual-core CPUs to six-core chips on the same motherboard, it's difficult to pinpoint a single baseline CPU that most upgrades will be coming from. So instead, we'll give you a range of top picks.

If you're upgrading from dual-core and living on a tight budget, the quad-core 2.8GHz Athlon II 630 is our choice. At $99, it's a steal and leaves its dual-core siblings in the dust.

The real steal, however, are AMD's new six-core processors. The hexa-core 3.2GHz Phenom II X6 1090T sells for less than $300 and the 2.8GHz Phenom II X6 1055T is going for $200. At those prices, this is a no-brainer upgrade, particularly given that with Intel, you're going to have to fork over $1,000 to get six cores.

Are AMD's hexa-cores capable of slaying their equivalently priced Intel counterparts? Yes and no. The Phenom II X6 1090T aced Intel's $562 Core i7-870 in some of our Lab tests involving heavily multithreaded apps. However, in most other applications, the top-end, $300 hexa-core AMD chip had a tough time beating even the $200 Core i5-750. Ultimately, since you can't buy the Intel chip without buying a new board, CPU cooler, and practically re-building your system from scratch, the comparison is moot. The takeaway is this: At $300, the Phenom II X6 1090T is a great upgrade for this platform.

LGA775 PICKS

We've long considered Intel's LGA775 platform to be a non-starter. Intel hasn't introduced anything new for it in a year. The chipmaker also hasn't cut prices on LGA775 chips enough to keep them competitive with AMD or even its own Core i3 parts. In many ways, these CPUs are simply not competitive with anything on the market today. Add the mish-mash of incompatible chipsets and motherboards to the mix and it's clearly time to bail on LGA775.

Despite all this, the LGA775 continues to make up the majority of new desktop PC sales thanks mostly to bottom-feeder boxes. So what do you do if you happen to have an LGA775 box that you're itching to upgrade? First, make sure you can upgrade it at all. If you have an old nForce 680i machine, for example, you probably cannot run a 45nm Core 2 Quad chip. (Thank you, Nvidia and Intel.)

Most of the lift you're going to get on the LGA775 will come from moving to quad-core. There, you'll get the best bang for your buck from the 2.83GHz Core 2 Quad Q9500. At $183, it resides squarely within the budget-minded price-performance sweet spot. We don't think the 3GHz Core 2 Quad Q9650 makes much sense at $316, but when you consider that the 3GHz Core 2 Extreme Q9650 cost $1,000 a few years ago, you might want to do it just for bragging rights. But you have to ask yourself: Does it make sense to buy a new chip for a platform that's a dead man walking when a new platform and a Core i3 or Phenom II will deliver a lot more power for your money?

CPU UPGRADE CHECKLIST

What To Do Before Swapping Your Proc

Don't be that guy. You know, the rookie who commits the cardinal sin of buying a non-returnable CPU just because it fits the same socket as his motherboard. That's something that will end in tears. Follow this quick list before you buy any upgrade CPU.

1 RTFM Yes, read the frakking manual and your motherboard maker's website to find out which CPUs will work on your board. If it's not listed, there's a good chance it just won't work no matter how much you wish that it would.

2 Update the BIOS OK, you've determined that the shiny new CPU will work on your board. Before you install the new chip and boot though, make sure you update the BIOS. Otherwise, you'll have to put the old chip in just to update the BIOS to POST your new chip.

3 Assess Your Cooling If you just went from a dual-core CPU to a six-core CPU, that $13 heatsink might not get the job done anymore. Think ahead. Be prepared to meet your increased cooling needs.

THE POWER OF Six

AMD's new Thuban hexa-core CPUs come out swinging with prices that belie their size

BY GORDON MAH UNG

If we've learned anything from years of watching action movies, it's that you never, ever count out the underdog. Such is the case with perennial underdog AMD.

Bloodied, beaten, and bruised by months and months of Intel chips that outpaced its parts, AMD isn't giving up. Instead, it's hitting back with its own hexa-core CPUs and doing everything just short of yelling yippie kai-yay!

And now for the shocker: These hexa-core CPUs are affordable. Hell, one of the parts is practically budget-priced. Intel's high-flying hexa-core Core i7-980X is $1,000. Contrast that with AMD's new 3.2GHz Phenom II X6 1090T at $295. Want more? The 2.8GHz Phenom II X6 1055T costs $200. Yes, $200 for a hexa-core processor. So, yippie kai-yay motherfrakker, indeed!

Want even more good news? AMD's new chip will be backward compatible with the vast sea of AM3, and even older AM2+, motherboards out there. We're quite glad to hear this, because at one point the company told us it planned to jettison DDR2 support, which would have cut off the AM2+ folks. Fortunately, the company changed its mind and both new chips include DDR3 and DDR2 support.

Just like with those Hollywood action movies, this story wouldn't be complete without an element of suspense: Are AMD's Phenom II X6 processors capable of whipping Intel's similarly priced quad-cores, or even its $1,000 wonder, the Core i7-980X? To find out, you're going to have to read on.

AMD'S TOP PROCS COMPARED

	Phenom II X6 1090T	Phenom II X6 1055T	Phenom II X4 965 Black Edition	Athlon II X4 630
CODE-NAME	Thuban	Thuban	Deneb	Propus
CLOCK SPEED	3.2GHz	2.8GHz	3.4GHz	2.8GHz
TURBO SPEED	3.6GHz	3.3GHz	N/A	N/A
CORES	6	6	4	4
L2 CACHE	3MB	3MB	2MB	2MB
L3 CACHE	6MB	6MB	6MB	N/A
SOCKET	AM3/AM2+	AM3/AM2+	AM3/AM2+	AM3/AM2+
TDP	125 watts	125 watts	125 watts	95 watts
PROCESS	45nm	45nm	45nm	45mm
DIE SIZE	346mm^2	346mm^2	258mm^2	169mm^2
TRANSISTOR COUNT	904 million*	904 million*	758 million	300 million
PRICE	$295	$200	$185	$100

**Transistor count not released. Figure is derived from AMD's hexa-core Opteron part.*


AMD Phenom™ II
HDT90ZFBK6DGR
ACBBE CB 1010BPMW
9E21533C00260
© 2009 AMD
D
DIFFUSED IN GERMANY
MADE IN MALAYSIA

UNDER THE HOOD OF AMD'S PHENOM II X6

In many ways, AMD's new Phenom II X6 isn't all that different from the Phenom II X4 processors. Both chips are built on a 45nm process, have the same 125-watt TDP rating, and feature the same microarchitecture. Of course, the Phenom II X6 is far larger than a Phenom II X4. A typical Phenom II X4 is 258mm^2. The Phenom II X6 is about 40 percent larger at 346mm^2. Oddly, AMD wouldn't disclose the transistor count of the Phenom II X6, but we'd guess it's the same or pretty close to the hexa-core Opteron 2435 at 904 million. On the surface, it would appear that AMD just took a Phenom II X4 and glued on two more cores. It's not quite that simple, though. The L2 cache of the new chip remains at 512KB per core. The 6MB of L3 cache is also the same as the older quad-core. That's actually a reduction, since the 6MB of L3 is shared among six execution cores instead of just four. Whether that plays into the performance of the chip isn't yet clear. But AMD has introduced some other changes to boost performance.

TURBO CORE OVERCLOCKS THREE OF THE CORES IN THE CPU WHEN THREADING LOADS ARE LIGHT

TURBO BOOST, MEET TURBO CORE

AMD's top-end Phenom II X6 1090T has a lower clock than the Phenom II X4 965 BE, but AMD makes up for that, and addresses the lack of applications and games optimized for six cores, by introducing a new Turbo Core mode. Turbo Core overclocks three of the cores in the CPU when threading loads are light. The 1090T will Turbo Core from 3.2GHz to 3.6GHz, and the 1055T will Turbo Core from 2.8GHz to 3.3GHz. The mode is transparent to the OS and works without the need for drivers. It's also not as discriminatory as Intel's similar Turbo Boost, where apps that hit only one core get more of a boost than apps that hit two or three. AMD said its tests show that the biggest benefits come from the three-core increase in games and apps that have not been optimized for four or more quads. Folks who want to mess with overclocking limits and ratios for Turbo Core can do so using AMD's OverDrive utility.

KEEPING THINGS COMPATIBLE

AMD's strength has been its ability to make new chips work in older motherboards. In the three years that Socket AM2+ has been out, Intel has retired Socket 775 and introduced two new sockets that are incompatible with each other. That AMD can get its newest Phenom II X6 to work in older AM2+ boards (and probably even a few AM2 boards, too) is a testament to good planning. The only limiting factor for upgraders with the new Phenom II X6 is likely a board's thermals: The mobo must support 125-watt processors for the user to expect long-term reliability. BIOS updates will also be required, but AMD says that at launch, no fewer than 160 boards will have BIOS updates available to support the Phenom II X6.

MOBO MATTERS

Make Way for the 890FX Chipset

With the new chip, AMD is also introducing a new chipset series: the 8 series. The top chipset from AMD is the 890FX, which will replace the 7 Intel counter 2.0 lanes the supports. On problems we into with the is the lack of PCI-E lanes t from USB 3.0 6 drives. Witl motherboard ing two GPUs x16 PCI-E slo handcuff the of USB 3.0 and SATA 6 devices. Not all boards are affected, but most are. With 42 PCI-E 2.0 lanes available in the 890FX chipset, ther devices without being constrained. Compare that to Intel's P55 chipset the CPU and s in the P55 B850 south for six devices. g from the .0. Most board s are instead integrating NEC control- lers to add the feature.

So, Is It Fast?

We match the Phenom II X6 up against Intel's best and brightest

As always, this is where the rubber meets the road, and to get a good feel for where Phenom II X6 1090T stands, we compared it to a spate of chips, including the $282 2.8GHz Core i7-860, the $200 2.66GHz Core i7-750, the $195 3.4GHz Phenom II X4 965 Black Edition, and, of course, the $1,000 3.33GHz Core i7-980X.

We used identical GeForce GTX 280 cards in all of our test platforms, along with the same graphics drivers. For the Athlon testing, we used the new MSI 890-FXA GD70 motherboard. A Gigabyte GA-P55A-UD6 powered the LGA1156 parts, and an Asus P6X58D Premium handled the chore for the LGA1366 procs. All of the dual-channel boards used 4GB of DDR3/1333 RAM, while the triple-channel boards had 6GB of DDR3/1333. Windows 7 Professional and matched 150GB Western Digital Raptor 150 drives were used in all the platforms.

For benchmarks, we ran more than two dozen tests to find the strengths and weaknesses of AMD's new chips. The results are interesting, to say the least.

VIDEO ENCODING AND VIDEO EDITING

In general, more cores mean better performance with multithreaded encoders and nonlinear editors, but there's more to the story than cores. Phenom II X6 chips continue to feature a three-issue execution core vs. the Core i3/5/7 processors' four-issue execution core. Against its natural competitor, the similarly priced 2.8GHz Core i7-860, the 1090T is at a big disadvantage, with encoding times in both Premiere Pro CS3 and Sony Vegas Pro 9 taking about 22 percent longer. HandBrake saw the 1090T do a little better, but AMD's hexa-core was still about 14 percent slower than Intel's quad. The closest the 1090T came to that chip was in our MainConcept Reference encoding challenge. The 1090T does a lot better against the Core i5-750, which doesn't have the advantage of HyperThreading, but is also $100 cheaper.

3D MODELING

This is where the Phenom X6 1090T shines the brightest. The 1090T aces both the Core i5-750 and Core i7-860 in Cinebench 10 and 11.5, and POV Ray 3.7. It even manages to outpace the $562 Core i7-870 in POV Ray and Cinebench 11.5. For cheap 3D modeling, Phenom II X6 is the go-to chip.

PHOTO EDITING

We saw hit-or-miss results here. Against the Core i7-860, the 1090T was faster in our Adobe Lightroom file-conversion test by a healthy 10 percent, but it got blitzed in our ProShow Producer slide show–creation test by 17 percent. Photoshop CS3 also saw the 1090T trailing by about 5.4 percent. In the heavily multithreaded Bibble 5.02, we saw both chips on equal terms, with the 1090T trailing just slightly by 2 percent. As you'd expect, the 1090T did much better against the Core i5-750 and much worse against the Core i7-870. We guess, as they say, you gotta punch your weight.

MEMORY BANDWIDTH

The 1090T actually has the highest memory bandwidth of the AMD CPUs, but all

THURBAN UP CLOSE

Inside the Phenom II X6

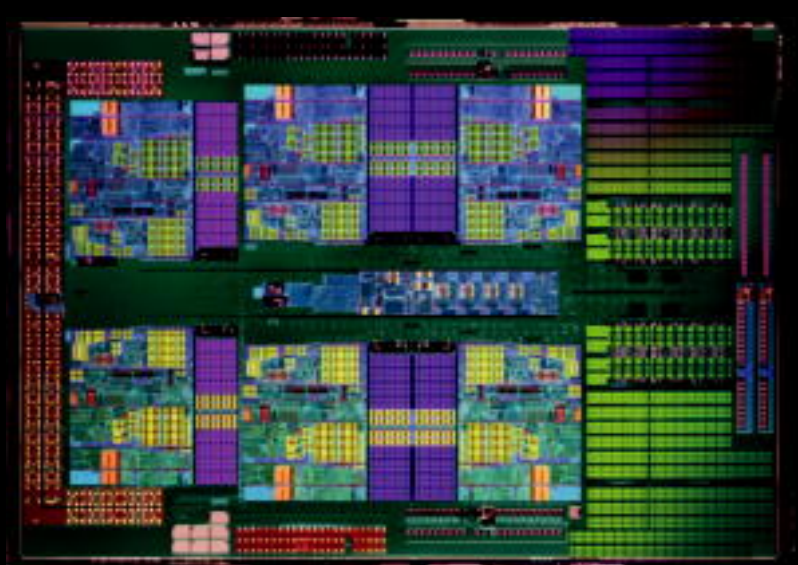

AMD's new Thuban core is essentially the Istanbul core used in the Opteron 2435. Packing roughly 904 million transistors, the monolithic core has controllers for both DDR2 and DDR3 embedded in it. While the Opteron that it's derived from has up to four HyperTransport links for multiprocessor configurations, Thuban uses just one HT link to connect to the chipset.

three Lynnfield parts and the two LGA1366 CPUs leave the 1090T behind. One thing to keep in mind: Memory bandwidth is apparently not the end-all, be-all as the Core i7-980X has only mediocre bandwidth yet is still the faster processor here.

GAMING

The area where the 1090T was most disappointing was gaming. Against the Core i7-860, we saw the 1090T trail by double digits in Far Cry 2, World In Conflict, Resident Evil 5, and the Valve Particle Test. The best that the 1090T did was in Dirt 2, where it beat the Core i7-860 by 28 percent. However, we suspect that our Core i7-860 score is the result of an error on our part as it doesn't line up with how the Core i5-750 and Core i7-870 performed. Also, keep in mind that we run our tests at very low resolutions. At normal gaming resolutions of 1920x1200, the performance really shifts to GPU-land, and you would find the margins closing up to the point where the CPU simply doesn't matter that much to the average gamer. The gaming margins would really only apply on a multi-GPU rig.

THE VERDICT

We didn't promise you an action-hero ending and you're not getting one. In the performance story, it's a mixed message. For the most part, Intel's Hyper-Threaded quad-cores still have the edge. However, Phenom II X6 gets far closer than AMD has been of late. AMD fans may be a little disappointed that the X6 doesn't spank the equivalent Intel parts, but there's still a lot of good here. The Phenom II X6 has a great price, it's probably compatible with the board you have in your machine now, and it gives you six cores.

It may not be the happy ending some were looking for, but think of it like the end of *Empire Strikes Back*. Sure, Luke got his hand cut off, Vader is his pops, and Han got frozen in carbonite, but maybe, just maybe, better things are coming on the horizon.

BENCHMARKS

	3.2GHZ PHENOM X6 1090T	3.4GHZ PHENOM II X4 965 BE	2.66GHZ CORE I5-750	2.8GHZ CORE I7-860	2.93GHZ CORE I7-870	3.33GHZ CORE I7-975 EXTREME EDITION	3.33GHZ CORE I7-980X
Premiere Pro CS3 (sec)	749	780	615	**581**	539	504	453
Sony Vegas Pro 9.0c (sec)	5,010	5,990	4,899	**3,863**	3,531	3,244	2,675
HandBrake DVD to iPhone (sec)	1,580	2,348	1,702	**1,360**	1,247	1,170	941
MainConcept 1.6 (sec)	2,816	3,286	3,092	**2,735**	2,486	2,308	1,827
Cinebench 10 64-bit	**17,892**	14,282	14,455	17,516	19,197	20,147	27,479
Cinebench 11.5 64-bit	**5.67**	4.00	3.83	5.15	5.54	5.99	8.92
POV Ray 3.7	**4,656.5**	3,313	2,810	3,883	4,497	4,236	6,557
Photoshop CS3 (sec)	130	133	118	**123**	100	91	89
Adobe Lightroom 2.6 (sec)	**426**	548	603	469	422	418	419
ProShow Producer 4 (sec)	1,669	1,912	1,425	**1,382**	1,290	1,208	1,092
Bibble 5.02 (sec)	145	202	186	**142**	122	120	97.2
PCMark Vantage 64-bit Overall	7,481	6,811	8,504	**8,903**	9,120	9,260	10,470
Fritz Chess Benchmark (KiloNodes/s)	**11,219**	8,196	8,407	10,997	11,995	12,738	12,733
Valve Map Compilation (sec)	132	126	**110**	116	106	100	99
Everest Ultimate MEM Copy (MB/s)	11,043	10,083	**15,445**	15,372	14,693	17,712	13,086
Everest Ultimate MEM Latency (ns)	51.6	55.2	54.3	**49.5**	52.5	59.8	61.3
SiSoft Sandra RAM Bandwidth (GB/s)	13	12	**17**	**17**	17	23	20
3DMark Vantage CPU	44,587	41,796	44,594	**46,064**	48,816	51,321	62,893
Valve Particle test (fps)	120	96	111	**148**	159	174	259
Resident Evil 5 low-res (fps)	100.3	91.6	110.3	**115.9**	126.6	130.7	134.1
World in Conflict low-res (fps)	162	179	**256**	253	253	317	358
Dirt 2 low-res (fps)	121	**174.8**	155	94	153.3	157	155.7
Far Cry 2 low-res (fps)	99	101.2	146.53	**150.2**	153.3	158.2	158.6

Best scores are bolded. The Core i7-870, Core i7-975 Extreme Edition, and Core i7-980X are shown for comparison only and are not weighted in determining the winner.


DELL

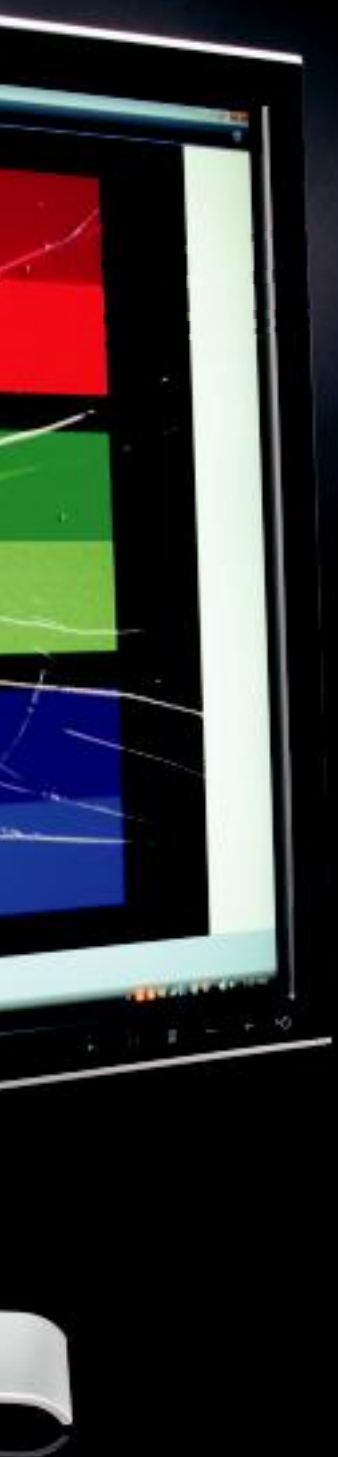

DISPLAY MYTHS: Shattered

Take everything you think you know about displays and throw it out the window. It's time for a clinic on what display specs really mean—brace yourself for the alarming truth

BY DR. RAYMOND SONEIRA
CREATOR OF THE DISPLAYMATE TESTING SUITE

Vision is our most amazing and complex sense, so it's no surprise that display technology is so amazing and complicated. It's also no surprise that most consumers don't have a good understanding of how displays function, or the best way to select them, buy them, use them, and adjust them.

Not only are displays getting more complicated and harder to understand, but the competition between manufacturers has gotten so brutal that marketing gimmicks—ploys that exploit the average consumer's technical ignorance—are playing an increasing role in driving sales. The goal of this article is to point out and explain some of the most important myths, misconceptions, and misunderstandings about display technology. Much of what you're going to read is like the classic tale of *The Emperor's New Clothes*. What you've been told about the latest and greatest thing really isn't there, or better, or meaningful, or even visible.

In the following pages, I'm going to discuss user controls, contrast ratios, pixel response time, and color gamut. These topics comprise just a portion of what a savvy consumer needs to know, so we'll be addressing other confusing display topics in future issues of the magazine and on MaximumPC.com. But for now, let's just start our journey with what should be the best question to ask before buying a new display: "What are the most important manufacturer specs to compare?" Unfortunately, the answer is none, because they're all exaggerated marketing specs rather than objective scientific specs. The only specs that are useful and meaningful are those in reviews that evaluate every display with the same consistent methodology—like the reviews in *Maximum PC*.

Confusing Users with User Controls

One reason why most consumers don't understand their monitors and TVs is because some of the most important user controls have misleading and technically incorrect names. No wonder folks can't figure out how to adjust them. In fact, they misadjust them, and then usually just leave them misadjusted permanently. Here are some highlights—well, lowlights really—of inane user-control engineering.

WHEN BRIGHTNESS ISN'T

On mobile displays with only a single user control, the control labeled "brightness" does in fact actually control the brightness of the image on the screen by increasing or decreasing the backlight intensity. However, on most monitors and TVs, the control labeled "brightness" does not control the brightness. It actually controls the signal-level setting for black on the display, which indirectly has a minor effect on brightness.

CONTRAST? NOT SO MUCH

The control labeled "contrast" has absolutely no effect on image contrast. It actually controls the brightness of the image, by increasing or decreasing the amplitude of the video signal. Monitors and TVs really should have a true contrast control, but the closest you'll find on some HDTVs is an obscure control labeled "gamma," and I have yet to see one that works properly. For more information on gamma, see my article on color and gray-scale accuracy here: www.displaymate.com/ShootOut_Part_2.htm.

CONTROLS OF A BYGONE ERA

Even more shocking, today's digital monitors and HDTVs still have the same basic user controls that were found in the original analog NTSC color TVs from 1953: brightness, contrast, tint, and sharpness. These controls only made sense for analog signals on the old NTSC television system. Brightness controlled the CRT direct-current bias, contrast controlled the video amplifier gain, tint controlled the phase of the color subcarrier, and sharpness performed analog high-frequency peaking to compensate for the limited video bandwidth of the old vacuum tube amplifiers. Today, none of these controls are necessary for digital signals.

Brightness and contrast controls shouldn't be there because, for digital video, the black level is fixed at level 16, reference white at 235, and peak white at 255. Similarly, tint and phase have no real meaning for digital signals. Finally, the sharpness control isn't appropriate for digital displays because in a digital image there's no transmission degradation—the image is received exactly as it appeared at the source. Sharpening the image involves digitally processing the pixels, which leads to artifacts and noise unless it's done at resolutions much higher than the final displayed resolution, which, of course, isn't the case inside your monitor or HDTV.

CONTROLS THAT DO WORSE THAN NOTHING

Most monitor and HDTV user-menu options are actually unnecessary features added for marketing purposes—gimmicks to suggest the display has unique features that other models lack. Even worse, most of these options actually *decrease* image and picture quality.

In many cases, it's not even clear what these sham controls really do. The documentation seldom explains them, and I even know engineers from high-level manufacturers who don't know what the controls do, either. When I test TVs, I spend an inordinate amount of time using test patterns to figure out what the options and selections really do, and in most cases, turning off the fancy options leads to the best picture quality and accuracy.

The following is a list of useless (or near-useless) menu options and selections from three HDTVs sold by major brands: Black Corrector, Advanced CE, Clear White, Color Space, Live Color, DRC Mode, DRC Palette, Dynamic Contrast, xvYCC, Color Matrix, RGB Dynamic Range, Black Level, Gamma, White Balance, HDMI Black Level, Fresh Contrast, Fresh Color, Flesh Tone, Eye Care, Digital NR, DNIe, Detail Enhancer, Edge Enhancer, Real Cinema, Cine Motion, Film Mode, Blue Only Mode.

Some of the terms sound impressive, but almost all of this is unnecessary puffery and jargon that confuses not only consumers but the pros, as well.

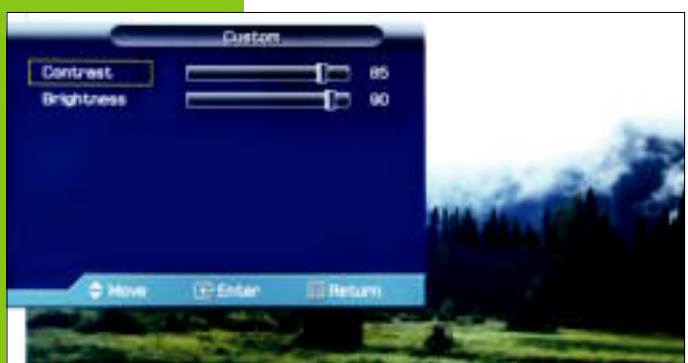

TOP: Rotary controls for a mid-century analog CRT. Those Contrast and Brightness controls are legit.

BOTTOM: Digital on-screen controls for a Samsung Syncmaster 242MP—that really have no business being labeled Contrast and Brightness.

Contrast Ratio, Ad Absurdum

Both manufacturers and consumers are obsessed with contrast ratios. Because many people choose the model with the highest number, manufacturers have developed new contrast ratio specs to win at this game.

It's a sordid business, but deserves exposure, so let's jump in.

WHEN CONTRAST RATIO ACTUALLY MATTERS

A careful, objective measurement of contrast ratio can be very revealing. After the display is accurately calibrated for optimum picture quality, the contrast ratio is determined by dividing the brightness of peak white by the brightness of black. In principle, the greater the ratio the better. Just be aware that contrast ratio is important only for low-ambient-light viewing, which is where black brightness values matter most. (In high-ambient-light settings, reflections off the screen abound, and they're all brighter than the display's own internal black.)

Even more to the point, a high contrast ratio really only matters when there is significant dark picture content, like you see in some movies. It's much less relevant for most TV shows because the picture seldom has much very dark content, and the image never really dips down to black except briefly between scenes. As for games, well, just consider which games really run a lot of "pure" black. If for some reason you're still playing in the dark, underground worlds of Doom 3, you need a high contrast ratio. But if the colorful world of Plants vs. Zombies is more your speed, you need not worry.

KEEPING UP WITH THE JONESES

Contrast-ratio specs are tremendously inflated. For the best LCDs, scientifically measured contrast ratios are actually between 1,500 and 2,000. But manufacturers almost never publish real contrast ratios anymore. You'll only find these true values in a small number of articles and publications. Yes, contrast-ratio values have been steadily improving over the years, but the year-to-year change is relatively small, which isn't good for marketing.

In their quest to quote ever-larger numbers, some manufacturers invented a completely meaningless spec called "dynamic contrast ratio," which is what's being prominently advertised now. Sometimes they don't even bother mentioning the "dynamic" part. Sadly, all manufacturers are now forced to play this game, as consumers wouldn't be interested in monitors and TVs that tout the true values. Meaningless contrast-ratio specs help substandard manufacturers by making their displays appear to be just as good as those from the best manufacturers, or even better, because the biggest liar wins. This not only hurts consumers, but it also hurts the better manufacturers because they're unable to publicize their superior technology.

BIG BUT NOT BIG ENOUGH

So what's really so "dynamic" about this bastardized contrast-ratio spec? It's really quite simple: When the display's video signal is entirely black or very close to black, the display's electronics go into a standby mode that significantly reduces the light output of the unit. This much darker standby value is then used when computing the contrast ratio—instead of the real value of when a picture is actually present.

Obviously, this trick doesn't change the true black or true contrast ratio for any picture that's not all black, so it's meaningless for picture quality. The primary reason for measuring the spec this way is that published contrast ratios can now go from about 1,500 up to, well, infinity. In 2008, many TVs were advertised with a "dynamic contrast" in the range of 15,000 to 35,000. Now, in 2010, some go into the millions and beyond. There's no real improvement, of course. It's just the same trick with a bigger exaggeration.

WHEN INFINITE MEANS NOTHING

At my local Walmart in Amherst, New Hampshire, the Sony KDL-52EX700, an LED-based TV, is listed by Walmart as having an "infinite contrast ratio" on its information label. First of all, "dynamic" was left off—it should say "infinite dynamic contrast ratio." This is then technically correct because the LEDs turn off when an all-black image is present. This results in a division by zero, and produces the infinite result. But this is also nonsense because the LEDs need to be on whenever an actual picture is present!

Labeling like this is intentionally misleading to consumers. Walmart should set an example for other retailers and refuse to show misleading manufacturer specs to its shoppers.

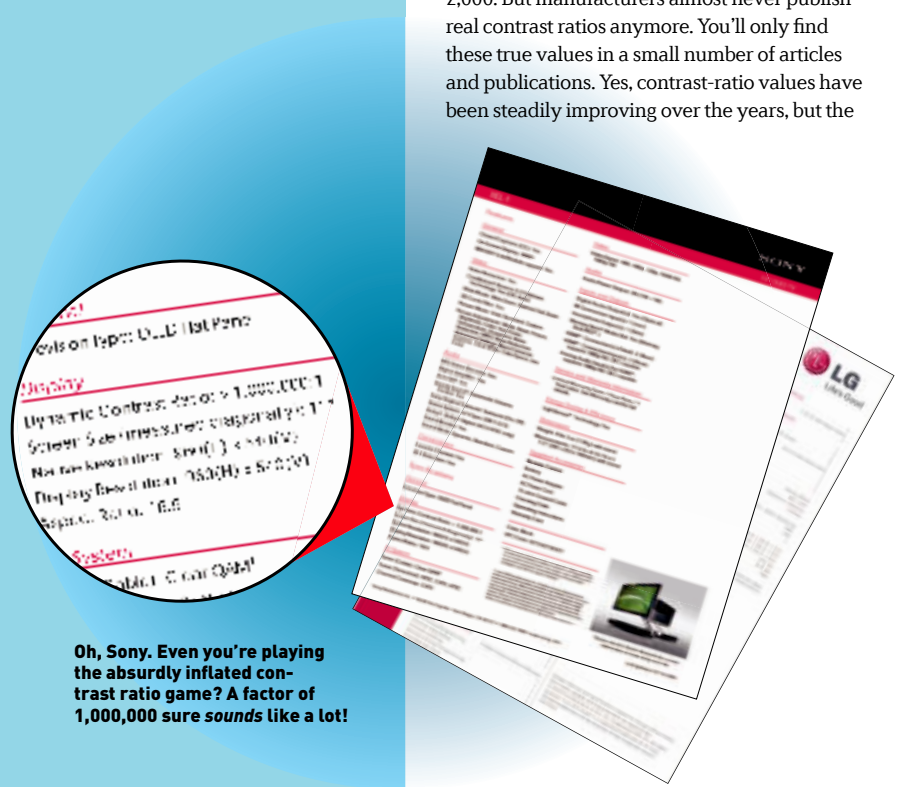

Oh, Sony. Even you're playing the absurdly inflated contrast ratio game? A factor of 1,000,000 sure *sounds* like a lot!

Response Times: How Fast Is Fast Enough?

All displays show artifacts of one sort or another when their screen images change rapidly. It's most easily detected with moving objects, or when the entire screen moves due to camera panning. In many cases, it's not the fault of the display. Rather, it arises somewhere in the signal path from the source, which can be caused by camera blur, interlaced scanning, MPEG compression artifacts, poor video processing, insufficient bandwidth, or insufficient CPU speed in the case of games. Further confusing matters, artifacts can occur for different reasons with CRT, LCD, LCoS, plasma, and DLP technologies. It can even occur with OLEDs, if switching speeds aren't sufficiently high.

But when people discuss motion artifacts, they are generally talking about LCD response time. And not surprisingly, the manufacturers' published specs for response time have become one of the major deciding issues for many consumers. As a result, in the last five years or so, manufacturers have somehow pushed response time numbers from 25ms (milliseconds) to an essentially untenable 1ms.

So what, if anything, do these specs really mean?

Witness two photos shot by Lauren Soneira. On the left, a static image—thus no motion blur. On the right, the image is moving at 1,018 pixels per second, and you can see six refresh cycles indicating a 50ms response time.

BEHIND THE BASICS OF BLUR

Motion blur arises when the liquid crystal—the active element within an LCD—is unable to change its orientation and transmission rapidly enough when the picture changes from one frame to the next. Because the standard video rate is 60 frames per second, a pixel is expected to fully update its light-transmission opacity within 16.7 milliseconds (that is, in one 60th of a second). If it takes any longer than that, the image will show some degree of lag, which appears as a trailing smear or blur whenever there is motion.

LCD motion blur is generally evaluated with an industry-standard spec called response time that (supposedly) measures the time that it takes for a pixel to go from black to peak-intensity white, and then back to black again. However, most picture transitions involve much smaller and more subtle shades of gray-to-gray transitions, which can take longer to complete.

But it gets even more complicated than that because every pixel is actually made up of independent red, green, and blue sub-pixels that have their own separate intensities, frame-to-frame transitions, and times. The upshot is that visual blur within a detailed, moving picture is a fairly complex and nebulous phenomenon.

MOTION BLUR: VISUAL PROOF

Motion blur is one of the most visually tangible display problems—the evidence speaks for itself in screenshots and photography, both of which can illustrate the relationship between response time and motion blur. In this article, I've included high-speed screenshots of moving DisplayMate test patterns, as well as moving test photos taken of a top-of-the-line, 120Hz Sony HDTV (shot with a Nikon DSLR using a fast shutter speed of 1/160 of a second). These images were taken in 2008, but the results wouldn't be much different today.

Sony's published response time for this XBR model is 8ms. Since this corresponds to a double transition (from black to peak white, and then back to black again), the single transition time (from black to white, or from white to black) should therefore be about 4ms.

But is the pixel response time really that fast? To find out, I ran DisplayMate tests in which black and white squares move across the screen at measured speeds. In the examples here, one photo shows the squares racing across at 1,093 pixels in a single second. The second photo shows the squares moving about 50 percent faster, covering 1,609 pixels in a single second. The white tips seen on the edges of the ghost images are artifacts resulting from electronic overdrive processing that's being used to try to improve the response time by exaggerating transitions.

As you can see from my screenshots—each a brief snapshot in time with a shutter speed of 1/160 of a second, which is less than the refresh rate—it's possible to make out at least eight individual screen refresh cycles on this 120Hz display. Indeed, in the screenshots, each square is shifted from the other by 1/120 of a second, which is 8ms, and those ghosted squares indicate that the older images haven't yet dissipated. The upshot is that you're looking at a true response time of about 65ms. In fact, a response time of much less than the 8ms refresh rate would be needed for there to be no visible blur. Obviously, 65ms blur in the screen shots doesn't jibe with the manufacturer's single-transition response time spec of 4ms.

The DisplayMate tests clearly demonstrate that the Sony's real LCD response time is considerably longer than its published spec would indicate. And by no means are we picking on Sony, as it actually had the best performance of all of the LCDs in our tests.

The screen on the left was grabbed from a motion speed test of 1,093 pixels per second, with the black and white squares indicating real-world pixel response times of 60ms and 40ms, respectively. The screen on the right shows the same lengthy response times during a faster test of 1,609 pixels per second.

BUT WHAT ABOUT MOVING PHOTOS?

It must be stated that DisplayMate test patterns are very sensitive to imaging effects—this is by

Behold our 11-display shootout—with the lights on. Photo by Dieter Michel.

design. Photographic images, on the other hand, typically consist of a very complex and varied admixture of blended picture elements. With so much going on in an image, motion blur is easily obscured and lost within the complex, variegated imagery of a typical photograph. In particular, photographs of real-world content lack uniform backgrounds, and uniform backgrounds—like those in my DisplayMate tests—make it easier to see the motion blur trails. For this reason, we would expect moving photographs to show much less visible blur than what's demonstrated with test patterns.

To wit: If you look carefully at the magnified marching band images on the previous page, you can see a total of at least six refresh cycles in the second photograph. This corresponds to a real-world pixel response time of 50ms. But the motion blur is still much less noticeable than what we see in DisplayMate's punishing test patterns.

Photographs are static images and moving them across a screen is quite different from live video, where images are a complex and varied mixture of continually blending picture components that are themselves constantly changing in both time and position. With all this screen activity going on, we would expect to detect much less motion blur in live video than with either of the moving static photographs or test patterns.

AND NOW FOR THE TESTS YOU'VE ALL BEEN WAITING FOR

To evaluate motion blur and artifacts in live video with lots of high-motion picture content, we set up a side-by-side comparison shoot-out with 11 HDTVs, and had both consumers and experts evaluate them. The top-of-the-line LCDs from Sony and Sharp had 120Hz screen refresh rates, the top-of-the-line Samsung LCD had strobed LED backlighting, and the other units had standard 60Hz screen refresh rates. Two of the units were plasma displays, and one was a pro-grade CRT studio monitor. The goal was to determine the degree to which this varied technology affected visible motion blur.

All of the HDTVs were fed identical, simultaneous digital video using an all-digital High Definition Tivo and a Blu-ray player. They were all compared side-by-side in the configuration as shown in the photo. The content included both daytime and nighttime sporting events, TV shows, and movies, all with lots of action. If any viewer thought he or she detected motion blur on any HDTV, we would repeatedly press the eight-second Tivo backup button and watch the sequence over and over again on all of the units until we fully understood exactly what was happening on each display. We did the same thing with the Blu-ray player and its content.

The conclusions from all participants were consistent across the board, and will likely surprise most consumers: There was essentially no visually detectable motion blur on any of the LCD HDTVs in any of the video content we assembled.

When people thought they saw motion blur, with only a handful of minor exceptions, the blur was either in the source video or a temporary visual illusion that disappeared when the segments in question were reviewed. Unlike what we empirically identified in moving test patterns and moving photographs, the eye is unable to detect the blur in live video because the images are much more dynamic and complex—and undoubtedly because of the way the brain processes and extracts essential information from visual images.

SO, IS BLURRING EVEN AN ISSUE FOR VIDEOS, MOVIES, AND GAMES?

For all of the tests—the DisplayMate test patterns, the moving photos, and the live video—we found that there was no visually detectable difference in motion blur for the mid- to top-of-the-line LCD HDTVs. This regardless of their claimed pixel response times, 60Hz or 120Hz refresh rates, strobed LED backlighting, or motion-enhancement processing. If you find this surprising then just re-read the classic tale of *The Emperor's New Clothes*.

The underlying reason why higher refresh rates don't mitigate blurring is that the true pixel response times of displays are considerably longer than the 60Hz video frame rate, so it doesn't matter whether the screen refresh rate is 60Hz or 120Hz, or whether the LED backlights are strobed off during the frame updating. Similarly, adjusting the electronic processing enhancements that some models offer—controls that are supposed to reduce motion blur—only served to introduce objectionable contours, edges, and other artifacts onto moving objects without reducing the overall motion blur.

So that's the story on video. What significance do these results have for PC gamers?

First, while motion blur isn't generally noticeable with live video, it's more likely to be seen by gamers who intently focus on particular moving objects. For this reason, the blur illustrated above with test patterns and test photos applies.

Second, don't pay much attention to a manufacturer's response time specs because they are so different from the real response time and motion blur that we have demonstrated here.

THE EYE IS UNABLE TO DETECT THE BLUR IN LIVE VIDEO BECAUSE THE IMAGES ARE MUCH MORE DYNAMIC AND COMPLEX

Third, while 120Hz refresh rate monitors and HDTVs don't inherently improve on motion blur over the 60Hz models, they are generally equipped with better performing panels and electronics, so they may still produce superior image and picture quality. And if you're a movie buff, the 120Hz units should offer better motion interpolation from the 24 frames per second used in all movies shot on film. The 60Hz models need 3:2 pull-down, which produces judder, but most people seldom notice it.

Fourth, be aware that the latest 240Hz displays don't offer any real picture-quality performance improvements, and are just a marketing gimmick taken to an absurd level.

For more information and details, see my article on LCD response time and motion blur here: www.displaymate.com/LCD_Response_Time_ShootOut.htm.

Color Gamut or Marketing Gambit?

Color gamut, which is the range of colors that a display can produce, is undoubtedly the most misunderstood and exploited spec—precisely because it's natural to believe that the range should be as large as possible. While that's true for most specs (even when they're exaggerated), it's definitely not the case for the color gamut.

The color gamut that you want on all of your PC monitors, laptops, HDTVs, and even smartphones is the same color gamut that was used when the content you're viewing was created. If a different gamut is employed, you'll see different colors than you're supposed to see.

Virtually all consumer content is created using industry standards that specify the exact color gamut to be used. For computers and digital cameras it's sRGB. For digital HDTVs, it's called ITU-R BT.709 (often referred to as Rec.709). Fortunately, both of these standards specify the same exact gamut. Yes, there are other color gamuts for specialized applications (more about that later), but sRGB and Rec.709 cover virtually all consumer content, and that is the color gamut you want on all of your displays. The color gamut in these standards specifies the exact color coordinates for the three red, green, and blue primary colors, which are used to produce all color mixtures on screen.

Now that you're versed on what color gamut is, let's share what it isn't—as illustrated by six examples of manufacturer misinformation.

BIGGER ISN'T BETTER

One common misconception frequently exploited by manufacturers is that a wider color gamut indicates a better display that produces more realistic colors. This is absolutely wrong. A larger gamut will simply make all of the screen colors for standard production content appear more saturated than they ought to appear. Indeed, displays claiming more than 100 percent of the standard color gamut simply can't show colors that aren't in the original source image. Expanded gamuts are just gimmicks that make consumers think they're getting something better.

THE PERILS OF RECALIBRATION

If you do get a display with a larger color gamut, it's necessary to reduce the gamut back to the sRGB/Rec.709 standard values by adjusting color saturation via a user control. Unfortunately, if the display isn't calibrated at the factory to match the standard color gamut, it's unlikely you'll be able to visually adjust it properly yourself. This kind of adjustment typically requires professional calibration using instrumentation.

NTSC? NEVER!

The often-quoted NTSC Color Gamut is from 1953. It's also obsolete and irrelevant. Computers, digital cameras, and HDTVs use the sRGB/Rec.709 color spaces, and specs should refer to them instead of NTSC. As explained above, values greater than 100 percent of the standard color gamut aren't desirable—unless you like punchy, unrealistic, oversaturated colors.

ADOBE RGB OR NOT TO BE?

As stated above, there are specialized color gamuts for specialized applications, and some of these are larger than the sRGB/Rec.709. Adobe RGB, one of the more common ones, is used by imaging professionals and you'll find it as an option on some digital cameras and scanners. Just be aware that if you use the Adobe gamut, you will also need a display that produces the Adobe gamut, and only a small fraction of consumer displays can do this. If you display an image produced with an Adobe gamut on a monitor with a standard sRGB/Rec.709 gamut, the colors will be incorrect and oversaturated.

HOW THE EYES PLAY TRICKS

Adobe RGB is a larger gamut than sRGB/Rec.709, but be aware that for most applications, gamut size doesn't matter very much. The further out you go in color space, the less frequent the colors appear in nature, so the human eye doesn't notice that they're not quite right except in rare circumstances (like when viewing a full-screen rendering of a very red tulip). When faced with a gamut beyond their rendering range, displays simply wind up reproducing the closest most saturated color they can under the circumstances.

BIT-DEPTH MISCONCEPTIONS

Manufacturers will also dupe consumers by advertising useless and misleading specs about the number of screen colors produced by their displays. Screen color counts have absolutely nothing to do a with display's color gamut, though manufacturers will attempt to tie them together. In reality, a display's maximum number of colors is a function of the total number of intensity-level combinations that the device can produce.

Let's do the math. Standard 24-bit color has eight bits per primary color, and eight bits generate 256 intensity levels. Because there are three primary colors, the number of possible color combinations is 256^3—16.8 million colors. Now, if a manufacturer uses 12-bit color processing internally within the same display, there are (in theory) 4,096 intensity levels and 68.7 billion possible colors. Sounds impressive, yes, but the display's color gamut remains the same as before and the additional number of colors doesn't mean anything visually.

EXPANDED GAMUTS ARE JUST GIMMICKS THAT MAKE CONSUMERS THINK THEY'RE GETTING SOMETHING BETTER

Still not convinced? First, remember that essentially all consumer content is 24-bit color. Thus, the source images have only 16.8 million colors, and the display can't "invent" intensities and color combinations that don't exist in the original.

Second, true onscreen 24-bit color does a good job of meeting the human eye's color and brightness discrimination abilities. You can read more about that here under "Digital Granularity": www.displaymate.com/ShootOut_Part_3.htm.

Third, be aware of the real reason why additional processing bits are necessary. Onscreen intensity levels are not supposed to be linear. Rather, they should follow a standard gamma curve with a nonlinear mathematical 2.2 power-law exponent (meaning the screen brightness for any sub-pixel varies as $s^{2.2}$, where s is the input signal intensity level 0-255). The extra processing bits are necessary just to get the display to produce the gamma curve accurately on screen.

Maybe Sharp Should've Consulted Mr. Spock Instead

By now, you've surely seen ads for Sharp's Quattron four-color technology. George "Mr. Sulu" Takei dons a lab coat, and fawns over Sharp's introduction of a yellow primary-color sub-pixel to the traditional three-sub-pixel, RGB primary-color arrangement. According to Sharp, this results in "expanding the color gamut and faithfully rendering nearly all colors that can be discerned with the unaided human eye, especially golden yellow."

If you have read this far, you already know that Quattron is just another shameful marketing gimmick. HDTV television and movie content is produced and color-balanced on three-color displays that are accurately calibrated to Rec.709. Sharp's fourth primary color is yellow, and there isn't anything for it to do because yellow is already being accurately reproduced with mixtures of the existing red and green primaries. More importantly, a Quattron display can't show colors that aren't in the original three-color source image. So what good is it? None, unless you like to see over-exaggerated yellows.

But could it be that existing consumer HDTVs are unable to reproduce the standard sRGB/Rec.709 color gamut, so Sharp's fourth primary color actually has something useful to do? We decided to find out.

Colors and color gamuts can be accurately measured and then plotted in a chromaticity diagram to compare values to the standard. What follows is from a 2008 article where I used a spectroradiometer to measure the color gamut of HDTVs in the DisplayMate lab. To the right, are the results for a Sony consumer LCD HDTV. The black triangle is the Rec.709 standard and the red dots are the measured values for the red, green, and blue primary colors of the Sony display. Notice that the Sony measurements all fall exactly where they should on the triangle vertices. It's perfect! In short, this Sony HDTV accurately shows exactly the same colors seen by, say, the director at a TV studio. Ipso facto, Sharp's fourth color is absolutely superfluous and can only decrease picture quality and accuracy! Undoubtedly, part of the Quattron's "Yellow Push" is being produced with simple video processing. Some people have been impressed watching the Sharp demos on the Quattron, but manufacturers demos are always fine-tuned to get a maximum wow response, so be careful before jumping to any conclusions about how it will perform displaying content at home.

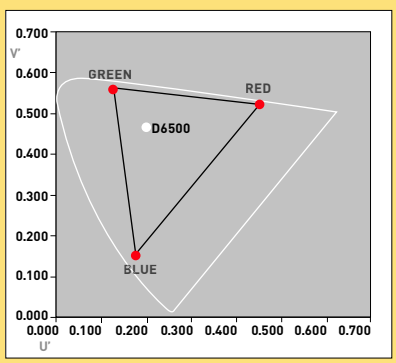

See that narrow strip of gamut real estate between the top of the black triangle and the inside of the white line? That's where Quattron technology would have to live.

Well, if Shatner can debase himself as a TV pitchman, can we really begrudge Takei?

Note that in our figure, the outer white curve represents the limits of human vision. While the Rec.709 standard is much smaller, it's important to note that the colors between the black triangle and white curve aren't common in nature. Yes, a display can only reproduce the colors that lie inside of the polygon formed by its primary colors, but because yellow falls between the red and green primaries, Sharp's yellow primary would need to lie somewhere *outside* of the red and green leg of the color triangle. But there isn't much room between the Rec.709 triangle and the human vision curve, is there? For this reason, it's difficult to see why a yellow primary sub-pixel is needed unless Sharp isn't able to put its red and green primaries where they belong.

Sharp shows its Quattron color gamut in some promotional material by using an old (x,y) distorted CIE Diagram from the year 1931 because it makes its extended color gamut look much larger than it really is. Our figure is the (u'v') 1976 Uniform CIE Diagram and shows the color gamut accurately.

ENDING DISPLAY FRAUD

It's both shocking and sad that display specs have been exaggerated to the point of meaninglessness. And you're not the only one who suffers—innovative manufacturers that develop new and better display technologies can't trumpet their hard work with superior performance specs. Instead, they're forced to play the game or lose significant business.

The National Institute of Standards and Technology (NIST) could help, but its display division was terminated in 2009. The only realistic solution that I see is the creation of an organization (that is completely independent of the manufacturers) to develop a set of straightforward, objective standards for measuring and advertising display specs. Manufacturers that meet those standards would be allowed to advertise their specs with a special controlled trademark, like the EnergyStar program. Consumers would learn to only trust specs with that trademark.

I proposed this back in 2003, but it went nowhere because too many manufacturers resisted the idea. But it's high time for this solution to finally be implemented—or just imposed. It's in everyone's interest except for the subset of manufacturers that can only compete using fraud.

Dr. Raymond Soneira is President of DisplayMate Technologies Corporation (www.displaymate.com), which produces video calibration, evaluation, and diagnostic products for consumers, technicians, and manufacturers. A research scientist with a career that spans physics, computer science, and television system design, Soneira was a Long-Term Member of the Institute for Advanced Study in Princeton, and a Principal Investigator in the Computer Systems Research Laboratory at AT&T Bell Laboratories.

30 MILLION MOBOS FROM NAN-PING

Inside Gigabyte's Taiwan-based motherboard manufacturing plant, tech reporter Dave James learns first-hand how the company's P55 motherboards come to be

Three million motherboards is a lot. Just think about the amount of solder alone that must take, the number of Japanese capacitors, the store of copper needed to fulfill Gigabyte's two-ounces-of-copper pledge for each of its mobos. But the thing I was most struck by during my visit to Gigabyte's Taiwan-based Nan-Ping factory was the incredible amount of human labor it takes to assemble and test each and every one of the 3 million boards the factory is capable of producing in a single year. That's 3 million boards assembled and tested by hand.

But the manufacturing process, with all its immense scale, is only part of the story. Before the Chinese-manufactured slivers of PCB even make it to the gates of Gigabyte's high-tech factory on the edge of Taipei's Xindian City, the work of designing the new chipset at the heart of the boards has been underway for a full year or more, while the research behind it stretches back even further.

Sitting among the clouds above Taipei's banking district, high up in the 101 Tower, I had the privilege of talking to Gigabyte vice president of motherboards, Henry Kao, and manager of the product planning division, Jackson Hsu, just as Intel's P55 chipset was all set to go global. I learned what exactly goes into putting a new chipset design out on the street.


GIGABYTE
MSI
ASUS
ACER
FOXCONN
GIGABYTE

Gigabyte vice president of motherboards, Henry Kao

Manager of the product planning division, Jackson Hsu

THE ORIGIN STORY

The Taiwan-based company began work on the latest P55 chipset a year in advance, as one of Intel's key development partners. "We don't need the Intel reference design board, because we make the reference design board for Intel," says Henry Kao, grinning broadly. "Intel has the idea for a new chipset and once it gets it to a certain level, the company brings the design idea to Gigabyte so we can codesign the reference board." It's this board that gets taken around to the other labs so Intel can show them the new design layout.

"With a new chipset, we usually spend one year co-developing with Intel. Once Intel feels comfortable, it releases that design to the other motherboard makers. From there we'll make our own specifications and it'll be another month before the first working samples are available," says Henry Kao.

The initial design team working on the co-development project will usually be a very small group of people. "At that stage you don't need many people involved, just a single team, dedicated to Intel," says Jackson Hsu. "Once the chipset is more mature, then we expand the team."

Once the reference design has been finalized, the various motherboard manufacturers can fashion their own different spins of the same board, creating high-, mid-, and low-end iterations of the chipset.

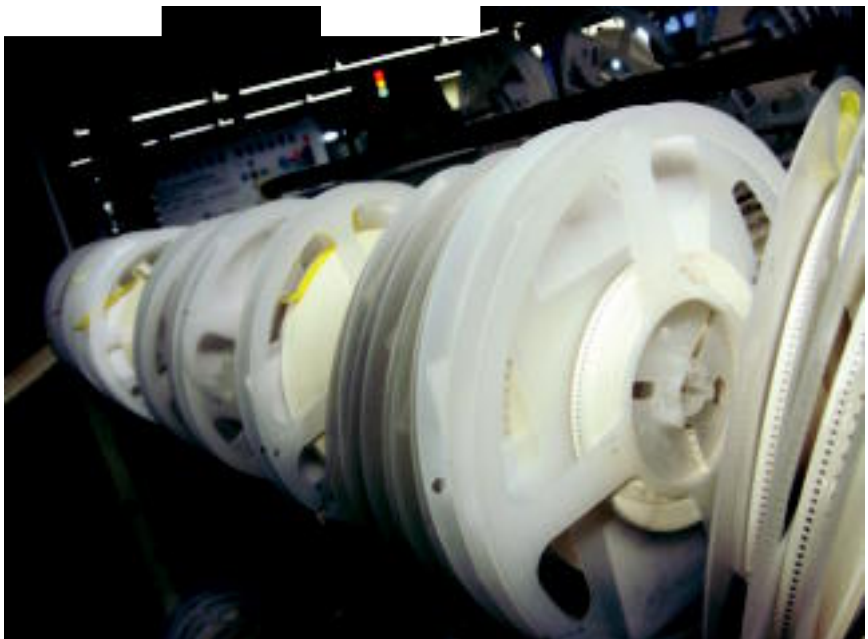

All of the smallest components for the motherboards are stored neatly on these racks.

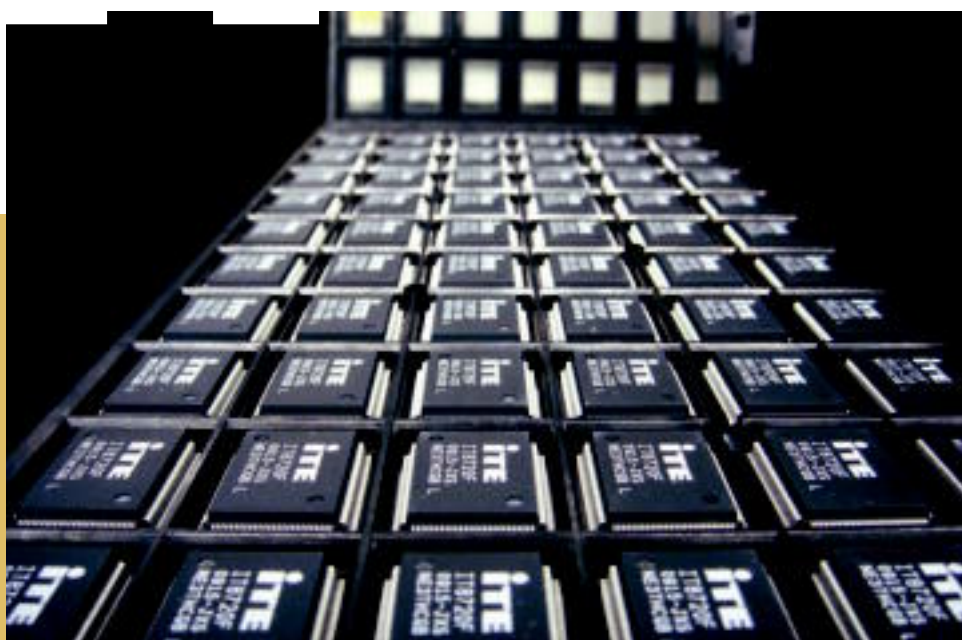

A big batch of controller chips stands by, ready for placement on the printed circuit boards moving down the line.

As a manager in the product planning division, Hsu has to be heavily involved in this stage: "As a planner you know the basic chipset and CPU layout but need to start thinking, 'What do I add? What would my competitors do and what does the customer need? And what is the latest technology?'

"Once you've identified all the possible elements, you have to start separating them out. What do you need for your highest high-end, then lower down the price scale, what do you take off? It's the same process every year, but the elements are different."

So, how does a manufacturer choose which components go into the higher-end boards and which get taken out? As ever, the overriding principle is cost. Sticking in the extra couple of USB slots that Intel asked for on the P55 boards isn't a problem across the full range of mobos, given the miniscule cost, but other parts and features inevitably need trimming in the move from the high-end.

GIVE AND TAKE

What's typically sacrificed in the lower-end motherboards is the newer technology, or the extent to which it's used: "I can't have everything 24-phase," says Hsu about the 24-phase power that's used for stabilizing

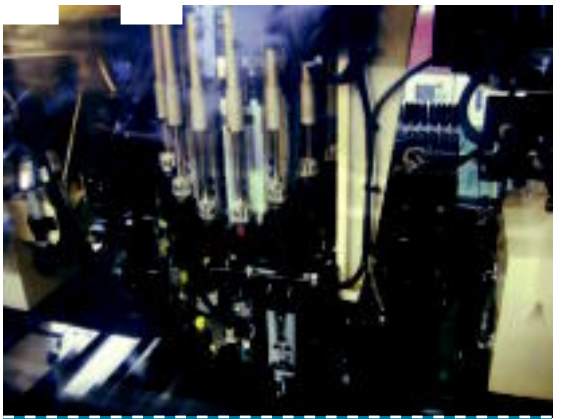

That blur is the high-speed placer nailing components into the waiting PCB.

the power signal on the top-end P55 motherboards. That has to be reduced as you move down the board pecking order, from 24-, to 12-, to 8-phase. Another example is with 6GB/s SATA: "Most of our P55 boards will have 6GB/s SATA; on the mainstream and high-end boards we'll have four ports, but lower down the line we'll have two ports. So even at the low-end, consumers will still have access to 6GB/s SATA, but it's limited," says Hsu.

But cost concerns don't necessarily govern Gigabyte's motherboard division as a whole. The current division CEO is Johnson Lin, who rose through the ranks right from the engineering floor, and thus has hands-on experience with what goes into both the design and manufacturing processes.

He's responsible for the aforementioned two ounces of copper that Gigabyte's performance boards are known for. Slotted between the layers of PCB is this relatively generous layer of copper, absorbing and distributing the huge amount of heat that is generated by the processor, chipset, and power-phase regulators. It has made the motherboards slightly more expensive to produce, but in the long run it has reduced the number of RMAs to such an extent that all of Gigabyte's motherboards are now adopting the same process.

Moving half of the company's motherboard manufacturing to this fairly new Taiwan-based factory isn't about cost, either. The factory workers expect, and get, a far higher wage than their Chinese counterparts in Shenzhen, for example. And with three shifts working eight hours each, that's a lot of man (or young girl, for the most part) power to pay for every day.

But this factory produces high-end motherboards and graphics cards, as well as notebooks, mobile phones, servers, and full desktop PCs. And having a high-tech factory of this sort on the doorstep of the research and development teams at Gigabyte means that any necessary product-design changes can be turned around rapidly. As Jackson Hsu told me later: "Some of our customers still like to have that 'Made in Taiwan' stamp on Gigabyte products."

PUTTING IT ALL ON THE LINE

So once all the design work has been done, and all the various iterations of a specific chipset have been hashed out, it's on to the manufacturing process itself.

The first thing that hit me—after I'd been decked out in sexy blue overshoes, a surgical mask, and had every speck of dust blown off my bloated form by the biggest hair dryer this side of Sassoon HQ—is the amount of motivational material dotted around the factory floor. Much of it, in fairly broken English, seems intended for the international visitor. A particular favorite of mine was a storyboard showing how a poorly fixed nut caused the explosion of a Chinese Boeing 737—a reminder to focus on even the smallest things or face the direst of consequences. Another favorite was a sign in both Chinese and English instructing workers to "Be more responsible, complain less, be more attentive, and make lesser mistakes."

WE DON'T NEED THE REFERENCE DESIGN BECAUSE WE MAKE THE REFERENCE DESIGN

The company is doing something right, given the Nan-Ping factory's ability to shunt out a quarter of a million motherboards every month, as one of four Gigabyte factories in Taiwan and mainland China.

There are four stages to the motherboard manufacturing process carried out in Nan-Ping. First is the surface-mount technology

Here's the result after the high-speed and multifunction placers have done their job.

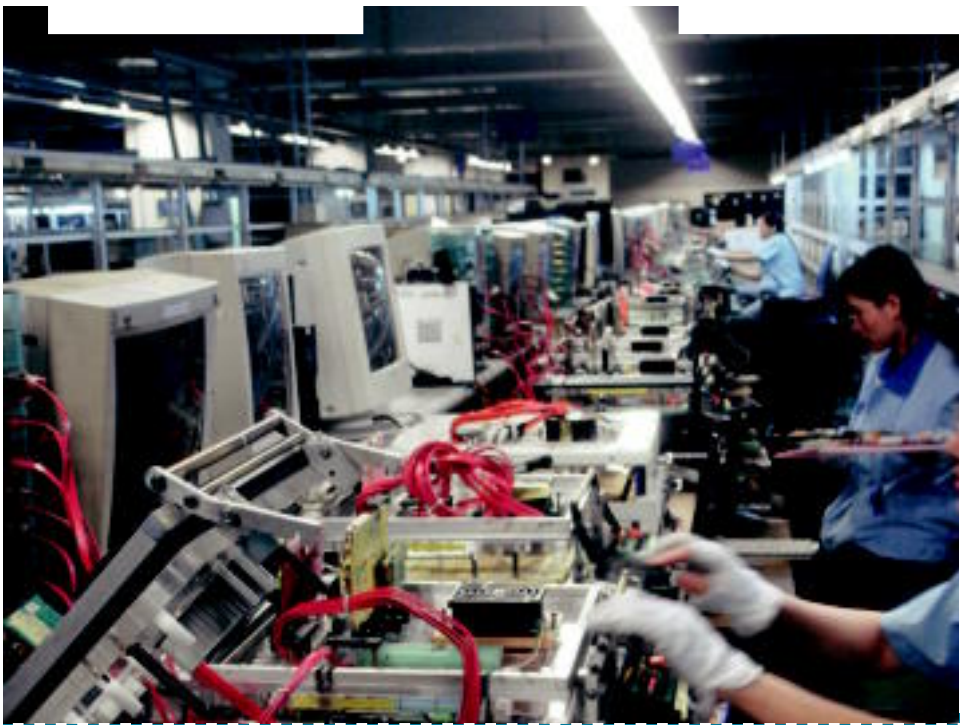

At a row of in-circuit testing benches, Gigabyte workers verify the electrical integrity of the newly placed components.

(SMT), where the smallest components are placed on the bare printed circuit boards (PCB). Next is the dual inline package (DIP), where the larger elements are attached. Then there's the testing stage, and finally, it's down to packaging. Each of these processes is carried out on separate production lines, and often on separate floors. For a single motherboard to go through the entire process, though, takes only around 15 minutes from the start of production to boxed finish.

The first step, the SMT process, is the most automated of the lot. All of Gigabyte's PCBs are manufactured in mainland China to its own specifications, and shipped over, bare, to the various factories. Once they hit one of the 11 SMT production lines in the Nan-Ping factory, each PC board is put first through a solder-paste printing machine. This machine lays out the areas of solder needed to hold the components, which are then placed in the following step.

The smallest components, right up to the north-bridge and south-bridge chips, are held in different-size rolls, feeding their chips into either the high-speed or multifunction placers. These placers are like component nail-guns, shooting chips, resistors, and capacitors into the motherboard with seemingly reckless abandon. In fact, they are incredibly precise, with the high-speed placer inserting a component every 0.1 seconds.

Once all of these smaller parts have been placed on the board, the entire PCB gets moved into an oven for what is called reflow soldering, which is basically melting and resetting the existing solder to lock the new components in place.

The final part of the SMT process is the testing. (At every step, the motherboard goes through the same testing procedure.) First comes the inspection, then in-circuit testing (ICT). In this case, with up to 1,300 components on a motherboard, the inspection is done automatically with an optical inspector, which uses a strobe to check that each component is correctly placed. This takes seconds. Then there's the first manual inspection, then the in-circuit testing, which checks the electrical performance of the newly placed components.

DIP AND COVER

In the DIP process, things become a whole lot more labor intensive. There are only four production lines in this part of the factory compared to the 11 of the SMT, but there are 50 operators on each of them manually inserting parts into the PCB. The components that are added in this stage are things like the PCI connectors, memory slots, and all the inputs and outputs.

Once all these larger components are in place, the board is then put through a wave-soldering machine to secure these latest additions. Then it's passed on to

Every single PCI-E slot gets mounted by hand—remember that the next time you install a new GPU.

At every step of the manufacturing process, workers inspect the boards to make sure everything is where it should be.

Within 15 minutes of first entering the assembly line, the once-bare printed circuit boards are now fully operational mobos, ready for retail.

another manual technician to touch up any missing solder or unsecured components. It's at this stage that all the heatsinks are attached to the chipset. Once attached, there follows an almost identical inspection and in-circuit test as in the SMT process.

On the same floor as the DIP procedure are the six testing lines, where all the boards get a complete checkup. First is the function test, where the motherboards are placed in testing benches for a full functionality verification. Then comes the burn-in test, where each board is subjected to a high temperature of 45 C and then a low of -10 C, just to make sure the final product can cope with the daily rigors it might encounter. From here it's on to the final manual inspection and then on to that final stage of packaging.

With such finely tuned electronics, it's tempting to think of the manufacturing process as being completely automated, something akin to a car manufacturing plant, but with far, far smaller mechanical arms. In fact, there's a surprising amount of human labor involved in every step of the process. So, the next time you purchase a motherboard (or graphics card), or even just peek inside your PC, take some time to think of exactly how many hands that bit of PCB has gone through on its way to you.

PCB Manufacturing Process Flow

The transition from bare circuit board to retail-ready motherboard happens in four steps, and takes about 15 minutes

1 SURFACE-MOUNT TECHNOLOGY	2 DUAL INLINE PACKAGING	3 TEST	4 PACKING	SHIPPING
Solder-paste printing	Manual insertion	Function test	Box forming	
Component placement	Wave soldering	Burn-in test	Put in accesories	
Reflow soldering	Touch up and heat-sink assembly	Inspection	Packing	
Test and inspection	Test and inspection			

WHITE PAPER

Internet Protocol Version 6

How the next-generation Internet protocol, IPv6, will save the day by increasing the number of available addresses and adding new features -ZACK STERN

Are you prepared for the next Y2K? We've dubbed it "The Great IP Address Shortage of 2011—or Shortly Thereafter." While "TGIPAS2011oST" doesn't roll off the tongue as easily as "Y2K," this creeping calamity could cause fundamental networking problems the day someone claims the very last public IPv4 (Internet Protocol version 4) address.

Fortunately, a solution is already here: IPv6 (Internet Protocol version 6), will create trillions (times trillions times trillions) more public addresses while introducing new networking features. Both IP technologies will coexist during the transition, so in many cases, IPv6 devices will switch to an IPv4 mode to communicate with old gear.

We'll explain the basics of how Internet layers shuffle information across worldwide networks. Then we'll discuss the changes and improvements in IPv6; it's about more than adding addresses. With these details, you'll be able to prepare for the dual-stack, hybrid near-future and the eventual transition to IPv6 exclusivity.

PEELING BACK THE LAYERS

The Internet uses four main layers to transmit data. You might think of them as nesting matryoshka dolls, with each layer in the progression encapsulating the next. (You can read more about this topic in the White Paper in the August 2009 *Maximum PC*. You'll find the issue in PDF form at http://bit.ly/gGSgC).

The Link Layer includes the hardware and software (or firmware) to handle the lowest level of communications, including Media Access Control (MAC). (While some define hardware connections as part of the Physical Layer below this, we'll side with those who combine the two.) The Internet Layer rides on top of that, consisting largely of IPv4 or IPv6. The Transport Layer transmits and delivers data to specific application protocols residing in the Application Layer, including HTTP (Hypertext Transfer Protocol) and FTP (File Transfer Protocol).

The Internet Layer is especially important since you'll sometimes set its addressing while configuring network equipment, such as a router. More often, the other layers are invisible to end-users; things just work.

CHANGE OF ADDRESS

A new addressing method marks the biggest change between IPv4 and IPv6. IPv4 technology uses 32-bit addresses that allow about 4 billion possible public nodes. Due to growing demand and inefficient allocation of existing IP addresses, the planet has consumed nearly all of these. Estimates vary as to exactly when we'll run out of IPv4 addresses, but most experts anticipate it happening soon.

IPv6 solves the problem in a big way. It uses a 128-bit addressing scheme that creates 2^{128} possible combinations: roughly 34 trillion trillion trillion (34 followed by 37 zeroes) IP addresses—enough to support every Internet-connected computer, phone, television, gaming console, refrigerator, and toaster humanity could ever want. Here's an example of what an IPv6 address looks like: 2001:0DB8:AC10:2F3B:9C5A:FFFF:3FFE:02AA.

And that's not all that IPv6 will provide. It also eliminates the need for Network Address Translation (NAT), a form of IP masquerading developed to get around the IP address shortage by hiding entire address spaces. The primary drawback to NAT is that a host residing behind a NAT-enabled router does not

HOW IT WORKS

Transition to IPv6

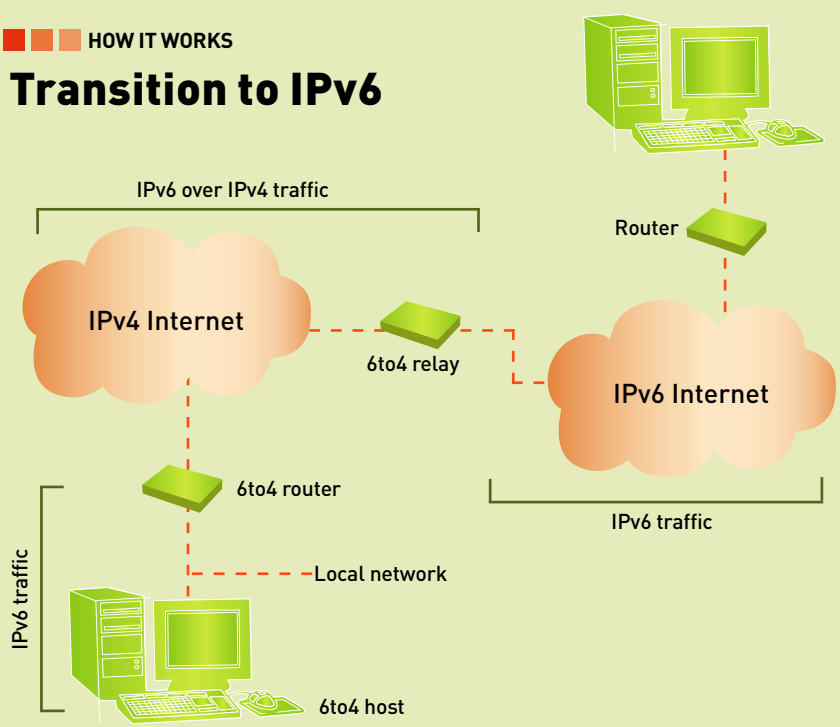

Since IPv4 and IPv6 will coexist for a time, IPv6 nodes will communicate with each other via tunnels through IPv4-only infrastructures using so-called "6to4" technology. This will enable devices (hosts, routers, etc.) with IPv6 address to reach IPv6 infrastructures without requiring a connection to the IPv6 Internet.

have end-to-end connectivity and so cannot utilize some Internet protocols.

Network devices connected to an IPv6 network will be able to auto-configure much more robustly than is possible with the existing Dynamic Host Configuration Protocol (DHCP). You can also simply reconfigure an existing network, keeping the address suffix—and subnet—while changing the public prefix. And since an IPv6 subnet can have 2^{64} addresses, ISPs and large institutions will no longer be forced to fragment their networks.

As with IPv4, data exchanged using IPv6 is contained inside packets. IPv6 packets consist of three elements: a fixed header with addressing information, an optional extension header that enables additional features, and a payload. IPv6 packets can be processed faster than IPv4 packets because their packet headers do not contain a checksum. Checksums are used in IPv4 to verify that data has been properly sent and received, but this task is performed at a higher layer with IPv6. In addition, some infrequently used processes have been moved out of the fixed header and into the optional header.

COUNTDOWN TO IPV6

All modern operating systems are ready for the transition to IPv6, including Windows XP with Service Pack 3. For the moment, operating systems can run in a dual-stack mode, juggling the two IP standards and often two IP addresses at the same time. For example, they might use IPv6 internally and route to an IPv4 destination externally.

Operating systems also support "6to4" translation. This method allows IPv6 devices and networks to communicate across IPv4 sections of the Internet. The technique stores an IPv6 packet inside the payload of an IPv4 packet, like a ferry transports cars across a river. Relay servers or a dual-stack destination on the other side can unlock the IPv6 data once clear of the IPv4 network. As long as you have a public IPv4 address, the technique should work even if your ISP hasn't yet adopted IPv6.

It's harder for IPv4 devices to communicate with IPv6. Translational gateways and proxies can help but they aren't reliable. Once your favorite Internet services switch over to IPv6-only, your IPv4-only equipment might not work. Thankfully, the dual-stack transitional phase should stave off that problem for many years.

3M Mpro150 LCOS Projector

We disassembled this pico projector to see how 3M's engineers packed a self-contained presentation system into a package that's just 5.12 inches long, 2.4 inches wide, and 0.98 inches thick. You can connect a PC to this device, but you don't have to—the projector hosts its own software and offer

LENS ARRAY The lens mounted to the LED module focuses the light coming from the LEDs. A set of lenses inside the black box takes the polarized light from the LEDs and directs it to the LCOS imager. These lenses then redirect the light from the LCOS imager to the focal lens inside the large gray module. The user adjusts the focal lens to focus the image on the screen.

LCOS IMAGER AND ASIC This component fits vertically into the slot in the lens housing next to it and hosts the LCOS (liquid crystal on silicon) imager and the ASIC (application-specific integrated circuit) that drives it. The LCOS imager creates the image that is projected onto the screen.

LEDS AND HEATSINK The LEDs provide the light source for the LCOS; the heatsink cools the LEDs.

ARM CPU This processor performs the projector's digital signal processing work and the bulk of its other functions.

FLASH MEMORY AND MICROSD CARD SLOT The projector is outfitted with 1GB of embedded flash memory and can host up to a 16GB MicroSD card (3M includes a 2GB card), so you can store documents, photos, and presentations right on the device.

ANALOG INTERFACE PROCESSOR This chip processes the analog video sources connected to the projector (VGA, composite, and component).

MICRO-CONTROLLER This chip monitors the projector's keypad and other digital control functions.

Please note that we have flipped the primary PCB over in order to expose some of the projector's most important components.

SUBMIT YOUR IDEA Ever wonder what the inside of a power supply looks like? Don't take a chance on destroying your own rig; instead, let us do the dirty work. Tell us what we should crack open for a future autopsy by writing to comments@maximumpc.com.

HOW TO

Step-by-Step Guides to Improving Your PC

THIS MONTH

PIN MULTIPLE FOLDERS TO THE WIN7 TASKBAR

ALEX CASTLE
MANAGING ONLINE EDITOR

Have you ever wished you could pin more than one folder to your Windows 7 taskbar? If you try to, it simply saves them all in a single stack, requiring two clicks when you only want to have to click once. Fortunately, there's a clever hack that gets around this problem.

Windows 7 allows you to pin shortcuts to the taskbar. Unlike folders, shortcuts to executable files never stack. So, make a shortcut to any executable that's not already pinned to the taskbar, and drag the shortcut to the taskbar. Once it's pinned, it'll stay pinned, even if you modify it. That means that you're free to right-click the shortcut in the taskbar, then right-click again on the name of the program and select Properties. This will allow you to change any of the properties of the shortcut, including its name, icon, and target. Select an appropriate name, point it at the folder you want, and voilà! This way, you can add as many folders as you want to the taskbar.

SUBMIT YOUR IDEA Have a great idea for a How To project? Tell us about it by writing to comments@maximumpc.com.

WINDOWS TIP OF THE MONTH

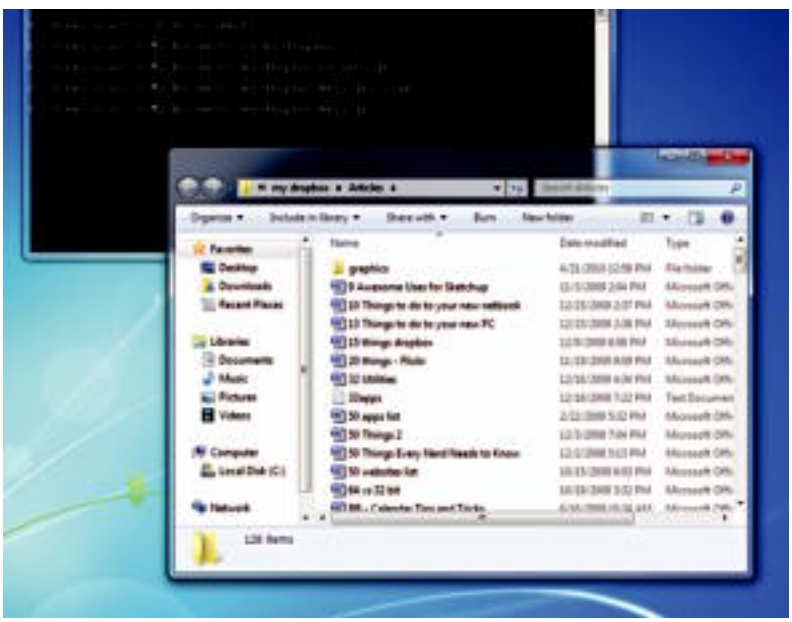

Use CMD and Explorer.exe Together

You probably already know that you can open a command prompt from the explorer in Windows 7 and Vista by holding shift and right-clicking, but did you know you can also do the reverse? To open an explorer menu at your current location in the command prompt, simply enter the following command (including the period): `start.`

Manage and Convert Your Ebooks with Calibre

With the arrival of the much-hyped iPad and the rest of tablet-mania, it seems like ebooks are about to have their "iPod moment," when they'll go from a favorite of early adopters and bibliophiles to a mainstream phenomenon. There's one problem, though: Unlike MP3s, there's not a single, near-universal standard for ebooks. Historically, this has made it difficult to organize your ebooks and transfer them between various reading devices.

Fortunately, there's one program that can help you solve nearly all of your ebook-related problems: Calibre. A free, open-source project, Calibre is one part iTunes-esque library-management program, one part batch-conversion tool, and one part file-transfer manager. In this article, we'll show you how to use Calibre to manage your ebooks and to get them working on any reader. —ALEX CASTLE

1 SET UP CALIBRE To get started with Calibre, visit the project homepage at http://calibre-ebook.com and download the latest version. Let the .msi do its job and install the program on your computer, and then launch the program. If this is your first time running Calibre, you'll be presented with a few choices.

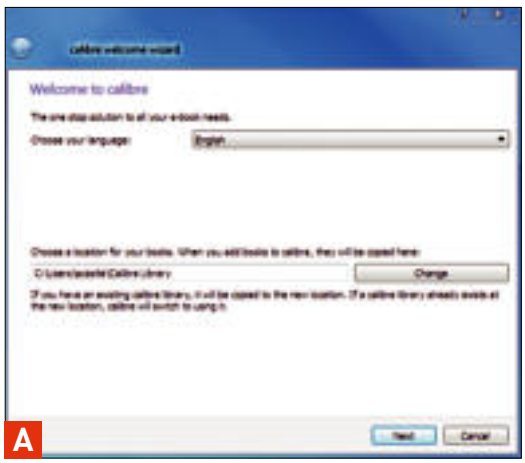

A

First, you'll be asked where you want to save your Calibre library (image A). Calibre will use whichever location you select as a home for its hierarchical file system, containing all your ebooks. It's worth noting that whenever you import a book into Calibre, it will actually create a copy of that ebook in its library, which can take up a lot of space.

Next, you'll be asked to select your preferred reader. If you have one of the readers on the list (and pretty much every reader *is* on the list), make sure to select it, because this will change the default output settings Calibre will use when converting ebooks. These settings optimize the outputted file for reading on the specific device you choose, and take into account screen size as well as file type and DRM restriction.

If you told Calibre that you have an iPhone or iPad, you'll be asked if you'd like to run the Calibre content server. We're going to show you how to use the content server to access your files, so go ahead and choose yes, but don't worry if this option didn't pop up for you—it's easy to turn the server on at any time.

2 ADD BOOKS TO CALIBRE

Now that you've finished setting up Calibre, you should be looking at a depressingly empty ebook library (image B). Unless you plan on reading the Calibre Quick Start Guide over and over again, you're going to want to add some books to your library. For use as an example, we downloaded a handful of classic books from the Gutenberg Project (and one classic comic book from goldenagecomics.co.uk) in different ebook formats.

Loading books into Calibre is easy. Start by clicking the Add Books button in the top left corner, then browse to a folder containing ebooks (image C). You can select multiple books, then click Open to import them all to Calibre. If you've already got an ebook library spread out over multiple subfolders in a directory, you can import all books from a directory tree by clicking the small arrow to the right of the Add Books button.

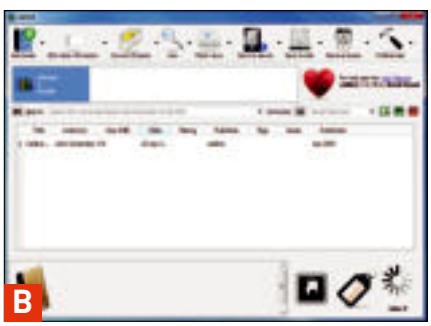
B

Hopefully, Calibre has properly read the metatags from the books you've chosen to upload, and you'll now be looking at a populated, searchable library. Unfortunately, many ebooks are missing proper metadata, so you may have to add it in yourself. There are three ways to do this.

First, you can simply click the book, click the Edit Meta Information button, and enter the information (such as title, author, and publisher) by hand (image D). This works fine if you only have a few books you need to add data for.


D

E

Second, you can click Edit Meta Information, enter the book's ISBN number (easily found on Amazon) and click Fetch Metadata from Server. This will automatically pull down all the information Calibre can find about the book, and is generally very accurate.

Finally, if you want to enter the same metadata for a whole bunch of books at once (say, an author for a series of novels or comic books) you can select all the books at once, and click the Edit Meta Information.

Once you've gotten the important metadata (title and author, at the least) entered for a book, you can open the metadata editor for that book and click Download Cover. This will attach a cover image to the book (image E), which allows you to use Calibre's slick Cover Viewer feature. If Calibre fails to download a cover, or downloads a different cover than the one you want, enter the ISBN for the version of the book that you want, and click Download Cover again.

3 BROWSE YOUR LIBRARY

There are a couple of ways to browse through your ebook library in Calibre. The default is a simple list view, with a search bar at the top, which lets you view only books with a certain title, author, or tag. Tag view, which lets you view your library in a branching, hierarchical structure, can be accessed by clicking the tag icon in the lower right corner of Calibre.

You can also view your library in a cool-looking Cover View mode (image F) by clicking the arrow button in the lower right corner of Calibre. Cover View

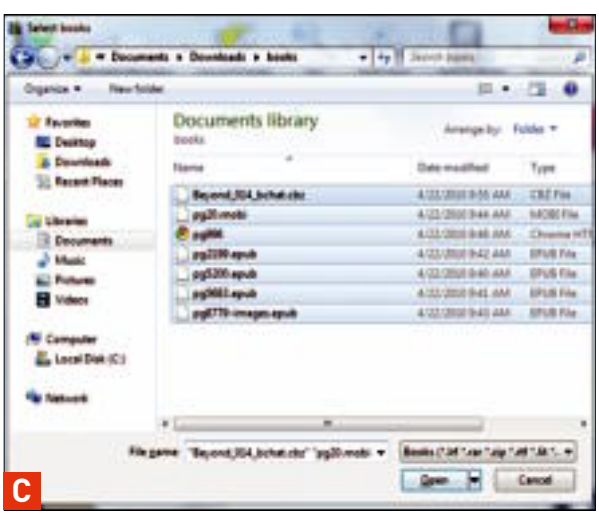

C


F

will show you the covers of whatever books are currently being viewed in either search or tag view. If you want search view to open in a separate window, or to view more covers at the same time, those options are available in Properties > Interface.

4 CONVERT YOUR EBOOKS

Calibre isn't just an ebook organizer, it's also the most powerful tool for batch ebook conversion. This is handy, because in the current, volatile ereader market, ebook standards change constantly, and it's easy to end up with a collection of books in multiple formats.

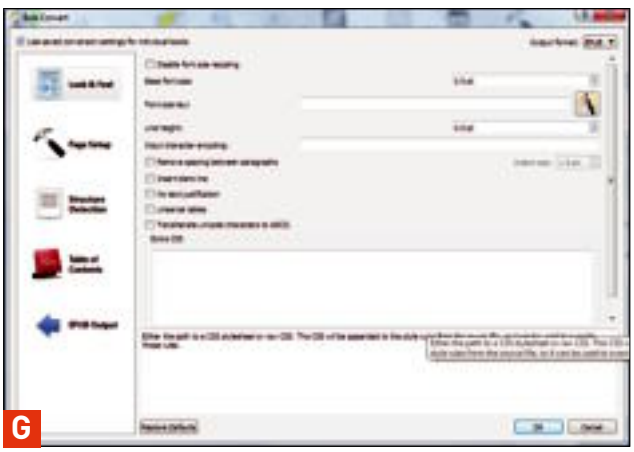
G

To convert ebooks, simply select them, and click the Convert button in the upper left of Calibre. A conversion window will open, with pages and pages of options you can dig through, to produce a perfect copy. For most books, however, the default settings will work and you can just click the OK button to start the conversion. If you want to convert to a different format than the default (which is based on the reader you selected when installing Calibre), you can do so using the drop-down menu in the upper right corner of the convert window (image G).

It will take Calibre a little while to convert each book (you can view its progress by clicking the Jobs button). When it's done it will store the created file in its library, alongside the original. You won't see the new file in Calibre by default, unless you click a title, then click the View drop-down list, then select View Specific Format.

5 TRANSFER FILES

Transferring files from Calibre to a reader is easy, though it works differently, depending on which reader you have. For most readers, such as the Kindle, Nook, or Sony reader, you simply connect your reader to your computer, select the books you want to transfer, then hit the Send To Device button. For the iPhone or iPad, you'll need to turn on the Calibre content server in Preferences > Content server, then download an app that can retrieve books from a remote source, such as readMe (for iPad) or Stanza (for iPhone and iPod touch).

If you want to be able to access your ebooks from anywhere using a web browser, that's also accomplished using the Calibre content server. You simply activate the content server and leave both your computer and Calibre running (image H). After that, you can view your library from any computer, by pointing a web browser to your static IP address (or, if you don't have a static IP, your dyn-dns.com account).

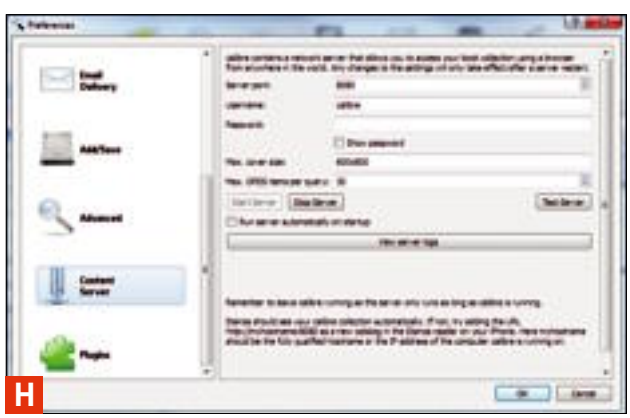
H

Use Your PC Remotely with LogMeIn

Mobile devices like smartphones, tablets, and netbooks have always had a trade-off: What you get in convenience, you lose in good old-fashioned power. Even as modern smartphones close the processing-power gap, you still can't do everything you could do at your primary PC. Or can you?

In this article, we're going to show you how to use remote-desktop software to control your PC from another PC or mobile device. There are several programs that let you remotely control a computer, but in our experience LogMeIn offers the most useful and consumer-friendly software in the category. In light of that, we're going to show you how to configure and use LogMeIn Free and LogMeIn Ignition to get desktop-grade power, anywhere. –ALAN FACKLER

GET STARTED WITH LOGMEIN FREE

LogMeIn Free is (you guessed it) a free version of LogMeIn's software that allows you to remotely access your PC or Mac computer from any computer or mobile device, anywhere, just so long as there is an available Internet connection (image A).

Getting started with LogMeIn Free is fast and simple. Visit www.logmein.com and select LogMeIn Free from the pull-down tab below Products, then hit Download Now. Like most web-based software, you'll have to create an account to use the service.

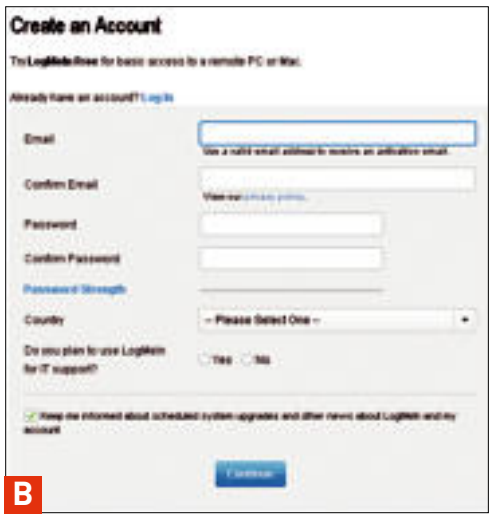

B

When prompted, fill out the information to create an account (image B). The next menu will ask you to choose whether you'd like to install the software on the computer you are currently using, or on a different computer altogether. In order to keep things simple, use the computer that you want to remote control when downloading this software. After clicking Add This Computer you will be asked to download and install the remote-access software and run the installation wizard.

C

Installation is very simple—LogMeIn walks you through every step. Keep in mind, besides creating an account name (just your email) and password, you may also need to type in your desktop's user name and password if your OS is password protected. The first time you log in, LogMeIn will ask you whether or not you would like to allow the program to make changes to your computer. Hit Yes to ensure the smoothest possible synch between your desktop and wireless device. Once a connection has been created, punch in your desktop's user name and password, and presto!

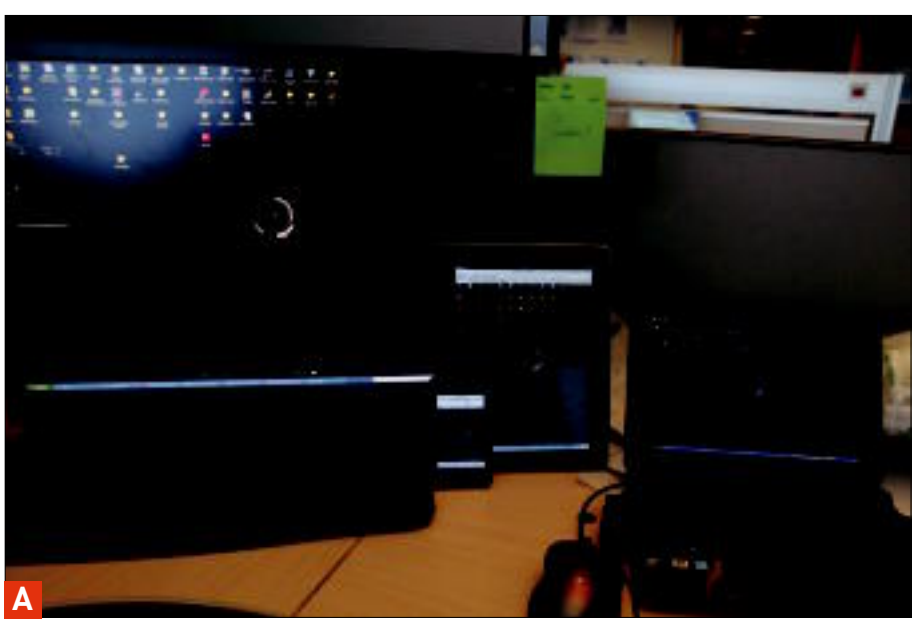
A

This will take you to LogMeIn's interface, which will give you a couple of options. In order to synch with your desktop, hit Remote Control (image C). This will launch the most important part of LogMeIn—the remote desktop.

LogMeIn's browser-based remote desktop lets you perform almost any task you normally would using your PC, over the Internet. You can save and make changes to files, reorganize your desktop, surf the net, and even remotely run heavy-duty apps like Photoshop, even on

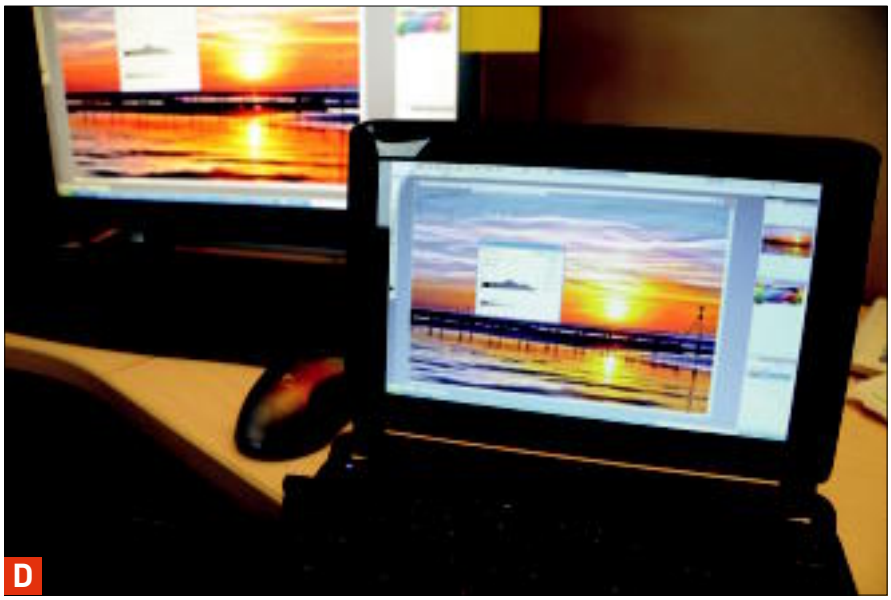
D

an underpowered netbook (image D).

Another cool feature is LogMeIn's Customize Toolbar, found at the top of the screen (image E). This drop-down bar offers customizable widgets to help you remotely control your desktop, including a magnifying glass (very useful for netbook use) and laser pointer, as well as a plethora of connection options if you'd like to sync your keyboard or change your screen settings. These helpful tools can really streamline your experience.

E

All in all, if you want to access your favorite computer remotely from a laptop or netbook, LogMeIn Free is a must-have.

SMALL DEVICES: READ ON

Using the free software with smaller devices, like a smartphone or iPad, however, is a different experience.

Remotely controlling a computer using LogMeIn Free works basically the same way on an iPhone or other smartphone as on a PC—you fire up the browser and navigate to LogMeIn.com. However, when using these smaller devices, LogMeIn Free acts as more of a "glimpse" than a workable interface. Your desktop will appear, and you will be given limited access to its content. You can open folders, Word documents, and the explorer, but the functionality ends there.

Navigation around your desktop is nowhere near as streamlined as with a more well-equipped device; dramatic movements cause lag, and double-clicking icons often doesn't produce the desired result. Instead, anytime you try to access a program or open software the page has to completely refresh, giving you a kind of slide show view of your computer, rather than an actual workable interface. Fortunately, there's another way, but it costs.

LOGMEIN IGNITION

If you have an iPhone or iPad, a program called LogMeIn Ignition addresses these issues, for a price. LogMeIn Ignition is available exclusively in the iTunes App Store for $30. What does that price tag get you? Near-desktop-level responsiveness, on a mobile device.

After downloading, simply launch the app and you will be asked to punch in your user name and password (image F). The process of logging into your account and beginning a remote session is exactly the same as with a mobile computer. The interface, however, functions differently.


F

Unlike LogMeIn Free's slow-moving, point-and-touch interface, your finger is now the mouse, which can be dragged anywhere on the screen.

Ignition for iPad takes advantage of the device's touch screen controls. Pinch your fingers to zoom in and out, single-tap items to select or drag them, double-tap to open them. The two-finger control scheme also allows you to easily right-click items or scroll through your web browser. The toolbar is simpler than LogMeIn Free's toolbar, but has all the important elements.

A keyboard icon allows you to type whenever you need to, as well as a Ctrl+Alt+Del icon, which allows you to remotely troubleshoot your computer if you're having issues offsite. There's also an in-depth settings menu, which allows you to adjust color quality, resolution, and network speed, just to name a few of the options.

REVIEWS

Tested. Reviewed. Verdictized.

INSIDE

ONLINE

- EVEN MORE REVIEWS!
- BEST OF THE BEST
- BREAKING TECH NEWS
- NO BS PODCAST

iBuypower Paladin F

Hexa-core and Fermi on the cheap

Most readers know the name iBuypower by now, but they don't know our nickname for the company: iStealpower.

OK, that's not really true, we just made that up to make this story sound sexier, but there is some truth to our jest. Over the years, we've often wondered how the hell these guys can offer PCs for less than the cost of the parts. You know, like getting $2,900 worth of parts in a machine that cost $2,200.

We're not sure if the cost of the parts in iBuypower's Paladin F exceeds the price of the machine, but it probably gets close. The Paladin F sports Intel's new hotness: the hexa-core 3.33GHz Core i7-980X (clocked up to 3.8GHz). Even with AMD's new hexa-core CPU now on the market, Intel's Core i7-980X is still clearly the recognized fastest CPU in *der verold*! To the 980X, iBuypower adds Nvidia's top-dog GeForce GTX 480 card, aka Fermi. Also aboard are 6GB of Kingston DDR3/1600, a 1-kilowatt PSU, an LG Blu-ray combo drive, a 1.5TB hard drive, and RAID 0 SSDs, along with Windows 7 Home Premium. The entire system is embedded in a Zalman GS1000 Plus enclosure.

Our first issue with the Paladin F stems from the storage configuration. In the rush to get a low-cost RAID 0 config in the PC, iBuypower used a pair of 30GB Kingston SSDNow V Series drives. That's 60GB for your primary boot partition. We started to get nervous about space just by installing our 20GB of benchmark files on the machine. Fortunately, there's also a 1.5TB Seagate drive in the mix—but come on, 60GB?

On the performance front, the Paladin F turned in good, but not insane, application numbers. Against our zero-point machine, a 2.66GHz Core i7-920 overclocked to 3.5GHz, the 3.8GHz Core i7-980 was about 35 percent faster in Vegas Pro 9 and Main Concept Reference. In ProShow Producer, the Paladin F only managed a 17 percent faster score. Lightroom saw a 15 percent boost.

The Paladin F turns in solid performance without the price—or pizzazz—of other gaming rigs.

In gaming, however, our zero-point dominated. Yes, the GeForce GTX 480 is the fastest single-GPU card in existence, but our zero-point uses a Radeon HD 5970 card, the single-fastest videocard in existence. That dual-GPU 5970 saw the Paladin F run roughly 23 percent slower in both Far Cry 2 and STALKER: Call of Pripyat. The asterisk on our 5970 card, though, is the cost and availability. You'd be lucky to find a Radeon HD 5970 for less than $700—if you can find one at all. It's not an easy card to get, folks.

Of course, neither the Paladin F nor our zero-point can hold a candle to Digital Storm's tri-CrossFire Hail Storm (reviewed in May) nor Main Gear's tri-SLI Shift (reviewed in June). But here's the kicker: Those two super PCs were $7,800 and $6,900, respectively. The Paladin F is $3,000.

Is it the fastest rig in *der verold*? No. We're not even sure it's the fastest in the Tri-State area, but for $3,000, you're getting Intel's best and a DX11 card that's actually almost as fast as the dual-GPU 5970 card.

In the end, despite a couple configuration issues, the Paladin F is a solid machine and a darn good value. **–GORDON MAH UNG**

SPECIFICATIONS

Processor	Intel 3.33GHz Core i7-980X (OC'd to 3.8GHz) with Acetek 550LC cooler
Mobo	Gigabyte GA-X58A-UD5
RAM	6GB Kingston DDR3/1600 in tri-channel mode
Videocard	GeForce GTX 480
Soundcard	Onboard Realtek
Storage	Two Kingston 30GB SSDNow V Series, 1.5TB Seagate Barracuda
Optical	LG Blu-ray combo drive
Case/PSU	Zalman GS1000 Plus / 1000 Watt PSU

BENCHMARKS

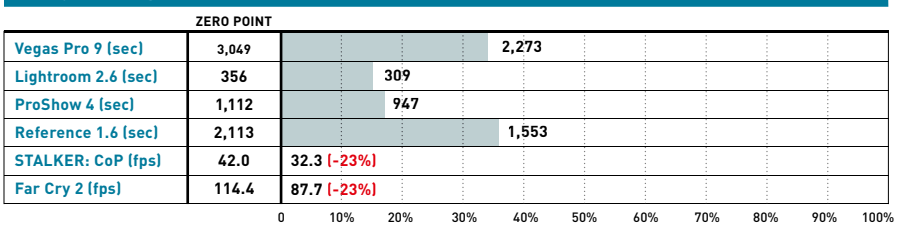

Our current desktop test bed consists of a quad-core 2.66GHz Core i7-920 overclocked to 3.5GHz, 6GB of Corsair DDR3/1333 overclocked to 1,750MHz, on a Gigabyte X58 motherboard. We are running an ATI Radeon HD 5970 graphics card, a 160GB Intel X25-M SSD, and the 64-bit version of Windows 7 Ultimate.

VERDICT 8

IBUYPOWER PALADIN F

+ RANGER	– ILLUSIONIST
Quiet; affordable; room to expand.	Incredibly small main partition; could use a second GPU.

$3,000, www.ibuypower.com

Sapphire Radeon HD 5870 Eyefinity Edition

Six monitors + one graphics card = eye candy heaven

Do gamers really need six monitors? Having two displays—maybe even three—on your desk certainly makes sense for a productivity boost. And having run some games on three displays, we can say that the added immersion in the game world can indeed be compelling. But you can run three displays with any Radeon HD 5000–series cards, provided you have at least one DisplayPort monitor.

Sapphire and AMD are betting that some gamers will lust after more than three displays, which is why Sapphire is shipping the Radeon HD 5870 Eyefinity Edition. This isn't just a stock 5870 with six monitor connectors; it also ships with a 2GB GDDR5 frame buffer. So even if you aren't planning on running six displays on your desk, the 2GB of VRAM might itself be attractive.

Eyefinity is flexible as to monitor configurations. You can have the six displays arranged in two rows, which can be configured as one huge surface or two 3x1 surfaces. Or, you can have five LCD panels in line for a wraparound gaming experience.

Sounds intriguing, right? But what about performance? We put the Sapphire card up against an XFX Radeon HD 5870 XXX Edition and the Asus GTX 480 card.

In single-display performance, what you immediately notice is that the Eyefinity Edition can't quite keep up with the overclocked XFX card, although the differences are pretty small. The Asus GTX 480 pretty much wins across the board. So if you're planning on using just one or two displays, the Asus GTX 480 is likely a better single-GPU card. But then, you don't want an Eyefinity Edition card for just a couple of monitors, right? To check the chops of the Eyefinity Edition, we ran the benchmarks in both three- and six-panel configurations, at 5160x1080 and 5160x2160, respectively. None of our three- or six-panel benchmarks had anti-aliasing enabled, which is fine since gaming performance is already so limited.

HAWX and Far Cry 2 seem to be playable on six panels, even with detail levels pumped up, but minimum frame rates will dip down into single digits. Three-panel gaming, though, looks viable across the board, particularly if you turn down some of the graphics detail.

The one real gotcha with this card is monitor support. The Eyefinity edition has six mini-DisplayPort connectors, but only five adapters ship in the box: two mini-DisplayPort-to-DisplayPort, two MDP-to-DVI passive, and one MDP-to-HDMI. So if you want to go with six-panel gaming, you'll need one to four more adapters, which cost about $10 each. That's in addition to the $500 cost of the card.

Is it worth it? If you want more than three panels with just one graphics card, this is the only game in town, short of some pro graphics cards that cost even more. So if you're lusting after—and have the room for—four or more LCD monitors on your desk, Sapphire's Radeon HD 5870 gets the job done. —LOYD CASE

VERDICT **8**

SAPPHIRE RADEON HD 5870 EYEFINITY

+ BUFFY	− WARREN
Four, five, and six displays with one card; 2GB VRAM.	Not the fastest card for the price; six-monitor gaming is performance-limited.

$500, www.sapphiretech.com

BENCHMARKS

	Sapphire HD 5870 Eyefinity	XFX Radeon HD 5870 xxx	Asus GTX 480
Unigine Heaven 2.0 (fps)	17	13	**26**
Battle Forge (fps)	47	49	**61**
Dirt 2 (fps)	71	71	**80**
Far Cry 2 / Long (fps)	75	78	**103**
Far Cry 2 / Action (fps)	63	65	**76**
Tom Clancy's HAWX (fps)	89	92	**104**
Crysis (fps)	32	**33**	31
STALKER: Call of Pripyat (fps)	37	38	**39**

Best scores are bolded. Our test bed is a 3.33GHz Core i7-975 Extreme Edition in an Asus P6X58D Premium motherboard with 6GB of DDR3/1333 and an 850TX Corsair PSU. The OS is 64-bit Windows 7 Ultimate. All games are run at 1920x1200 with 4x AA.

EYEFINITY RESULTS

	Three displays 5160x1080	Six displays 5460x2160
Unigine Heaven 2.0 (fps)	18	10
Battle Forge (fps)	30	18
STALKER: Call of Pripyat (fps)	27	18
Dirt 2 (fps)	42	25
Far Cry 2 / Long (fps)	62	35
Far Cry 2 / Action (fps)	47	30
Tom Clancy's HAWX (fps)	59	37
Dawn of War 2: Chaos Rising (fps)	46	29
Crysis (fps)	15	8

The Eyefinity Edition sports six mini-DisplayPort ports but only comes with five adapters.

Alienware M11x

At last, real games on a little laptop!

The terms petite and gaming notebook are about as incongruous as self-restraint and Wall Street, so our curiosity about Alienware's M11x was naturally mixed with skepticism. Could this sub–five pound, 11-inch rig do much more than play aged or anemic titles?

Small as it is, the M11x indeed has substance. The first sign of hope was the GeForce GT 335M graphics card—a slightly faster kin to the GT 325M we found in Asus's N61J 16-inch notebook (reviewed in May). Also stuffed into the wee chassis: a 1.3GHz Core 2 Duo overclocked to 1.73GHz (which can be turned off in the BIOS), 4GB of DDR2/1066 RAM, and a 7,200rpm 500GB hard drive. That's a lot of gear to cram into an 11.25x8.25x1.25-inch body—so much, in fact, that there's no room left for an optical drive.

Deciding on which notebook to use for our zero-point was tricky, as the M11x straddles the gaming and ultraportable categories. We settled on our most recently reviewed notebook, the aforementioned Asus N61J, an all-purpose mainstream machine. Like that notebook, the M11x features hybrid graphics so you can enable or disable the videocard to suit your needs. Whereas the N61J uses Nvidia's new Optimus hybrid graphics, however, the M11x features the old-school method, requiring a manual switch and reboot to implement. Be that as it may, when the M11x's discrete videocard was activated, we were impressed with its chops.

You can change the color of the M11x's various backlights by using the AlienFX utility within Alienware's Command Center software.

SPECIFICATIONS

CPU	1.3GHz Intel Core 2 Duo SU7300 OC'd to 1.73GHz
RAM	4GB DDR3/800MHz
Chipset	Intel GS45
Graphics	GeForce GT 335M / Intel GS45
Hard Drive	500GB Seagate ST9500420AS (7,200rpm)
Ports	VGA, HDMI, DisplayPort, Ethernet, three USB, FireWire, two headphone, mic, 3-in-1 media reader
OS	Windows 7 Ultimate 64-bit
Lap/Carry	4 lb, 8.6 oz / 5 lb, 4.3 oz

Sporting an additional 24 CUDA cores and a 20MHz core clock boost, the M11x's GT 335M performed 40 percent faster than the Asus's GT 325M in both of our gaming tests. And while the M11x's benchmark score of 18.2fps in Far Cry 2 might not strike you as special, consider that it hit a respectable 33.5fps when playing at the notebook's 1366x768 native res. Yes, on an 11-inch notebook! Other notebooks of this size buckle under Quake 4.

In the applications benchmarks, the M11x's proc was overpowered by the Asus's 2.26GHz Core i5-430M, not surprisingly. (For the record, although the BIOS reports the M11x's proc as 1.73GHz, we never saw it exceed 1.6GHz in CPU-Z during our benchmark tests). But compared with other ultraportable machines, the M11x is competitive. In fact, it beat out the premium Toshiba R600, our ultraportable zero-point, by 55.9 percent in ProShow and 18.4 percent in Photoshop. On top of that, the M11x surpassed that machine and others of its kind in battery life. With the graphics card disabled and the power-saving mode selected, we ran a video file full-screen on the M11x for five hours. Of course, the M11x's 8-cell battery, not to mention its videocard, make this machine a pound or more heavier than most so-called ultraportables. But even with the power brick included, we're talking less than five and a half pounds.

Aesthetically, the M11x sports touches befitting a gaming notebook: an Alienware-esque contoured body and LED backlighting behind the keyboard, power button, and front grills. It also comes fairly well-equipped with ports, including both HDMI and DisplayPort outputs. It lacks an Express Card slot, but mobile broadband is possible via the internal WWAN option (an extra $125).

Frankly, we're amazed at Alienware's accomplishment. Are there sacrifices? Sure. The biggest being the lack of an optical drive. But all things considered, that seems a fair trade for an otherwise thoroughly outfitted ultraportable that can actually play today's games.

—KATHERINE STEVENSON

BENCHMARKS

	ZERO POINT	
Premiere Pro CS3 (sec)	1,260	2,400 (-47.5%)
Photoshop CS3 (sec)	164	245 (-33.1%)
Proshow Producer (sec)	1,482	1,198 (-25.8%)
MainConcept (sec)	2,453	5,040 (-51.3%)
Far Cry (fps)	13.0	18.2
Call of Duty 4 (fps)	25.5	35.8
Battery Life (min)	125	300 (140%)

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

Our zero point notebook is an Asus N61J with a 2.26GHz Core i5-430M, 4GB DDR3/1066 RAM, a 500GB Seagate hard drive, a GeForce GT 325M, and Windows 7 Home Premium 64-bit. Far Cry 2 tested at 1680x1050 with 4x AA; Call of Duty 4 tested at 1680x1050 with 4x AA and anisotropic filtering.

MAXIMUM PC KICK ASS!

VERDICT 9

ALIENWARE M11X

+ AREA 51	− STUDIO 54
Uberportability, robust gaming, and long battery life in one kick-ass package.	No optical drive; no Express Card; old-school hybrid graphics solution.

$1,100, www.dell.com

Western Digital VelociRaptor 600GB

A much-needed refresh for WD's flagship speedy drive

For years, if you wanted the speediest consumer hard drive you could get your hands on, you went with a 10,000rpm Western Digital Raptor. Its first incarnation, released in 2003, was a 37GB single-platter drive using a PATA-to-SATA bridge. The next year brought a 74GB SATA 150 drive, and thereafter the drives roughly doubled in size (and went up a SATA spec) every couple of years. Last time WD refreshed the line, it bumped the capacity to 300GB, named the resulting 100MB/s-plus drive the VelociRaptor, and promptly won our Kick Ass award. But that was 2008—several hard drive generations ago. And though Western Digital's latest VelociRaptor ups the ante with 600GB of space and a 6Gb/s SATA controller, the drive now has to compete with solid state drives and high-capacity, high-performance drives like WD's own Caviar Black series.

Make no mistake: The new VelociRaptor, with its 32MB of cache and 6Gb/s transfer rates, is the fastest mechanical SATA drive we've ever tested. With average sustained read and write speeds greater than 130MB/s, it's fully a third faster than the last-gen VelociRaptor, which averaged around 100MB/s for both. Random-access times hit around 7.1ms—about the same as the last-gen VelociRaptor, and about twice the speed of a fast 7,200rpm drive.

But both solid state drives and 7,200rpm mechanical hard drives have come a long way since the last Raptor reigned supreme, and for sheer performance, the VelociRaptor isn't top of the line any more. And it's not even as far ahead of the rest of the mechanical pack as it once was—a 2TB WD Caviar Black costs $280 and offers nearly 115MB/s average sustained reads and writes, though random-access times are slower.

On the other end of the price/capacity/performance triangle, of course, are solid state drives. A 100GB SandForce-controlled solid state drive offers read and write speeds north of 200MB/s, random-access times measured in the tenths of milliseconds, and read and write IOPS in the 5,000s—for $400. With SSDs much more expensive but much faster, and 7,200rpm mechanical drives nearly as fast but with more capacity per dollar, where does the VelociRaptor fit in?

We still think there's a niche for the VelociRaptor. It does deserve big props for being the fastest mechanical SATA drive on the market, and its sustained read and write speeds are faster than the fastest first-gen SSDs. It's the best single-drive compromise between the capacity and price per gigabyte of a 7,200rpm drive and the raw speed of an SSD. The 6Gb/s SATA support helps bump up burst speeds, if not much else, and 600GB is enough for an operating system and every program and game you're likely to install. Power users are likely to prefer a solid state drive for OS and games and a 7,200rpm drive for storage, but that adds complexity: worrying about OS TRIM support, default library locations, managing two drives, and so forth—things you don't have to stress over with the new VelociRaptor.

The VelociRaptor isn't the fastest drive you can get, nor the most capacious, but if you just want one drive that does it all, the newly refreshed VelociRaptor is still a compelling choice. **–NATHAN EDWARDS**

The label reads Enterprise Storage, but we suspect that plenty of VelociRaptors will find homes in enthusiast rigs.

BENCHMARKS

	600GB VelociRaptor	600GB VelociRaptor	300GB VelociRaptor	Caviar 2TB Black
Controller	6Gb/s SATA	3Gb/s SATA	3Gb/s SATA	3Gb/s SATA
HDTune 4.01				
Min Read (MB/s)	**97.8**	**97.8**	62	68.7
Avg Read (MB/s)	**132.6**	**132.6**	99.9	114.2
Max Read (MB/s)	157.8	**158.2**	121.2	147.8
R.A. Read (ms)	**7.1**	**7.1**	7.2	11.7
Burst Read (MB/s)	**270.8**	201.2	210	197.6
Min Write (MB/s)	95.9	**96.4**	76.4	68.3
Avg Write (MB/s)	130.2	**130.4**	100.9	113.4
Max Write (MB/s)	**163.4**	155.9	121	146.9
R.A. Write (ms)	**7.1**	**7.1**	**7.1**	11.8
Burst Write (MB/s)	**270.8**	196.4	214.6	203
4KB Read (IOPS)	**145**	139	145	81
4KB Write (IOPS)	300	321	**384**	201
Premiere Pro (sec)	**285**	291	300	299
PCMark Vantage HDD	**7,881**	7,467	6,042	5,781

Best scores are bolded. All tests performed on an Asus P6X58D Premium motherboard with a Core i7-X980 CPU @3.33GHz with 6GB DDR3/1600 running Windows 7 Professional 64-bit. SATA 3Gb/s tests performed using the ICH10R south bridge; SATA 6Gb/s tests performed using onboard Marvell 9123 controller.

VERDICT 9

WD VELOCIRAPTOR WD6000HLHX

+ VELOCIPEDE	– CENTIPEDE
Fastest hard drive; 6Gb/s SATA; welcome capacity bump.	Occupies awkward performance/capacity middle ground.

$330, www.wdc.com

The iPad's wide black bezel is unsightly, but provides a good resting place for smudge-inducing fingers.

Apple iPad Wi-Fi

Touch navigation works—it just plain works. Netbooks beware!

Even if you reject the iPad on an intestinal level—you know, because you don't want to be associated with mock turtlenecks and man bags—then you should still view Apple's device as a referendum on the looming wave of tablet computers. The bottom line is that the iPad is damn useful. The referendum has passed.

The LED-backlit display clocks 1024x768 pixels across 9.7 diagonal inches. Those are netbook-like specs in a physical formfactor that's more attractive (both aesthetically and functionally) than any netbook. The best part about the screen is that it defines the iPad in toto—without the baggage of a hinged physical keyboard, track pad, or pointing stick, the iPad thrives when typing, web-surfing, or doing similarly simple tasks while lying on your back.

For couch-based computing, just prop the iPad in the angle of your lap. In landscape mode, the virtual keyboard stretches 7.75 inches edge-to-edge, providing easy typing targets. But it's really the iPad's touch navigation that makes the device so convenient when working in a prone position. Scrolling, panning, or just inserting a cursor is much easier when all you have to do is tap your finger. Simply put, touch-based navigation trumps every mousing alternative we've ever encountered on a netbook. And thanks to the iPad's 1GHz ARM processor—a system-on-chip that also handles 3D, audio, power management, and bus traffic duties—the GUI is extremely fast with speedy screen redraws. The iPad eschews mechanical storage for flash, aiding app responsiveness and load times, as well as battery life (which consistently beat Apple's 10-hour estimate for "normal use" scenarios).

When you're ready to hit the road, the iPad provides a winning mixture of convenience and functionality. Compared to most any netbook, it's lighter, slimmer, and more packable. It also boots faster—near instantaneously. For plane trips, you can watch HD movies in spectacular color and clarity, read ebooks, listen to iPod tunes, and play an ever-growing arsenal of addictive casual games. The touch interface is an Achilles' heel when it comes to shooters and other action games, but pool, pinball, board games, card games, and various PopCap ports translate wonderfully to the touch screen.

When you hit your hotel room, you can web-surf and email without compromise, but you'll finally hit some snags when you attempt more complex activity. The iPad runs the iPhone OS, which won't support multitasking until later this year. But the real bugaboo is the closed system of the OS. Neither Windows nor Mac software is welcome to the party. For 90 percent of all computing activity, Apple's App Store has an app for the job. But for the remaining 10 percent, you're on your own.

In the coming months, a host of PC makers will be rolling out their iPad killers, and we can't wait to see one that runs Windows apps, multitasks, supports Flash, and has a front-facing camera for video chat—all features the iPad currently lacks. HP's Slate appeared to be the front-runner in the non-Apple tablet race, but speculation has the Slate running Palm's webOS instead of the resource-hogging Windows 7.

Regardless of which OS they're running, all the new tablet competitors should dominate both netbooks and old-school, stylus-controlled tablets. Touch interface is the key, and it doesn't need to be Apple's. **–JON PHILLIPS**

SPECIFICATIONS

Footprint/Weight	H/W/D: 9.5x7.5x0.5 inches; 1.5lbs
CPU	1GHz Apple A4 (ARM-based SoC)
GPU	PowerVR SGX 535
RAM	256MB DRAM (integrated into SoC)
Storage	Flash memory (16, 32, or 64GB)
Display	9.7-inch, LED-backlit IPS LCD, 132ppi
Battery	Non-user-replaceable, 6600mAh lithium ion polymer
Input/Output	30-pin dock connector, 3.5-inch headphone jack, built-in speaker and mic

VERDICT **8**

APPLE IPAD WIFI 32GB

+ GRANNY CLAMPETT	– GRANNY SMITH
Speedy touch-based GUI; super-efficient SoC; lightweight and attractive; killer display and virtual keyboard; tons of killer apps.	Can't multitask; nonstandard OS; no Flash support; no camera; limited as a gaming platform.

$600 (for 32GB Wi-Fi only version), www.apple.com

Bowers & Wilkins MM-1 Computer Speakers

Yes, they're that good

B&W tells us it didn't need to move a lot of MM-1s to earn a profit—no surprise there—but would the system sound as sweet if it were produced en masse?

We've reviewed a number of great speakers, but we haven't been this excited about a set of boom boxes since we laid hands on the eponymous Cambridge SoundsWorks MicroWorks in the very first issue of *boot* magazine—way back in September 1996.

Listening to Peter Gabriel's new album *Scratch My Back* on Bowers & Wilkins' MM-1 computer speakers sent chills down our spines, a sensation rapidly followed by slack-jawed awe. We downloaded the album from B&W's Society of Sound music club, which has the exclusive rights to distribute the album in studio-master quality: losslessly encoded in FLAC with 24-bit resolution at a 48Hz sampling rate. The MM-1 delivers audio quality that's so exquisite, so pristine, that it makes the mighty AudioEngine A5—our previous favorite 2.0-channel speakers—sound almost muddy in comparison.

We were particularly impressed with the MM-1's bass response: There's no subwoofer, so this system will never rattle your floor joists; in fact, frequency response bottoms out at 57Hz. But those lows are absolutely clean, lusciously well rounded, and perfectly balanced with the most scintillating highs and midrange frequencies we've ever heard from desktop speakers. In fact, this system is already pushing the Klipsch Pro Media 2.1 off our Best of the Best list—and that product took up its perch only last March (you'll find our review at http://bit.ly/bP2aP6). And yes, we're taking games and movie soundtracks as well as music into consideration.

The MM-1 sounds so magnificent that we don't even mind that it's a USB speaker system that bypasses the PC's audio circuitry; in fact, that's a major plus—unless you've invested in a high-end soundcard like the kick-ass Asus Xonar Essence STX we reviewed in June 2009 (read our review at http://bit.ly/dPlQv). Our only other mild criticism is that B&W decided to use a proprietary cable hardwired to the right speaker (which contains the amplifier and DAC) to connect the left speaker, and the cable is just long enough to position the speakers on either side of a 30-inch display.

If you ever needed proof that sheer power doesn't necessarily ensure great sound, give the MM-1 a listen. The amp delivers just 18 watts to each of the four drivers (two one-inch Nautilus tube tweeters and two three-inch glass-fiber woofers), and while B&W considers them strictly near-field monitors, we found they had no difficulty filling our 80-square-foot home office with music. The speakers look as gorgeous as they sound, too, wrapped in stretchy black fabric with spun-aluminum tops and surface-mounted volume and power buttons embedded in the aluminum band around the right cabinet. You'll find a line-in jack for a digital media player and a headphone jack there, too. A flat, ovoid remote controls volume and the transport mechanism for your favorite music software (we tried it with Media Monkey, iTunes, and Foobar 2000). Thick rubber bases on both speakers isolate them from your desktop.

We'd like to tell you that the MM-1 is comparably priced to speaker systems made by the likes of Logitech, Altec Lansing, or even Klipsch, but 500 clams puts it at the opposite end of the spectrum. If you're as passionate about audio quality as you are about the rest of the components in your rig, on the other hand, that's not such a high price to pay. —MICHAEL BROWN

SPECIFICATIONS

Amplifier Power	18 watts x 4 channels
Frequency Range	-6dB at 57Hz and 22kHz
Tweeter	One-inch Nautilus tube (two)
Woofer	Three-inch glass-fiber (two)
Inputs	1/8-inch line-in; USB 2.0
Outputs	1/8-inch headphone

MAXIMUM PC KICK ASS!

VERDICT **10**

BOWERS & WILKINS MM-1 SPEAKERS

+ SIGNAL	− NOISE
Exquisite audio reproduction; elegant design.	Very expensive; short-ish, hardwired cable on right speaker cabinet.

$500, www.bowers-wilkins.com

Asus ENGTX480

The fastest single GPU can be made even faster

Just about everyone knows that Nvidia's hot new Fermi graphics chip is literally hot. So, when Asus bundled its new ENGTX480 card with a custom voltage tweaker for overclocking, we wondered if it was such a good idea.

After all, do you really need the card to run hotter? And with the speed of the ENGTX480, you probably don't need the higher clocks anyway. The ENGTX480 ships with 32 shader processors (what Nvidia calls "CUDA cores") disabled, yet the card still manages to be the fastest single-GPU card you can buy today.

In addition to 480 shader processors, the GPU also offers 60 texture units and a whopping 48 render output (ROP) units in its render back end. This incarnation of the Fermi GPU is built onto a board with 1.5GB of GDDR5 VRAM. The GPU core clock is 700MHz, while the shader units run at 1,401MHz. The GDDR5 memory frequency is 924MHz, suggesting that Nvidia's GDDR5 memory controller is less efficient than AMD's, which runs GDDR5 at 1,200MHz on the Radeon HD 5870.

All this graphics goodness comes with a double whammy: a high price and high power consumption. The Asus card typically costs $520–$550, somewhat higher than the expected $500. System power consumption at idle is 159W (versus 134W for a stock Radeon HD 5870) and a whopping 399W at full throttle (compared to 268W for a stock HD 5870).

We popped the Asus ENGTX480 into our standard Windows 7 Ultimate 64-bit test system, which is built on an Asus P6X58D Premium motherboard and an Intel Core i7-975 Extreme Edition. We ran six very different games, all at 1920x1200 and 4x AA, with graphics detail maxed out, plus the Unigine Heaven 2.0 benchmark, with tessellation set at extreme. We compared the GTX 480 results to three AMD cards: the dual-GPU Radeon HD 5970, a factory overclocked Radeon HD 5870, and the similarly priced Sapphire Radeon HD 5870 Eyefinity Edition (which offers 2GB of 1,200MHz GDDR5.)

While the Radeon HD 5970 won most of the benchmarks, bear in mind that it is a dual-GPU card, so it may have issues running some games in dual-GPU mode. The 5970 is also hard to find for $680. Most 5970 cards break the $700 mark. Even then, the ENGTX480 card won a couple of benchmarks. If you compare just single-GPU cards, the ENGTX 480 won every contest except Crysis, which was a dead heat.

Despite the added power consumption and higher cost, the ENGTX480 impressed us with its overall performance. It's very much a forward-looking card—performance in recent games is excellent, indeed. It is pricey, though, at up to $150 more than a stock Radeon HD 5870. So while performance is excellent, the price differential is pretty substantial. Bear this in mind when deciding between AMD and Nvidia. **—LOYD CASE**

BENCHMARKS

	Asus ENGTX480	HIS Radeon HD 5970	XFX Radeon HD 5870 XXX	Sapphire HD 5870 Eyefinity
Unigine Heaven 2.0 (fps)	**26**	21	13	17
Battle Forge (fps)	61	**73**	49	47
Dirt 2 (fps)	80	**89**	71	71
Far Cry 2 / Long (fps)	103	**114**	78	75
Far Cry 2 / Action (fps)	**76**	75	65	63
Tom Clancy's HAWX (fps)	104	**128**	92	89
Crysis (fps)	31	**44**	33	32
STALKER: Call of Pripyat (fps)	39	54	38	37

Best scores are bolded. Our zero point uses a Core i7-975 Extreme Edition, an Asus P6X58D Premium, 6GB of DDR3/1333, 64-bit Windows 7 Ultimate, and a Corsair TX850 PSU. All games tested at 1920x1200 with 4x AA.

VERDICT 9

ASUS ENGTX480

+ PLANTS	− ZOMBIES
Excellent performance, particularly with DX11 titles.	Very expensive; sucks vast amounts of power.

$550, www.asus.com

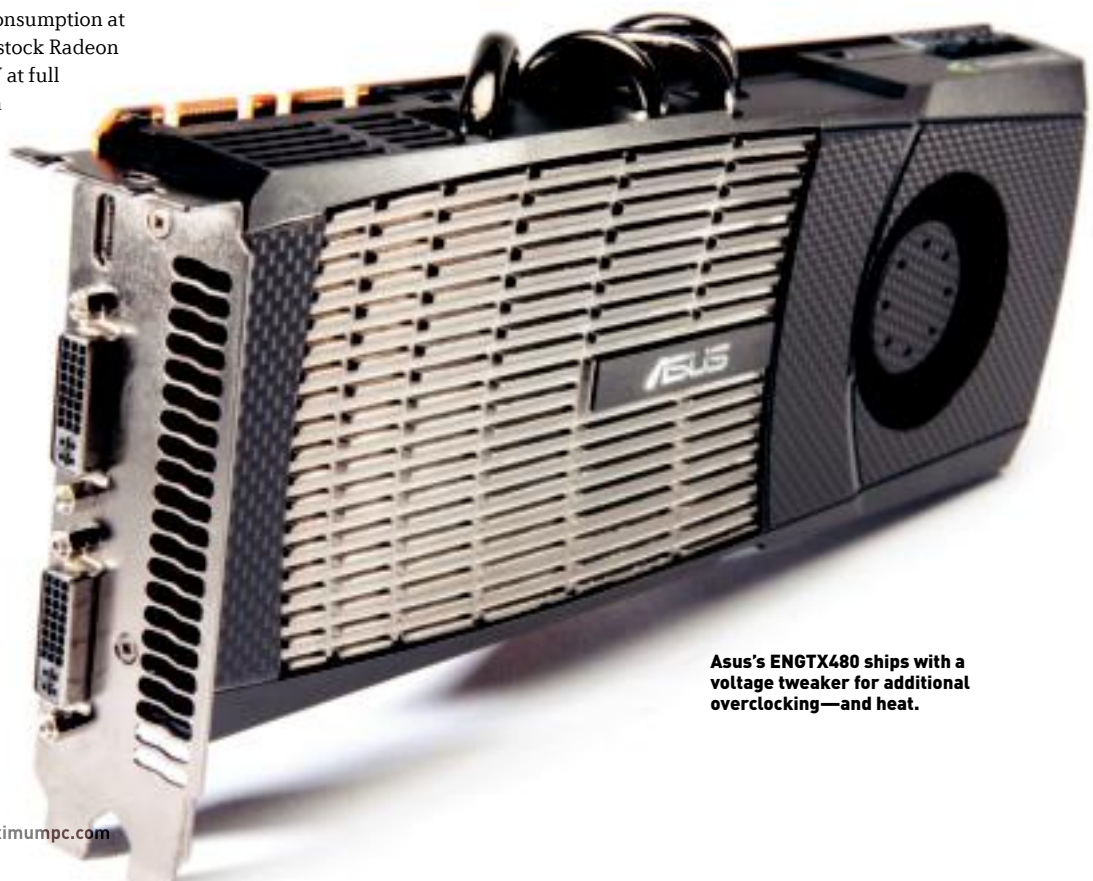

Asus's ENGTX480 ships with a voltage tweaker for additional overclocking—and heat.

Linksys E3000 Wi-Fi Router

Return of the flying saucers

Cisco's new Linksys E-series routers look a lot like the WRT-series routers they replace, complete with the weird flying-saucer motif and internal antennas. All the new features are under the hood and in the setup software.

As befits a flagship product, the E3000 is a dual-band (2.4GHz and 5GHz) router that enables you to operate two discrete wireless networks simultaneously. You can also operate a virtual guest network on the 2.4GHz band that limits clients to Internet access, isolating them from the rest of your network.

We typically ignore a router's setup software, but we're of two minds when it comes to Cisco Connect. It's the most foolproof tool we've ever used to set up a router, but it's very fragile unless you do things exactly Cisco's way. The utility assigns the router a unique, easy-to-remember SSID (e.g., RubyPenguin) and a complex password. Connecting additional PC clients to the network is a simple matter of using Cisco Connect to create a USB key and then plugging the key into each client.

You can use Cisco Connect to change the router's SSID or password, but doing so from within the web interface breaks Cisco Connect because the software doesn't have this new information. That wouldn't be such a big deal if the web interface exposed all the router's functionality, but it doesn't. If you want to configure the guest network or tweak the router's crappy parental controls, Cisco Connect is the only means. If you want to forward a port or accomplish other complex tasks, the web interface is the only means. Changing the router's SSID and password to their original values restores Cisco Connect, but we'd prefer to use one user interface for everything.

In our tests, the E3000 delivered much less range than our current Best of the Best pick, the Netgear WNDR3700; and it delivered about 15 percent lower TCP throughput at close range. At most of our other test locations, however, the E3000's TCP throughput was very close to that of the WNDR3700. The Linksys E3000 is a good router, but we think Netgear has a better product. —MICHAEL BROWN

BENCHMARKS

	Linksys E3000	Netgear WNDR3700
Kitchen, 20 feet (Mb/s)	68.8	**83.0**
Enclosed Patio, 38 feet (Mb/s)	45.7	**46.9**
Media Room, 35 feet (Mb/s)	29.8	**30.9**
Bedroom, 60 feet (Mb/s)	43.9	**45.7**
Outdoors, 85 feet (Mb/s)	0.8	**2.7**

Best scores are bolded. TCP throughput measured using IPerf. N/C indicates no connection at that location. Read more about our testing methodology at http://bit.ly/16w270.

LINKSYS E3000

VERDICT 8

+ A SAUCERFUL OF SECRETS	− A SAUCER FULL OF SPIT
Good performance overall; USB port for storage; can limit the number of clients on the guest network.	Two configuration tools; can't share a printer over the network; no BitTorrent client.

$150 street, www.linksys.com

The Linksys E3000 has a four-port Gigabit Ethernet switch onboard, but only one USB port, which is limited to hosting a mass storage device—you can't use it to share a printer.

Corsair 700D

Our favorite full-tower case gets a little more affordable

If you've been paying attention at all to case reviews lately, Corsair's 700D should look familiar. That's because it's a slightly stripped-down version of the 800D, Corsair's debut chassis (reviewed March 2010). We awarded the 800D 9 out of 10 points and a Kick Ass award, lauding its roominess, features, and design. The 700D only differs from the 800D in two respects: Its side panel has no window, and the 800D's hot-swap SATA bays have been replaced with four HDD trays.

Like its predecessor, the 700D is huge—24 inches high, 24 inches deep, and 9 inches wide—and painted in a matte powder-black inside and out, except for the brushed-aluminum faceplate. It has five toolless 5.25-inch bays and six hard drive bays with slide-out trays, which can accommodate 3.5-inch hard drives without the use of screws, or 2.5-inch drives with screws. Two of the hard drive bays are in the case's lowest compartment. The 800D had two 3.5-inch bays there too, but they were less accessible and did not feature slide-out trays. The remaining four hard drive bays take the place of the 800D's hot-swap bays.

The 700D is divided into thermal zones, just like the 800D. The lowest part of the case holds the power supply and filtered intake ports (with room to mount fans). The power supply vents directly out of the case, so hot air from the PSU doesn't go to the rest of the case. Air from the rest of the case's lower compartment is drawn into the motherboard compartment via a 14cm fan and exhausted via a 14cm rear exhaust fan. A third 14cm fan draws air up from the base of the chassis to cool the hard drives. The air from the hard drive compartment exhausts behind the motherboard tray and exits through the rear panel of the case. The top panel of the case includes mounting brackets for three 12cm fans or a triple-fan water-cooling radiator, though no top exhaust fans are included.

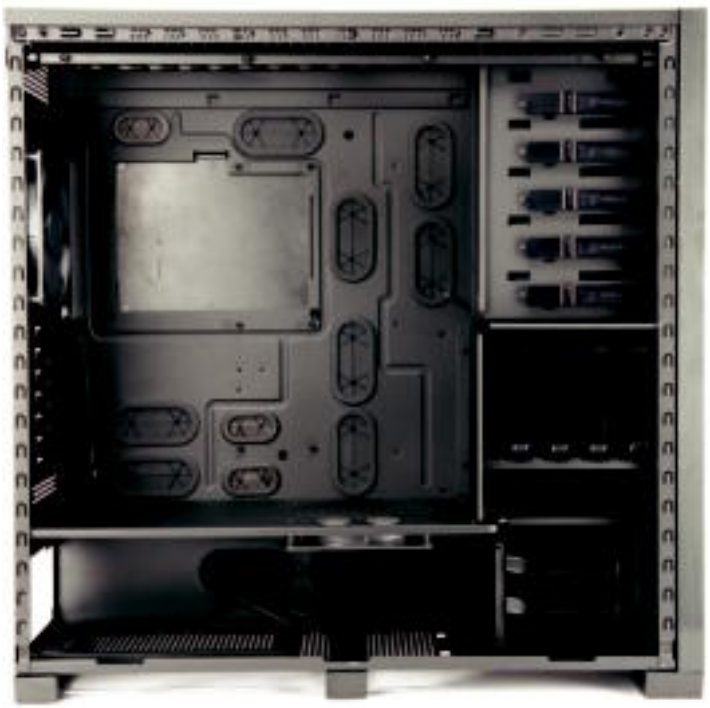

The 700D eschews hot-swap bays in favor of sliding hard drive trays and a lower price point.

The Corsair 700D's front-panel connections (four USB, FireWire, audio) are hidden in the compartment next to the power button.

The 700D supports micro-ATX, ATX, and EATX motherboards, and makes the most cluttered configs look spartan with its generous interior dimensions. It'll easily accommodate the longest videocards on the market with room to spare.

Corsair's commitment to well-managed cables continues to shine through—with so many rubber-grommeted cable-routing holes in the motherboard tray, and such a roomy interior, it's actually harder to achieve a bad-looking wiring job than it is to achieve a good one. That's one reason we kinda miss the 800D's side window—what's the point of a pretty build if you can't ogle it? But getting rid of the window and the hot-swap SATA bay helped Corsair knock $50 from the 700D's price, so we aren't complaining too much.

The one awkward thing about the 700D is its airflow. By default, the case only intakes air from the bottom, and for that reason the bottom of the case is raised an inch above the ground. If your computer room has thick carpet, you might need to keep your rig on a desk. And who has room on their desk for a monster like the 700D? Other than that, the 700D is a hell of a chassis, succeeding where the 800D does, at a lower price. If you can do without the hot-swap ports and the side window, and want to save $50 on a great case, here ya go. **—NATHAN EDWARDS**

MAXIMUM PC KICK ASS!

CORSAIR 700D

VERDICT 9

+ 7-UP	− FAYGO
A kick-ass case at a lower price.	Huge; we miss the side window.

$220, www.corsair.com

Acer H5360 3D Video Projector

Just slightly ahead of its time

3D Vision is one ace up Nvidia's GPU sleeve that AMD doesn't have an answer for. And if you enjoy 3D Vision on the small screen, you'll really dig it writ large with Acer's H5360 DLP video projector—provided you can accept the shortcomings inherent to the projector, Nvidia's GPUs, and 3D Vision in general.

When it comes to putting an Nvidia GPU in your home-theater PC, the biggest drawback is that none of the green team's cards provide the protected audio path needed to route encrypted HD audio to the card's HDMI output. So if you want to hear the Dolby TrueHD and DTS-HD Master Audio soundtracks on Blu-ray movies—and you do—you'll need to invest in a specialized soundcard, such as Asus's Xonar HDAV (read our review at http://bit.ly/2U3HAB) or Auzentech's X-Fi Home Theater HD (read our review at http://bit.ly/bAL65g).

The projector itself has two weaknesses: Most significantly, its native resolution is 1280x720, versus 1920x1080. That matters most when you're watching Blu-ray movies, because they look so much better splashed across a wall-size screen at 1080p. On the other hand, the limited resolution becomes something of a perverse benefit when it comes to gaming, for reasons we'll explain in a moment.

The projector's second deficiency has to do with the fact that it's a single-chip DLP model: We'd heard about the DLP rainbow effect, but this is the first time we've actually experienced it. DLP projectors produce color by spinning a wheel with red, green, and blue segments in front of the light source (this one spins at 9,000rpm). The projector displays each alternating color of the frame at the same location for the entire length of the frame. If your eyes move quickly across the entire screen—as happens frequently during gameplay—the three colors might appear to separate into a rainbow. This doesn't happen to everyone, and it didn't happen to us while we were watching movies; your mileage may vary.

Now let's get back to the resolution issue: When the GPU is rendering a game in 3D Vision, it produces two versions of each frame—one intended for the right eye and one intended for the left—so your net frame rate for a given game will generally be half what it would be while running in 2D mode. For the sake of convenience, we connected the projector to an AVADirect X1800 desktop-replacement notebook PC, which is outfitted with two GeForce GTX 285M GPUs in SLI. With 3D Vision enabled, that rig was able to produce Batman: Arkham Asylum at a respectable 45 frames per second. Had the projector enabled us to run the game at 1920x1080, we likely wouldn't have been satisfied with the results.

Nvidia is doing a good job of evangelizing 3D Vision to game developers, and game developers are doing an increasingly strong job of incorporating the technology into their code. What's more, the effects are much more powerful on the big screen than they are on a small monitor. Be that as it may, the images that drew the most oohs and aahs from the crowd that assembled for show time in the Lab were static scenes from Nvidia's demo slide show. Early-adopter gamers won't be disappointed with this projector, but we'll hold onto our wallets until Hollywood releases some Blu-ray 3D movies. —MICHAEL BROWN

VERDICT 7

ACER H5360 3D VIDEO PROJECTOR

+ STEREOVISION	− STRABISMUS
Bright; compact and portable; minimal price premium over conventional projectors.	Limited to 720p; uses DLP technology; needs a fast Nvidia GPU for gaming.

$700, www.acer.com

To take full advantage of Acer's H5360 120Hz DLP projector, you'll also need an Nvidia-powered videocard ($200 plus), Nvidia's 3D Vision kit ($200), additional pairs of glasses ($150 each), a projection screen ($100 plus), and an HDMI soundcard ($150 plus).

Samsung N210

It's OK to judge this book by its cover

The N210 is pretty on the outside but bloated on the inside. Be prepared to uninstall some nagware.

By now, if you're buying a netbook, you know what you're getting: All the models of a given generation are the same on the inside. So with the internals out of the decision tree, how do you choose which of dozens of near-identical netbooks is worthy of your purchase? Sure, the old standby differentiator of battery life still applies. But how about aesthetics? Can you actually choose a netbook based on design?

We think so. The Samsung N210's internals could be those of any current-gen non-Ion netbook—a 1.66GHz Atom N450 Pine Trail processor, 1GB RAM, a 250GB 5,400rpm hard drive—but it's what's on the outside that counts. The device has an embossed cream-color lid covered with a clear plastic coating. The interior is all matte white; and with its chrome edge trim and crisp gray lettering, it's almost retro-futuristic. The keyboard puts every other netbook keyboard to shame—the chiclet-style keys aren't cramped at all and the keyboard doesn't feel mushy. We could type on it all day. The track pad's multitouch capabilities help make up for its small size, and the LED-backlit screen is readable even at low brightness levels. Cranked up, the backlighting is quite bright for an office environment.

Beyond the excellent aesthetics of the case, we have a few complaints. Though the RAM is easily accessible, the hard drive is not, and the N210 has 802.11b/g/n Wi-Fi, but no Bluetooth.

But we're more perturbed about the N210's software loadout. The Windows 7 Starter desktop contains more than a dozen icons for software trials, nagware, and bloatware. Useful shortcuts, like the one that turns on the N210's sleep-and-charge USB port, are lost in the clutter. Even more annoying: The installer for the FailSafe trial (a LoJack-like theft-prevention tool) kept auto-launching when we were trying to run benchmark tests and wouldn't stop until we uninstalled the application.

One useful feature included with the N210 is the Linux-based Hyperspace instant-on OS. Though "instant-on" may be a misnomer, the OS does boot to a usable state in just 20 seconds (the Win7 boot, by contrast, takes more than a minute). Hyperspace offers an office suite, Twitter and Gmail integration, web browsing, and a list of app buttons, though some of them are misleading—the Photoshop button leads to Photoshop.com, and the Bloggers button, which uses the icon for Google's Blogger, links to the unrelated Bloggers.com. Despite these quirks, Hyperspace is useful, though we didn't find ourselves using it much.

The N210's performance in our benchmarks was almost indistinguishable from the other Pine Trail netbooks we've tested, like the Acer Aspire One we reviewed in April and the Toshiba NB305 we reviewed in May. The only benchmark where the N210 lagged behind was battery life—it lasted just five hours on our video-rundown test, while both the Acer and Toshiba netbooks topped seven hours on similar 6-cell batteries.

We expect any *Maximum PC* reader to be able to strip out the bloatware and nagware that comes with a new PC (and may we suggest PC Decrapifier to help with that?). But we don't appreciate Samsung's cluttering up the OS install. On the one hand, the N210's design is fantastic and its performance is exactly what we'd expect. And we appreciate the (albeit quirky) Hyperspace instant-on OS. But the less-than-stellar battery life and bloated Win7 Starter desktop lose the N210 some points in our book. —NATHAN EDWARDS

SPECIFICATIONS

Processor	1.66GHz Intel Atom N450
Chipset	Intel NM10 Express
Graphics	Intel GMA 3150
Display	10.1-inch LED-backlit matte LCD@1024x600
RAM	1GB DDR2/667
Storage	250GB HDD (Seagate Momentus 5,400rpm)
Ports	Three USB 2.0, audio in/out, SD/MMC reader, VGA, 10/100 Ethernet
Wireless	802.11b/g/n
Lap/Carry	2 lb, 14 oz / 3lb, 8 oz

BENCHMARKS

	ZERO POINT	
Premiere Pro CS3 (sec)	708	675
Main Concept (min)	251	243
Quake 3 (fps)	60.9	56.2 (-7.7%)
Quake 4 (fps)	3.6	4.1
Battery Life (mins)	255	300

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

Our zero-point netbook is a Lenovo IdeaPad S12 with a 1.6GHz Intel Atom N270, 1GB of DDR2/667 RAM, a 160GB hard drive, Intel GMA950 integrated graphics chipset, and Windows XP Home SP3.

VERDICT 7

SAMSUNG N210

+ SAM'S SONG	− SON OF SAM
Great design; bright screen; good touch pad.	Bloated Win7 Starter install; no Bluetooth; middling battery life.

$380, www.samsung.com/us

CA Internet Security Suite Plus 2010

This AV app takes aim at inexperienced users but misses the mark

When we ran our annual antivirus roundup in the May 2010 issue, many of you wrote in asking why we didn't include Product X or Product Y. Fair question, so here's the deal: We could have filled an entire issue reviewing just AV products, but that would have grown old by about page 32. Rather than do that, we're devoting space each month to cover apps that didn't make the cut, and CA Internet Security Suite is first up to bat.

After we installed CA ISS, it quickly became apparent that power users are not the target demographic. CA took a wrecking ball to last year's version and completely redesigned the UI in an attempt to "eliminate the technobabble that makes PC security difficult to understand and control," but in doing so, it made it needlessly tedious to poke around under the hood. The main interface consists of four index card–shaped menus that you can cycle through like a tie rack. Sounds easy enough, but if you want to set up a scan schedule, for example, you'll need to bring up the My Computer card, click the Update Settings link, highlight the Threat Settings tab, and then scroll to the bottom. You'll fumble around like this until you get accustomed to the interface, and when you do, you'll discover there's not a whole lot to play with. Strike one.

Strike two came when we tested for performance. CA ISS added 30 seconds to our test bed's startup time, trailing only Trend Micro (32 seconds) as the worst offender we've analyzed all year. When we fired up PCMark, CA ISS came in dead last and was the only security suite to turn in a score lower than 5,000.

But our biggest gripe is that CA ISS leaves your system vulnerable to malware, though not because the scan engine is weak. By default, a couple of critical modules are left disabled, including Program Protection, which prevents unknown programs from spawning other programs when executed, and Code Injection Protection, which prevents programs from gaining access to system privileges. Left turned off, our collection of malware was allowed to run amok on our test bed, and that's inexcusable for a security suite that touts ease of use over everything else. Strike three, batter's out.

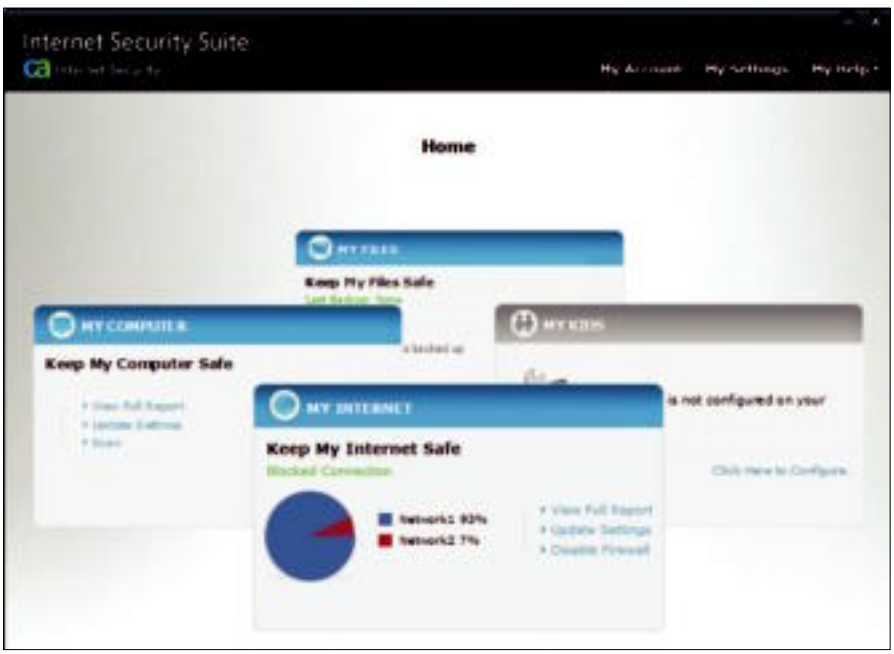

Navigating CA's seemingly simple menu system is harder than it looks.

There were a few redeeming qualities that leave us hopeful for next year's version, such as CA ISS's stellar subsequent scan times. As is becoming standard practice in the AV industry, CA ISS skips over files that haven't changed since the last time they were scanned, and this resulted in a subsequent sweep of our system in just three minutes and 19 seconds. We also dig the Link Advisor tool, which tells us if a URL is a gateway to a dangerous site before clicking on it. But these are just a couple of diamonds, the rest is pretty rough. —PAUL LILLY

If you see this screen (as we did after CA ISS failed to protect our test bed), then something has gone terribly wrong. The real Windows Security Center will never direct you to a paid security suite.

DARE TO COMPARE: PERFORMANCE

	CA ISS	Norton	ESET	MSE	McAfee
Scan 1 (min:sec)	12:09	16:18	**7:45**	16:56	13:33
Scan 2 (min:sec)	**3:19**	4:47	7:43	16:56	6:45
PCMark	4,833	5,760	**6,067**	5,622	5,645
Boot (seconds added)	+30	+18	+12	**+9**	+13

Best scores are bolded. Our test bed is a Core 2 Quad Q9400, 8GB DDR2/800, a Seagate Barracuda 320GB 7200.10 (~60GB filled across two partitions), a Radeon HD 3650, and Windows 7 Professional 64-bit. The reviewed app is compared to the top-performing apps from our AV showdown in the May 2010 issue (see http://bit.ly/cB6sqN).

VERDICT **4**

CA INTERNET SECURITY SUITE PLUS

+ PRINGLES	− SHINGLES
Relatively fast scan times; effective link scanner.	Greatly slows down boot time; default protection levels are too low.

$70 (1 yr, 3 PCs), http://shop.ca.com

Metro 2033

Not quite the ride of your life

We liked Metro 2033. We really did. But we wanted to love it. Its dusty, downtrodden, nuked-to-oblivion vision of a post-apocalyptic future is a thing of perverse beauty. At once terrifying and unsettlingly believable, it threatened to suck us in like no game before it. "Half-Life 2, who?" we asked ourselves frequently during the game's opening moments—that is, when we weren't left completely breathless.

Then the game made the mistake of putting a gun in our hands.

At best, Metro's shooting is serviceable. The weapons—while compulsively upgradeable—are crafted in such a way as to be realistic, which in this case means "boring." That would be fine and dandy if the other two pillars of first-person-shooter fun—level design and enemy AI—did enough heavy lifting to make up for it. Sadly, they don't.

Levels, of course, are visually spectacular, but at their core are linear corridor crawls. Fittingly enough for a game called Metro, the entire thing is almost glaringly on rails. Now, this isn't necessarily a bad thing, so long as you stick to the beaten path. Think of the game as a roller coaster ride: If it stays on track, it's the ultimate thrill ride. If it goes off track, everyone dies. That's Metro 2033 in a nutshell.

Those flaws, however, are almost speck-like in their insignificance compared to the plank that is Metro's enemy AI. Sure, when it's powering hordes of monstrous Dark Ones, it does the job well enough, but human opponents are a different story. On multiple occasions, we were able to stand in plain sight with our silenced weapon and pick off enemy after enemy, and all the while our confused quarry mounted no noticeable offense. Other times, enemies ran right into us and carried on as though we were nothing but a particularly thick cloud of air.

This screenshot—believe it or not—barely does Metro's immersive qualities justice. Metro's sound design is absolutely fantastic, and plays a huge role in sucking you into the game world.

It's a shame, too, because the game's certainly not all about shooting. In fact, its most intense moments come creeping in when the action grinds to a halt. Metro makes excellent use of ambient sound, bringing the proverbial roller coaster to its absolute peak before the inevitable climax. Also a bit disappointing is the game's storyline, which is surprising considering that it was actually adapted from a novel by Russian author Dmitry Glukhovsky. While Metro's world design is almost uniformly spectacular, its plot relies too heavily on convenient coincidences and ultimately resolves with a twist gamers have seen hundreds of times before.

In spite of all that, though, we're still willing to recommend Metro 2033. The game's world design, attention to detail, and dread-inducing pace are second to none. Seriously. Rarely do such an uneven game's high points overshadow its glaring flaws, but Metro's an exception to the rule. The game may have an unfortunate knack for killing its own buzz, but even that's not quite enough to ruin the overall experience. That's what Metro is, really: an experience. And a damn good one at that. Its shooting—while admittedly lacking—is only a part of that whole, and Metro 2033 is greater than the sum of its parts. **—NATHAN GRAYSON**

Actually fighting the creatures that go bump in the night is a bit of a letdown, thanks to shoddy AI and by-the-rails play mechanics.

METRO 2033

VERDICT 7

+ ON-RAILS	− OFF-TRACK
Breathtaking atmosphere and world design; unique setting.	Mediocre-to-awful enemy AI; sometimes stifling linearity.

$50, www.metro2033game.com, ESRP: M

LAB NOTES

Revisiting SSD Scores

A bug in our old test bed prompts a re-test of a drive reviewed in May

NATHAN EDWARDS
SENIOR ASSOCIATE EDITOR

During the testing for last month's SSD roundup, I identified a problem with my SSD test bed: Some scores, especially read speeds, were proportionally slower on that test bed than on any other system. Fortunately, only one SSD was reviewed on the new test bed—the Kingston SSDNow V+ (May 2010)—before I caught the bug. We gave that drive an 8 verdict, citing its TRIM support and strong performance. When we re-tested the SSDNow V+ on the test bed we used for last month's feature, we found read speeds of around 200MB/s and writes around 167MB/s, rather than 180MB/s and 175MB/s, respectively.

Upon retesting the Kingston SSDNow V+ under the same conditions as last month's SSD roundup, we decided to keep the 8 verdict but update the benchmarks. The updated benchmarks (including three of last month's drives for reference) appear on the right.

BENCHMARKS

	Kingston SSDNow V+	Patriot Torqx	OCZ Vertex LE	Corsair Nova V128
Controller	Toshiba	Indilinx	SandForce	Indilinx
Capacity	128GB	128GB	100GB	128GB
HDTune Avg. Sustained Reads (MB/s)	202.2	**220**	197.5	210.8
HDTune Avg. Sustained Writes (MB/s)	166.9	162.3	**223.5**	163.9
HDTune 4KB Random Read (IOPS)	4,630	**7,084**	5,050	7,439
HDTune 4KB Random Write (IOPS)	3,700	3,435	**5,271**	2,829
HDTune Random Access Writes (ms)	0.6	**0.2**	0.3	**0.2**
Premiere Pro CS3 (sec)	**356**	364	381	361
PCMark Vantage x64 HDD	20,105	23,674	**35,655**	24,796

Best scores are bolded. All tests performed on an Asus P6X58D Premium motherboard with a Core i7-X980 CPU @3.33GHz with 6GB DDR3/1600 running Windows 7 Professional 64-bit.

GORDON MAH UNG
SENIOR EDITOR

I've never been a fan of voice chat for myself (it's fine for others in the party) in online gaming, but Dolby's new Axon audio enhancer has me interested. Instead of sounding like the booming voice of God, players sound positionally accurate with Axon. This has the potential to greatly enhance the immersiveness of gameplay.

ALEX CASTLE
ONLINE MANAGING EDITOR

This month, I got to try out a couple of new levels from the StarCraft 2: Wings of Liberty single-player campaign, and it made me even more excited about the game than before. Blizzard is promising a game where no two missions are alike, and I really think that it's going to deliver.

ALAN FACKLER
ONLINE ASSISTANT EDITOR

After spending tons of time figuring out Adobe's new Photoshop CS5, I now see the world through different eyes. New features like Content Aware Fill and Puppet Warp will make it harder for even the most die-hard photo enthusiasts to gauge what's real and what's been altered with Photoshop's refined magical tool box. Adobe, it seems, is poised to do it again. Again.

MICHAEL BROWN
REVIEWS EDITOR

An audiologist fitted me for a set of custom in-ear monitors from JH Audio for an upcoming review. She warned it might hurt, but I shrugged it off because my pain tolerance is fairly high. *Ay, caramba!* I almost cried when she shoved the impression material into my ear canal. Those monitors better rock my world.

KATHERINE STEVENSON
DEPUTY EDITOR

First, Alienware managed to pack real gaming into an 11-inch notebook, now HP is boasting a reasonable-size notebook with support for AMD's Eyefinity. Equipped with VGA, HDMI, and DisplayPort ports, and a Radeon HD 5850 graphics card, the new 17-inch Envy will let you run three 1080p monitors simultaneously. And the notebook remains portable at 7.5 pounds.

We tackle tough reader questions on...

LCD TV vs. Monitor
UAC
Maximum PC iPad Subscription?

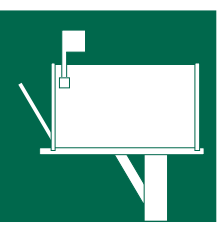

Dept. of Display Justifications

I'm looking to upgrade from my 24-inch Dell LCD monitor. My choice is between a 30-inch LCD monitor from LG that goes for a little more than $1,100 and a 32-inch Samsung LCD TV that costs $600. I understand that the monitor offers higher than 1080p resolution. But the TV is cheaper. What other advantages and disadvantages should I consider? How do I justify the more expensive monitor? Will the TV burn my eyeballs out like mom always said it would?

—Robert Burnham

Reviews Editor Michael Brown Responds: If you're looking for a display to connect to your PC, and you'll be using your PC for productivity apps (email, word processing, photo editing, and so on) as well as gaming and watching movies, you'll be happier if you keep screen size and resolution in proportion. You don't mention exactly which models you're considering, so we can make only generalized recommendations here. A 32-inch television with 1920x1080 resolution will look just fine when you're watching it from across the room; but when it's sitting right in front of your face, you'll be driven to distraction by the gaps between the pixels. And when it comes to rendering color, monitors tend to be both more consistent and more accurate than TVs. Consider the more expensive display justified.

In Favor of UAC

I have to disagree with your recommendation to turn down UAC in Windows 7 in your June cover story. In fact, it's better to turn it up to max so that it also requires Windows components to get permission. If you don't do that, malicious software can use parts of Windows to do whatever it wants—basically rendering UAC worthless.

I've never understood the objection to UAC. Even turned to max, getting a prompt should be very rare. You're not going to get any in normal day-to-day use. I've also heard people claim that power users don't need it, which is actually backward. "We" know if something we're doing should require admin privileges, so it's actually better protection for us than it is for a typical user who just randomly clicks OK to every prompt they get. I love that NT 6 finally lets us basically run as normal users except when we really need admin! (All without actually having to create a normal user account… and come on, how many of us ACTUALLY do that?)

—Andrew

Senior Editor Gordon Mah Ung Responds: This is really a philosophical discussion as there is probably no right answer for someone who is security conscious, keeps his or her OS, apps, and plugins updated, and runs high-quality AV software. But I'm actually with you, Andrew: I leave UAC turned on (although not to maximum) and don't mind the messages. But I am paranoid. I actually know many power users who prefer to hang loose. Some run with UAC turned down *and* without antivirus software installed because they know, *just know* how to avoid getting infected. If you ask me, that's simply foolish given the sophistication of the exploits today.

Long-Term SSD Testing?

I have been doing my research in preparation for the switch to a solid state

CUTCOPY**PASTE**

- In our review of the Netgear WNDR3700 Wi-Fi router (May 2010), we neglected to report that the parental controls can be customized.
- In our antivirus roundup (May 2010), we mistakenly listed BitDefender as not including rootkit detection. It does. We also neglected to include the performance numbers for Avast. Those numbers can be found in our updated chart online: http://bit.ly/cXwjNj.

NOW ONLINE

Five Apps to 3D-ify Your Desktop

Google recently bought up Bumptop, the most widely known 3D desktop software project, and is ending its free beta. But that doesn't mean that there aren't still plenty of ways to give your desktop some 3D flair for free. In this week's freeware files, we round up five of the best apps for desktop 3D-ification: http://bit.ly/bI3WbD.

drive. While reading tons of user reviews, I have seen many people complaining about their drives dying or suffering performance degradation over time. I was wondering if *Maximum PC* has done any long-term testing of SSDs? Six months of reads and writes and then re-running your benchmarks?

—Steve

Senior Associate Editor Nathan Edwards Responds: There's really no way to do long-term testing of SSDs other than to use them in the real world on a day-to-day basis. Due to the way NAND flash works—which you can read about in our white papers on SSDs (http://bit.ly/19I0GC) and TRIM (http://bit.ly/7GhVfo)—using an SSD without the benefit of the TRIM command and/or a garbage-collection program *will* result in slowdowns after all flash blocks are written to at least once. That's why we recommend buying an SSD with TRIM support (if your OS supports it) or regularly running a garbage-collection utility if your SSD and OS don't support TRIM—each drive manufacturer has its own. In short, TRIM allows constant garbage-collection on the OS level, enabling Windows to erase deleted blocks as soon as the data is deleted, rather than waiting until the next time the deleted block is written to.

As for drives dying out-right, though reports in early 2008 pegged failure rates of first-gen drives at close to 30 percent—a claim that was disputed by SSD manufacturers—current-gen SSDs are at least as reliable as mechanical hard drives, and much more durable.

Build Your Own Notebook?

After reading the issue where you listed three different-priced PCs (September 2009), I decided to build my own. But now I'm wondering if there has been an issue where you built a notebook? Is building one more complicated than a desktop?

—Pablo Sanchez

Senior Editor Gordon Mah Ung Responds: We actually did a build-a-laptop story in April 2006. Believe it or not, at the time, Intel was even pushing the idea of DIY notebooks, and many vendors were also pushing notebook kits. However, the initial excitement soon fizzled as notebook makers realized that the delicate assembly of putting together a notebook and the warranty problems for blown builds simply wasn't worth it. The process isn't any more complicated than building a desktop, but the scale is much smaller and, frankly, the flexibility isn't quite there either—you can't, for example, put any GPU you want in a notebook. You can still find some bare notebook kits on the market, but they're certainly not being heavily promoted.

Pad Me, Bro

Long, long-time print subscriber. I love reading the archive PDFs on my iPad. Any chance you'll do an iPad-specific subscription plan so we can get the current issue right away on the iPad?

—Tom Bush

Editorial Director Jon Phillips Responds: First off, thanks for the long-time support. As for your request, we're not convinced that enough *Maximum PC* readers own (or plan to buy) iPads to warrant an iPad-specific version. The iPad is a great tablet, but I'm not convinced it's going to dominate the tablet universe (as the iPhone dominates smartphones), let alone ebook devices.

I'm a firm believer that for long-form, leisurely reading, folks want their content on either printed paper or electronic paper (à la the Kindle). Backlit LCD screens are just too harsh on the eyes for extended leisure-time reading. And my opinion seems to be supported by the fact that no magazine publishers are making much money from digital subscriptions—people just aren't interested in buying them. This could change when the wide-format digital reader with color e-ink hits the market.

Tell you what: Check back with us in 2011. By then, the tablet universe should have expanded to the point where print and web publishers have a clearer idea of what users really want from these devices.

I HAVE TO DISAGREE WITH YOUR RECOMMENDATION TO TURN DOWN UAC IN WINDOWS 7

LETTERS POLICY Please send your questions and comments to comments@maximumpc.com. Include your full name, city of residence, and phone number with your correspondence. Letters may be edited for space and clarity. Due to the amount of mail we receive, we are unable to respond personally to all queries.

Unleash your PC's Potential...

Try MAXIMUM PC

2 FREE Trial Issues

Each issue of Maximum PC features:

- **Brutally honest product reviews**
- **Hard-hitting editorials**
- **Tips to blast your machine's performance**
- **Insightful and innovative How-To's**
- **A CD loaded with new software, utility and game demos**

Reserve your 2 FREE Trial Issues today! There's no obligation.

To order, head to:
www.maximumpc.com/archive

BUDGET PROCESS

AMD Phenom II X6 T1055

Our pick for the budget CPU was not an easy one to make. We were downright split on the decision, but we ultimately decided on AMD's phenomenally great 2.8GHz Phenom II X6 chip. In most of today's games and apps, the $200 Intel Core i5-750 still eclipses AMD's hexa-core parts. However, when you get to apps that exploit more than four threads, the 2.8GHz Phenom II X6 takes the lead. Since we think the future will see increased thread loads, the Phenom II X6 will actually get faster as games and applications are updated. And hell, it's a hexa-core CPU for $200!
www.amd.com

THE REST OF THE BEST

High-End Processor
Intel 3.33GHz Core i7-980X
www.intel.com

Midrange Processor
Intel 2.8GHz Core i7-860
www.intel.com

Budget Processor
AMD Phenom II X6 T1055
www.amd.com

LGA1366 Motherboard
MSI Eclipse SLI
www.msi.com

LGA1156 Motherboard
Asus Maximus III Formula
www.asus.com

Socket AM2 Motherboard
MSI K9A2 Platinum
www.msi.com

High-End Videocard
ATI Radeon HD 5970
www.ati.com

Midrange Videocard
ATI Radeon HD 5850
www.ati.com

Budget Videocard
ATI Radeon HD 5770
www.ati.com

Capacity Hard Drive
Western Digital Caviar Black 2TB
www.wdc.com

Performance Storage
Patriot Torqx 128GB SSD
www.patriotmem.com

DVD Burner
Samsung SH-S223
www.samsung.com

Blu-ray Drive
Plextor B940SA
www.plextor.com

Full-Tower Case
Corsair 800D
www.corsair.com

30-Inch Display
NEC LCD3090
www.nec.com

Gaming Mouse
Logitech G9x Laser Mouse
www.logitech.com

2.1 Speakers
Bowers & Wilkins MM-1
www.bowers-wilkins.com

Games we are playing

Just Cause 2
www.justcause.com

The Settlers 7: Paths to a Kingdom
http://thesettlers.us.ubi.com

Tom Clancy's Splinter Cell: Conviction
http://splintercell.us.ubi.com

Battlefield: Bad Company 2
www.badcompany2.EA.com

For even more Best of the Best entries, such as speakers and budget components, go to http://www.maximumpc.com/best-of-the-best.

MAXIMUM PC (ISSN 1522-4279) is published 13 times a year, monthly plus Holiday issue following December issue. 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 South San Francisco, CA and at additional mailing offices. Newsstand distribution is handled by Time Warner Retail. Basic subscription rates: one year (12 issues) US: $20; Canada: US$32; Foreign: US$44. Basic subscription rates including monthly CD, one year (12 issues/12 CD-ROMs) US: $30; Canada: US$42; Foreign: US$54. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R128220688). PMA #40612608. Subscriptions do not include newsstand only specials. POSTMASTER: Send changes of address to Maximum PC, PO Box 5159, Harlan, IA 51593-0659. Standard Mail enclosure in the following edition: None. Ride-Along enclosure in the following editions: B1, B2, B3, B4. Returns: Bleuchip International, PO Box 25542, London, ON N6C 6B2, Canada. Future US, Inc. also publishes Mac|Life, Nintendo Power, PC Gamer, The Official Xbox Magazine, PlayStation: The Official Magazine, World of Warcraft Official Magazine, NVISION, Guitar World, Revolver, Guitar Aficionado, Windows: The Official Magazine, MOM, Crochet Today! and Pregnancy.
CUSTOMER SERVICE: Maximum PC Customer Care, PO Box 5159, Harlan, IA 51593-0659. Phone: 1-800-274-3421. Email: maxcustserv@cdsfulfillment.com. Web: www.maximumpc.com/customerservice. Back issues can be purchased by calling 1-800-865-7240.
REPRINTS: Reprint Management Service. Phone: 717-399-1900 ext. 100.