DE

Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

Put Together a Virtual Cluster

Combine mulit-core workstations to create a virtual machine P. 16

Where Do I Go for *Data Acquisition* Products?
omegamation.com, *of Course!*
Your single source for process measurement and control products!

Economical Solution for Facilities with an Existing Wi-Fi Network

© Kutt Niinepuu / Dreamstime.com

**Wi-Fi *Wireless* Sensor System
802.11b/g Wireless Ethernet**

wireless

wSeries
Starts at
$195

MADE IN
USA

Alarms by
Email or
Text Message

Visit
omega.com/wseries

**Shop Online at omega.com®
CONTROLCAT™**
New Horizons® in Temperature, Pressure, Flow and More!
• Automation • Sanitary • Temperature Measurement • Heaters
• Wireless • Data Acquisition • Pressure Transducers • Flow Meters

HOT NEWS®

Version No. 28

Includes 105 Dilbert Cartoons

FREE!

For Sales and Service Call
1-800-TC-OMEGA® OMEGAETTE®

DILBERT

Visit **omega.com**
to order your **FREE** copy of
Dilbert™ CONTROLCAT™ New Horizons®
in Temperature, Pressure, Flow and More!
Version No. 28 Featuring 105 DILBERT Cartoons!

**Expandable Modular Data
Acquisition System**

iNET-400
Series
Complete
Starter
System
$990

Visit
omega.com/inet-400_series

**Temperature, Humidity and
Dew Point Data Logger with
LCD Display**

OM-EL-USB-2-LCD
$97

Visit
omega.com/om-el-usb-2-lcd

*UWTC Series
Starts at
$125

MADE IN
USA

*Temperature-to-Wireless
Connector/Converter System*
Visit *omega.com/uwtc*

**Temperature and Humidity
Virtual Chart Recorder
Web Server**

iTHX-SD Series
Starts at
$395

MADE IN
USA

Easy and Intuitive to
Operate, 3-Point
Calibration Option
Visit
omega.com/ithx-sd

Portable Data Logger
Visit *omega.com/om-sq2010*

OM-SQ2010
Starts at
$1540

Voltage Data Loggers, NOMAD®

OM-CP-VOLT101A
Starts at
$179

Visit
*omega.com/
om-cp-volt101a_series*

MADE IN
USA

For Sales and Service, Call TOLL FREE
1-888-82-66342®
1-888-TC-OMEGA

*PATENT PENDNG

VISA   MasterCard

OMEGA®

The HPC Results Are In

Recently we surveyed you on your interest and use of high-performance computer (HPC) systems in engineering applications. I would like to thank the 651 subscribers who took the time to fill out an extensive survey. You let us know how HPC is changing the way you work, where you turn to get the information you need to evaluate HPC, and what would justify your purchase of an HPC cluster. The results will help us provide you with the kind of content you are looking for, as well as to understand how important different channels of information are to making the decision that could cost you anywhere from a few thousand dollars to millions.

You Said It

Interestingly, 59% of respondents either have an HPC cluster, have evaluated purchasing one, or are in the process of evaluating the purchase of a cluster computer. Not surprising, most of the installed base of HPC clusters today are in mid-to-large size companies, while about half of small-to-mid sized companies are currently evaluating the purchase of their first HPC system.

> **It's not just about doing things faster, it's about doing them better.**

So, where do you go to find information that will help with the decision on acquiring an HPC system? Ninety-eight percent of those considering a purchase said that independent published information was very or somewhat important. Hardware vendors came in second at 93%. I found it interesting that only 15% used another engineering company as a source of information on the purchase of HPC systems.

Of those who have been involved in the purchase of an HPC cluster, performance was the key factor in the purchase, followed by support, with price lagging substantially. Just as important: In-person training from their CAE vendor would have helped respondents in their purchase decision process.

Tracking Trends

For four years now, *DE* has been surveying our subscribers on their use and views of high-performance computing. The trend of using these enabling technologies to perform analysis and simulations has increased dramatically as the costs have gone down and the performance has increased. While we did not directly ask about ease of use, clusters have become more accessible to the small- to mid-size engineering companies that need them most.

DE's recent coverage of how Parker Aerospace used a workstation cluster to speed up throughput on CFD simulations demonstrates that most engineering teams have the compute resources for some level of multi-core system. These systems can be implemented without a large IT department and used by any engineering team with the available cores in their workstations. With the advent of multi-core processors, a workstation can have 12 or more cores available.

Building a dedicated cluster is much easier than even a few years ago. Depending on the application, a reasonably scaled cluster can be achieved without the problems inherent to assembly and testing that were once common. See our continuing coverage of HPC clusters on page 16.

Performance ROI

So what is the next question? How much performance do I need? This is where the software vendor participation becomes crucial.
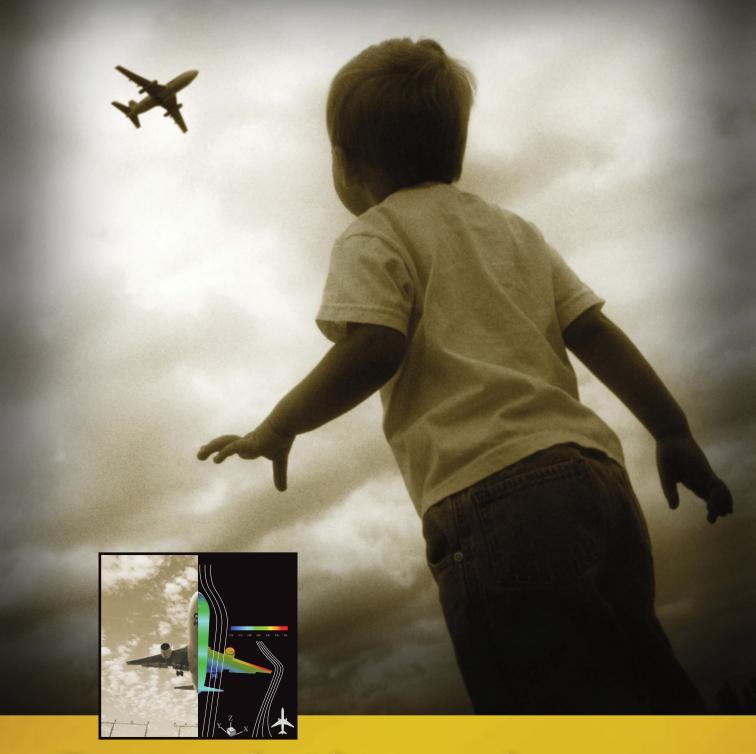
I know management and the financial guys find return on investment hard to quantify, but I've been at *DE* a long time. For the first five years of this magazine, the average engineer—those of you not working in automotive or aerospace—had to justify the purchase of very expensive software. Whether finite element analysis, computational fluid dynamics, or some other analysis application, you were questioned on why you needed it, what it would do that wasn't already being done, and how it would create a ROI. Those same questions will be asked of HPC. The answers are obvious: When justifying the purchase of a high-performance cluster, it's not just about doing things faster, it's about doing them better. Being able to run more simulations creates a better design.

HPC is here to stay. What usually happens, and this is what happened in the Parker Aerospace story, is once implemented, the first cluster shows enough benefits to justify another one. This is why 10% of our subscribers are already considering the purchase of more HPC systems. It might be a small cluster on the department level, or a remote one that will still be faster than waiting for large amounts of data to return from the main office.

Of course, not everyone needs a high-performance computing system, but for most engineers running simulation and analysis jobs on their workstation, I see one in your future.

I'm off to Supercomputing 2011 this month. I'll report back on what's new. **DE**

Steve Robbins *is the CEO of Level 5 Communications and executive editor of DE. Send comments about this subject to DE-Editors@deskeng.com.*

Every product is a promise

For all its sophisticated attributes, today's modern product is, at its core, a promise.

A promise that it will perform properly, not fail unexpectedly, and maybe even exceed the expectations of its designers and users. ANSYS helps power these promises with the most robust, accurate and flexible simulation platform available.

To help you see every possibility and keep every promise.

Realize Your Product Promise™

ANSYS®

COVER STORY

Put Together a Virtual Cluster

16 A virtual cluster can be built from your existing resources and dedicated to completing computationally intensive engineering tasks, such as simulation or rendering. With the right software, a virtual machine—one that doesn't physically exist but uses the untapped resources of one or more workstations—can be created. The more workstations that contribute resources to the virtual machine, the more powerful it becomes. Even a small cluster of four 12-core workstations or eight six-core workstations can create a virtual resource of 32 cores.

ON THE COVER: Multi-core processors come together to form a virtual machine. *Images courtesy of iStockphoto*.

A stator blade in the turbine stage of a jet engine is heated by the combustion gases.
To prevent the stator from melting, air is passed through a cooling duct in the blade.

Capture the Concept.

With COMSOL Multiphysics® you are empowered to build the simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. This multiphysics approach delivers results—tangible results that save precious development time and spark innovation.

Watch Tutorial

Joule Heating of a Circuit Board Fuse
www.comsol.com/showcase

DE
Desktop Engineering®
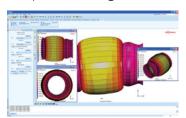
deskeng.com

ENGINEERING ON THE EDGE BLOG
See the future of engineering technology at engineeringontheedge.com. **NEW**

VIRTUAL DESKTOP BLOG
Read Kenneth Wong @ deskeng.com/virtual_desktop for a closer look at lifecycle components via articles, podcasts and video reports.

NEWSLETTER REGISTRATION
Newslink; Editor's Pick of the Week; Check It Out; Virtual Desktop; Focus on Analysis and Simulation; Focus on Engineering IT & Computing; Focus on MCAD; and Focus on Rapid Technologies.

DE MOBILE BETA
Check out our mobile app for Android and iOS by visiting wbxapp.com/de-mobile on your mobile device, or scan the QR code to the right.

***DE* EXCHANGE**
Join *DE'*s online community @ DEexchange.com

***DE* ON TWITTER**
Follow us @ DEeditor

LIKE *DE* ON FACEBOOK
Visit *DE* @ deskeng.com/facebook

BPA WORLDWIDE

We didn't reinvent the wheel...
We just thought it should be round.

Autodesk Subscribers
Get Cloud Club Membership for Free

If you're an Autodesk subscription customer, access to Autodesk Cloud—roughly a dozen cloud-hosted functions embedded in the company's design software, along with 3GB of storage space—is now yours.

Data storage is part of the recently launched Autodesk Cloud. Autodesk's strategy also includes the deployment of remote servers—equipped with far more processing power than what's typically available in desktops and laptops—to deliver certain software functions.

The list of initial web services available to Autodesk subscribers includes:

- **Autodesk Cloud rendering**, which allows you to render photo-realistic scenes and 3D images of products through remote hardware;
- **Autodesk Inventor optimization**, which allows you to experiment with design alternatives and select the best;
- **Autodesk Revit Conceptual Energy Analysis**, which allows you to analyze building energy consumption and energy-related costs;
- **Autodesk Green Building Studio**, which allows you to obtain environmental footprints of architectural designs and experiment with alternatives; and
- **Autodesk Buzzsaw software as a service (SaaS)**, which allows you to collaborate, manage data, and manage projects online.

Some functions, such as remote rendering and Inventor optimization, are available only to subscribers with Premium or Ultimate editions of Autodesk product suites. Autodesk plans to work with third-party infrastructure providers, including Amazon and Citrix, to operate a scalable cloud computing setup.

The initial round of web services are not available to Autodesk customers without a subscription. The plan may be the company's strategy to lure more people to sign up for its annual subscription, making it easier to justify the added cost with cloud-hosted benefits.

Critics often cite security and encryption—or a lack of them—as vulnerabilities of the emerging cloud-computing model; however, the use of virtual vaults is clearly gaining momentum among consumers and professionals alike.

The use of web services, which bypasses the need to install and master professional software packages, has led Autodesk to tiptoe into consumer-friendly offerings with applications such as Autodesk Homestyler software and Autodesk Seek. Autodesk business segments now include a Consumer Group, formed less than a year ago.

Products introduced via the Consumer Group showed "mixed success," but "something started to change with tablets," said Andrew Anagnost, Autodesk's vice president of suites and web services. Among Autodesk's tablet and smartphone apps, SketchBook Mobile stands apart as one of the most successful.

Some desktop clients that used to be a standard part of Autodesk products, such as Autodesk Design Review, may also be transitioning to become mobile apps. Design Review Mobile, released on September 25, is now available for iPhone, iPad, and iPod Touch at the Apple App Store. Design Review Mobile gives tablet and smartphone users access to drawings and files saved on Autodesk Cloud.

As Autodesk expands its web services, the company may have to consider new licensing schemes. For example: giving away a basic level of rendering rights for free, but charging subscribers additional fees for excessive demands. In addition, Autodesk must now take on additional responsibilities to ensure its cloud-hosted storage and services are reinforced with robust fail-over, replication, encryption, and data synchronization plans—something the company previously didn't need to worry about. **DE**

SolidWorks in Transition

Jon Hirschtick leaves board; Gian Paolo Bassi takes over as VP of R&D.

Gian Paolo Bassi takes over as executive vice president of R&D for SolidWorks, filling the role previously held by Austin O'Malley.

In October, as the company got ready to launch SolidWorks 2012, Jon Hirschtick, SolidWorks cofounder and former CEO, resigned from the board. He announced his decision to employees by email, followed by an update on Twitter: "Today [October 4] was my last day at SolidWorks and [Dassault Systèmes]. A tough decision for me, but one I feel is right," he wrote.

Hirschtick, an MIT graduate, was among the visionaries who believed in and pioneered Windows-based CAD for personal computers in the early '90s. After serving as CEO, he continued to influence the company's technology development, culture, and vision as a member of the executive board. For many loyal users, Hirschtick's down-to-earth, approachable character was the personification of SolidWorks.

Hirschtick's departure coincided with that of Austin O'Malley, one of the original developers of SolidWorks (dating back to version 95). As executive vice president of R&D, O'Malley shepherded SolidWorks through eight major product releases. The post vacated by O'Malley was filled by Gian Paolo Bassi, founder and CTO of RIWEBB. Bassi's past roles included vice president and CTO of ImpactXoft, director of think3, and program manager at Computervision.

Kristen Wilson, SolidWorks public relations manager, clarified that Hirschtick's resignation and O'Malley's departure were "completely unrelated." Wilson noted, "We were fortunate that [Hirschtick] stuck around for much longer than most founders typically do."

As the new captain charting the company's course, Bassi will have to negotiate between loyal users' proud legacy and proficiency in parametric CAD, and parent company Dassault Systèmes' desire to explore cloud-hosted data sharing, social media-powered collaboration, mobile apps, and other emerging trends. Responding to questions submitted via email, Bassi offered some clues to how he would reshape SolidWorks:

Q: What is your vision for SolidWorks for the near- and long term?

A: My vision is to provide a solution with more performance and modeling flexibility and to elevate the design process to be more conceptual. Models need have more freedom in the way they are designed, modified, and behave. I will bring hybrid methodologies beyond geometric parameters to functional design.

Q: What do you consider to be the most important technology trends reshaping the future of engineering and design software? Mobile computing? Cloud computing? The resurgence of Mac? Sustainability? Social media-driven collaboration?

A: I think the biggest trend is availability of massive computational and data transmission capabilities. Everything else is enabled by such technological advances. For example, we couldn't have cloud computing without scalable access to computing power and IBM's Watson [a supercomputer capable of processing questions posed in natural language] is also made possible by these capabilities.

Q: What is your personal philosophy on technology?

A: To think outside the box and stay open to any wonder. One needs to keep looking around without preconceptions. Your mental structure can limit you. **DE**
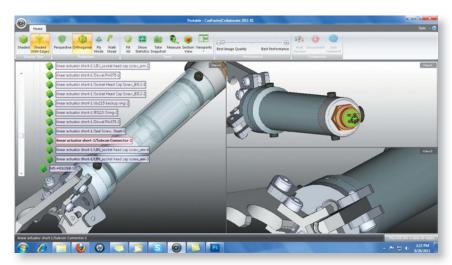
Bridge the Desktop and iPad with CadFaster Collaborate

You can add CadFaster, developed by a relative newcomer, into the mix of early iPad apps for CAD users. The software works as a plug-in to several 3D modelers, including SolidWorks. It is complemented by an iPad app, called CadFaster Collaborate.

Once the plug-in is installed, you may use the CadFaster tab in Solid-Works to convert a 3D assembly model into an executable file—a lightweight 3D file with a self-contained program for viewing and markup. You do not need a desktop client to view the exported file; a simple double-click will launch the file, along with its own interface for viewing and annotation.

You can also use CadFaster to upload the lightweight model to a cloud server, accessible from the iPad app. Initially, you get 100MB of storage space. If you're working with assemblies with hundreds of components, you will undoubtedly fill up the 100MB quota. CadFaster believes users will willingly pay a nominal fee to get additional storage.

Launching the executable file on a local desktop machine, you can view your model in cross-sections, study its internal components, add comments to it, or take measurements on it. You may also initiate a co-viewing session, which allows you and a collaborator to go online and look at the same model together in real time. With this method, you can easily invite someone to review and comment on a design from an iPad, in a collaboration session that facilitates both the desktop user and the tablet user. In a co-viewing session, the host takes control of the operations (so the host can, for example, set the transparency of a part to reveal its internal structure), but the guest sees the same view the host sees. At any given mo-

CadFaster's executable file works as a self-contained environment for viewing and annotating lightweight CAD files.

ment, the host can pass control over to the guest. (This method is a bit different from AutoCAD WS's co-viewing, where both parties may zoom and pan a shared DWG file, add comments, and use editing tools simultaneously.)

In CadFaster for iPad, you can do pretty much everything you can do with the executable file. Perhaps more important, CadFaster synchronizes comments between desktop and mobile apps, so if you add a comment on a shared document from your desktop, your collaborator who launches the same file on his or her iPad receives an alert, along with your comment.

CadFaster developers have discovered that, with certain antivirus software, the executable file might trigger unwarranted security alerts. With some overprotective antivirus software, the file may not launch at all. For those reasons, CadFaster is considering replacing the executable file with a thin client for the desktop.

The first release of CadFaster for iPad contains a few quirks in its interface. For

instance, the display doesn't reorient itself when you switch the iPad's orientation from portrait to landscape mode. The comments currently appear as a series of thumbnail images with texts. To make it easier for people juggling multiple discussion threads and models, CadFaster may need to give users the option to organize comments by associated models, by dates, by comment-author names, or some other means.

In its use of cloud-hosted data syncing, CadFaster is following the example of successful eBook reader apps, like the Amazon Kindle app. When reading an ebook, readers may use a mix of platforms (Kindle for PC, Kindle for iPad, Kindle for iPhone, or on the Kindle device itself). Kindle apps' ability to synchronize among themselves in the cloud ensures that readers see the same annotations and highlights, regardless of the device they're using to access their digital content. Similarly, CadFaster is bridging users' desktop operations and mobile operations, paving the way for collaboration across multiple devices. **DE**

Rock-Solid Innovation.

Get the Workstation Advantage

The Intel® Xeon® processor-based workstation family delivers rock-solid dependable performance for entry, essential, and expert workbench needs:

- **Reliability –** Workstations let you know your time investment in design will not be met with a blue screen.

- **Stability –** Workstations are certified for select software applications, and validated to ensure they deliver the optimum experience you need.

- **Performance –** Workstations deliver the performance you demand to get the job done, whether you need fast rendering, processing power for compute-intense applications, or memory capacity for large workloads.

Matching Your Needs to the Right Workstation

	ENTRY WORKSTATION Intel® Xeon® Processor 3400 Series	ESSENTIAL WORKSTATION Intel® Xeon® Processor 3600 Series	EXPERT WORKBENCH Intel® Xeon® Processor 5600 Series
Key Features	Integrated graphics for optimum performance and energy-efficient solutions. Support for ECC memory.	Increased compute, ECC memory, and system bandwidth capacity to support larger models.	The maximum compute and ECC memory, capacity you need for high-speed digital prototyping, simulation, and rendering.
3D CAD only	●	●●	●●
3D CAD + Rendering	○	●	●●
3D CAD + Rendering & Simulation	○	○	●●

●● Performance headroom for growth and innovation   ● Ideal system   ○ Not recommended

Get the Intel® Xeon® processor-based workstation advantage today.
For more information, visit **intel.com/go/workstation**

(intel)

by Kenneth Wong

Active Workspace in Code;
Solid Edge Free for Students

Active Workspace, Siemens PLM Software's new product currently in development, is now in code, and is expected to arrive in November, according to Chuck Grindstaff, president of Siemens PLM Software. The revelation was part of Grindstaff's presentation at the company's analyst conference, held in Boston in September.

In late 2009, Siemens PLM Software introduced the notion of HD3D, a way to display product lifecycle data (costs, delays, supplier info, compliance records, etc.) inside a 3D assembly model. The company spent 2010 refining and advancing the concept. Active Workspace, HD3D's manifestation as a product, was quietly previewed to attendees at PLM Connection this May, as a single screenshot that appeared briefly in Grandstaff's keynote address (see "Siemens PLM Connection: Clues to Active Workspace," May 9, 2011 at deskeng.com/virtual_desktop/?p=3653).

"Search is a very important part of Active Workspace," according to Grindstaff. The latest screenshot of Active Workspace shown to the analysts includes search-engine controls and breadcrumbs to navigate through data repositories.

According to Grindstaff's presentation, Active Workspace's characteristics include:
- context-sensitive search and navigation;
- rapid access to data;
- personalized interface;
- viewing large amount of data with preconfigured settings;
- photo-realism;
- customized reports;
- shape-based query; and
- suggested contacts relevant to your project.

Autodesk, one of Siemens PLM Software's rivals, is exploring a similar concept with its data management product, Autodesk Vault. Dassault Systèmes, another rival, has also acknowledged the importance of search in the future of lifecycle data management. In June 2010, Dassault spent $166 million to acquire enterprise search engine Exalead, described by some as the French Google.

At the same conference, Karsten Newbury, Siemens PLM Software's senior vice president and general manager of Velocity Series, broke the news that the company is offering a student edition of Solid Edge for free. "This no-charge 12-month license is available only to students through an easy, instant download," the announcement reads. "It contains not only the latest Solid Edge functionality necessary to create detailed product designs, but also Siemens PLM Software's market leading synchronous technology, which represents the next generation of design software."

Penetration into the education market is seen by leading CAD software makers as a way to cultivate potential users among the next generation. Siemens PLM Software and its rivals — including Autodesk, Dassault Systèmes,

Siemens PLM Software's upcoming product, **Active Workspace**, as previewed at a recent analyst conference in Boston.

and PTC — continue to promote their 3D design software products with heavily discounted student licenses and free licenses for participants of FIRST Robotic Competition.

Newbury's presentation shows Velocity Series licensing is enjoying a global growth. "[It is] continuing strong growth trajectory in 2011 and for (fiscal year) 2011 outlook," he noted. There is a 40% increase in technology partner app submittals, according to him. Siemens PLM Software has also begun holding a series of user gatherings for Velocity Series customers. **DE**

Kenneth Wong *writes about technology, its innovative use, and its implications. One of DE's MCAD/PLM experts, he has written for numerous technology magazines and writes DE's Virtual Desktop blog at deskeng.com/virtual_desktop. You can follow him on Twitter at KennethwongSF, or email him via de-editors@deskeng.com.*

Calibration of Engine Performance at Mercedes-AMG

MathWorks' solution optimizes engine calibration to help engineers design new powertrains.

BY HASAN UZUN

The Mercedes-Benz SLS AMG Gullwing.

The calibration of today's complex, high-degree-of-freedom (DOF) engines can be challenging, even for the most experienced automotive engineers. There are too many variable combinations to test individually, and trade-offs among performance, efficiency, emissions and reliability are complex.

At AMG, we developed a custom engine calibration tool using MATLAB, Model-Based Calibration Toolbox and Parallel Computing Toolbox that enables engineers at all levels of expertise to extract the highest possible performance from AMG powertrains. The tool supports the entire calibration process, from design of experiments (DOE), through model generation and selection, to optimization.
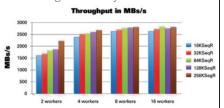
Applying DOE Methodology

Engine performance depends on precise control of a large number of engine control parameters. For example, today's engine control units (ECUs) enable multiple injections per engine cycle—each with different amounts of fuel, adding more parameters to the problem.

MORE ➜ deskeng.com/articles/aabcmf.htm

Accelerating iSCSI SAN Storage

StarWind iSCSI SAN software, coupled with Mellanox Ethernet-based high speed infrastructure, accelerates storage networking connectivity.

StarWind iSCSI SAN turns Windows and Linux servers into iSCSI SAN appliances.

Current data centers require frequent high-speed access between server and storage infrastructures for timely response to customers. Web-based service providers, such as infrastructure as a service (IaaS) or platform as a service (PaaS), provide high-speed access to their customers while keeping infrastructure costs low. For a cost-efficient and higher speed requirement, 40GigE infrastructure combined with software-based iSCSI storage provides a solution.

StarWind SAN iSCSI software running over a Mellanox ConnectX-2 40GigE networking solution provides better performance, high availability (HA) and redundant iSCSI storage solutions at 40Gb/s bandwidth and higher IOPs, according to the company. The installation time for StarWind iSCSI SAN takes only a few minutes. It requires no reboot and is entirely plug-and-play with no downtime.

A Solution Emerges

Under Mellanox Enterprise Datacenter's initiative, StarWind used Mellanox Zorro cluster to run StarWind iSCSI Initiator on HP DL380 systems with Mellanox ConnectX-2 40GigE NICs. Performance benefits were realized over the maximum bandwidth allowed by the PCIe Gen2.

MORE ➜ deskeng.com/articles/aabcme.htm

Stow-and-Go for Space Travel

Caltech researchers use Abaqus FEA to design self-deploying composites booms for satellites.

A MARSIS satellite nears Mars, showing one composite antenna boom already deployed and a second boom in the process of unfolding. *Image courtesy European Space Agency.*

Satellites have to fit a lot into a compact package. Briefly protected inside a rocket while blasted through the atmosphere, a satellite is launched to continue its unmanned mission alone. It uses gyroscopes, altitude thrusters and magnets to regulate sun exposure and stay pointed in the right direction. Once stable, the satellite depends on solar panels to recharge its internal batteries, mirrors and lenses for data capture, and antennas for communications.

Every satellite is susceptible to static electricity buildup from solar wind, the freeze of the Earth's shadow (or deep space), and tiny asteroids that pepper the route like killer bees. In such a hazardous environment, the functional longevity of the average satellite is limited.

While more than 2,000 satellites are estimated to be in Earth's orbit at any one time, the countries and private enterprises that own them must keep sending up replacements at a steep price tag. As a result, aerospace engineers continually strive to come up with smaller, lighter satellites that are cheaper to make, less expensive as rocket payloads, and still capable of fulfilling their high-tech jobs.

MORE ➜ deskeng.com/articles/aabcmg.htm

➜ For the complete application stories visit deskeng.com/fastapps

Build a Virtual Cluster

A virtual cluster can be built from your existing resources and dedicated to completing computationally intensive engineering tasks.

BY PETER VARHOL AND JAMIE J. GOOCH
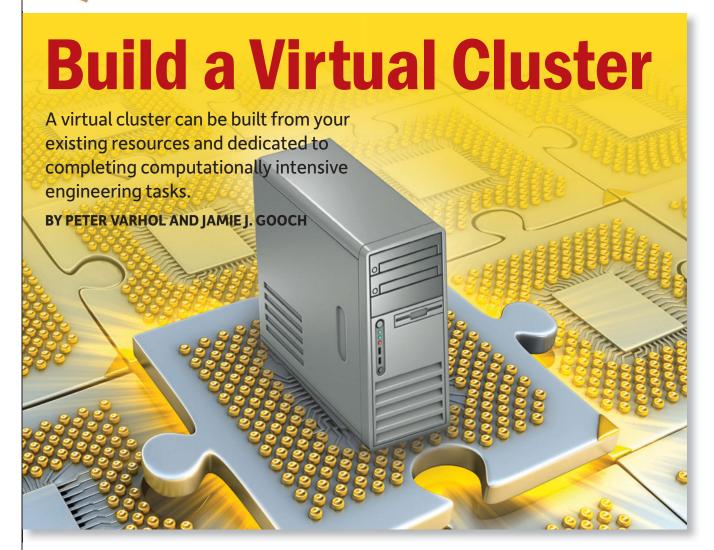
When it comes to today's multi-core workstations, the sum of the parts can be greater than the whole. Many everyday computing tasks—checking email, online research, even CAD modeling—don't begin to tax multi-core processors, and only use a fraction of the memory and hard drive space available. There's a lot of untapped potential sitting in that unassuming computer case. Why let it go to waste?

That's where virtualization comes in. With the right software, a virtual machine—one that doesn't physically exist but shares the resources of one or more workstations—can be created. Even a single workstation can be partitioned to hand over some of its resources to a virtual machine. The more workstations that contribute resources to the virtual machine, however, the more powerful it becomes. Even a small cluster of four 12-core workstations or eight six-core workstations can create a virtual resource of 32 cores. That's enough processing power to significantly speed intensive tasks like simulation or rendering jobs.

And the best part, aside from the fact that you may already have the hardware needed to create a virtual machine, is that the physical workstations are not noticeably hampered by dedicating some of their resources to the virtual machine. So how do you go about creating this unseen processing powerhouse?

Step 1: Make an Assessment
First, determine if you already have the resources available to create and use a virtual cluster. You'll need enough unused workstation resources—specifically processor cores, memory and hard drive space—to dedicate to a virtual workstation cluster.

You should also consider how you'll connect the workstations. You'll need a high-speed interconnect such as gigabit Ethernet or InfiniBand to handle the traffic. You may also need to install additional networking cards so that the workstations can communicate with each other on one network, and with the rest of the company on another to avoid congestion.

Step 2: Choose the Software
Next, decide on the virtualization software you'll need. Virtualization software from companies such as Parallels Holdings

Ltd. and VMware, Inc. have steadily improved over the past few years. It's now common to run two operating systems on one workstation via virtualization. The same concept—using one computer's resources to virtualize another—holds true when creating a virtual cluster from multiple workstations.

Virtualization software offers even better performance when paired with hardware-assisted virtualization technologies. In most virtualization technologies, you use software to map the input and output among the virtual machines, the physical devices, and the computer ports. This is satisfactory for most usages; however, the high input/output (I/O) rate and performance required in many engineering applications is greater than what's possible with typical virtualization technologies. In most virtualized environments, drawing to the screen is significantly slower, and network speed is less than it is from the native operating system.

Hardware virtualization technologies provide a more direct path for input and output for better efficiency.

Step 3: Choose the Hardware

Some workstations are now being designed for virtualization. They may bundled with virtualization software that can take advantage of hardware virtualization.

In the Intel Xeon chipset, Intel Virtualization Technology for Directed I/O (VT-d) provides for direct access from a virtual machine to I/O devices and ports, eliminating the need for a higher-level—and slower—software interface. Instead, devices can deliver I/O just as they would in the native operating system's environment, so performance is comparable to the native environment. Such a setup is necessary in graphics display and network access, where high performance is imperative.

AMD calls its virtualization technology AMD-V. According to the company, it is a set of hardware extensions that allow you to make better use of your resources. All new AMD processors, including the 12-core Opteron processors, are designed to use AMD-V features, such as hardware-based virtual machine memory management and direct device access by virtual machines.

In addition to hardware virtualization, you should ensure the workstations you plan to use have the ability to accept multiple networking cards and can be configured with a sufficient amount of memory.

Step 4: Divide and Conquer

The next step is to allocate resources from each workstation that will be used by the virtual cluster. It's possible, for example, to allocate four cores and 8GB of memory on one workstation to a design engineer's interactive tasks, and eight cores and 16GB of memory for the same purpose on another system in the same cluster.

Some engineering applications can make use of all of the cores and memory available, while others require a great deal of network bandwidth. Cluster administrators have to take into account the needs of the application and the type of work being done in order to appropriately configure each workstation.

Users of CAD software, for example, typically only use a small number of cores when running their applications. As reported in the September issue of *DE*, airplane component and system supplier, Parker Aerospace built a virtual cluster to relieve simulation and analysis bottlenecks. The company allocated individual cores and memory for the desktop operating system, productivity applications, and any other applications. It was able to leave six or eight cores of each of the four 12-core workstations available for clustering.

This one-time setup of each individual application (e.g. ANSYS, Microsoft Office, LS-Dyna, CFD Design) on each clustered workstation took anywhere from a few minutes to a few days for Parker Aerospace, depending on the characteristics of the jobs and underlying applications. Some were very predictable, while others varied a great deal. But this only had to be done once to determine how many cores to initially al-

Taking a Deeper Look

This article concludes a series that looked at Parker Hannafin's engineering practices, specifically how Parker Aerospace used a virtual cluster to speed simulations. *DE* also produced a white paper on workstation cluster computing, which was sponsored by Intel and HP. It can be downloaded for free at deskeng.com/workstationcluster. If your company would like to be the subject of a "*DE* On the Case" story, email us at de-editors@deskeng.com.

locate. Once the computational behavior of a task was ascertained, Parker Aerospace could allocate the appropriate amount of computing resources to its purpose.

Making the Decision

If you've performed a hardware upgrade in the last few years, you may have already invested in a fair number of multi-core workstations. Since complex simulation and analysis jobs have become standard practice in the industry, creating a virtual HPC cluster to process these tasks out of existing workstations gives an organization the benefit of an HPC server without the cost and expense of purchasing and maintaining one.

If you're looking at upgrading hardware, keep virtualization options in mind. Even if you don't need a virtual cluster now, investing in the right hardware will give you the option in the future. DE

Peter Varhol *is a Contributing Editor for* Desktop Engineering *covering high-performance computing.* Jamie Gooch *is Managing Editor of* Desktop Engineering. *Contact them via de-editors@deskeng.com.*

Cloud *Control*

Document control in the cloud can improve quality management systems and productivity for manufacturers.

BY GREG MILLIKEN

The efficiency and productivity of all businesses are dependent on team members consistently and repeatedly executing important business processes, as well as always having access to the latest and most up-to-date information, such as the correct versions of important documents. However, access to the right information and correct documents is particularly critical in certain industries, especially when compliance with industry and quality standards is involved, such as ISO 9001 related to Quality Management Systems (QMS), commonly employed in manufacturing organizations.

Quality management systems utilized to help achieve ISO certification encompass the notion of Good Manufacturing Practices (GMP), or more broadly, "Good Practice" quality guidelines (GxP) that emphasize records of events, secure access permissions, and standard operating procedures and processes (SOPs), rather than just managing documents or files. This includes ensuring that outdated file versions are managed appropriately, that correct forms and templates are used, and that documents follow mandatory review policies.

Other important aspects of a comprehensive QMS include capturing and recording signatures at the necessary time, as well as making sure personnel get the right training, read required documentation and have clear to-do lists for important tasks. In addition, an effective QMS ensures that team members only do and see what their authorization allows, and that timed activities actually take place on time. Lastly, the ability to corroborate that these tasks have been completed is a fundamental requirement.

Manufacturing, design and engineering present particular challenges: Not only can there be specific quality management standards in play, but the nature of the work and related documentation can lead to errors quickly propagating to related parts and assemblies, as well as to expensive tooling necessary for manufacture. Changes to a detailed CAD file can be hard to spot, often consisting of merely a slight change in a dimension or tolerance value.

The problem only increases as teams increase in size and become distributed with the addition of external contractors and suppliers, resulting in the risk of greater financial losses, lower quality—and worse, a design failing in use. This results in downtime, and possibly damage to other equipment or injury to employees or customers.

Simply stated, manufacturing, design and engineering environments require that team members are always working on the most current and accurate project files and documentation, including, but not limited to CAD files, to eliminate errors and improve quality and efficiency. The result is a more profitable business that produces high-quality products according to processes and procedures that comply with industry standards.

Seeking Solutions

According to the Association for Information and Image Management (AIIM), professionals can spend up to 50% of their time looking for the right information. And professionals in manufacturing and engineering are no exception; engineers and designers often spend an inordinate amount of time tracking down the current version of a CAD file or other project-related document, which may reside on an individual workstation, or lie buried somewhere within the internal network in a home-grown, and often chaotic, folder structure.

It's not uncommon for subsequent design changes to be made to the wrong version of a document or CAD file, resulting in significant additional rework when and if the problem is finally identified. If the error is not caught before manufacture, the result is often wasted material and scrapped parts, or low-quality products making it out the door to customers.

This isn't a new revelation. Manufacturing and engineering businesses have been trying to solve this problem for years, and numerous specialized solutions exist. All of the large CAD vendors have their own product lifecycle management (PLM) or product data management (PDM) solutions (naturally, focused on their particular CAD system), and there are numerous independent vendors offering solutions as well. Most of these solutions are traditional on-premise systems for managing product-related files like CAD files, bills of material (BOMs), process plans, and so on. On-premise solutions such as this require the purchase, configuration and maintenance of internal servers, and the solutions themselves are traditionally expensive and often complex, requiring sophisticated IT resources to install, configure and maintain them. Further, CAD-centric solutions that are implemented on the engineering side of the house can create information silos that hinder data flow across an organization to other groups such as accounting, marketing and sales.

Traditional "On-Premise" Software License Model

SaaS "Cloud" Subscription Model

Upfront Costs	
• Server hardware • Server platform software licenses • Document management licenses • IT staff time for implementation	• None

Recurring Costs	
• IT staff time (routine maintenance, backup, troubleshooting)	• Monthly subscription fee (storage, software, backup and support included, although costs increase when more is needed)

Budgetary Categorization	
• Capital assets	• Operating expenses

Future Costs	
• Upgrade or expansion of server hardware • Upgrade/add platform licenses • Upgrade/add document management licenses	• Monthly subscription fee (typically include upgrades) (rates scalable based on users/storage)

But with the emergence of the cloud, new options have become available that offer businesses flexibility in terms of how and what they pay for such solutions, the need for IT resources, and data access across the organization and with external partners and customers.

Cloud Benefits

The cloud makes it possible for essentially any business to deploy a document/data management system to alleviate the common problems of misplaced documents and lack of version control, and to facilitate the implementation of a comprehensive QMS.

Cloud-based document management systems provide design engineers and other team members with a framework for centrally organizing, managing and tracking the important documents, information and processes typical in manufacturing firms. In addition, moving to the cloud can often better support lean manufacturing practices and quality management programs to ensure efficient and ISO-compliant manufacturing processes.

In addition to ensuring that documents are centrally organized, tracked and easily located, and proper versions of files are being accessed and edited, cloud-based document management also provides a number of other benefits:

• **Convenient online access to all company documents from any location.** Cloud access makes sense for today's notebook- and mobile device-equipped engineers who can retrieve any document from home, while waiting for a flight, or while in a meeting at a client's office. Similarly, organizations with multiple office locations can now share data from a single repository, rather than separate, unrelated systems such as an internal network, FTP site, or web file sharing service.

• **Freedom from the responsibilities of server ownership and management.** Running servers requires the tasks of monitoring, maintenance and administration, as well as expert troubleshooting when a server goes down. Manufacturers that leverage cloud solutions are liberated from not only the initial costs of buying dedicated hardware and software, but also all the sometimes-unpredictable costs of ownership.

• **Reliability** because an increasing number of cloud vendors store data in modern, well-designed data centers with large numbers of servers and multiple levels of redundancy, such that when a server goes down, another is always available to take its place.

• **Better protection against data loss.** Notebooks can be lost; hard drives can fail; servers can crash. Having information in the cloud secures manufacturers against accidents that can occur to physical devices. Cloud service providers also typically provide automated regular backups of your data, often with storage and archival in a safe, secondary different location.

• **More control over information access.** Multiple firewalls protect against unwanted access from the outside, while document management permissions can place controls on specified files or document types based on password-protected user accounts and groups. Multiple levels of access can be implemented—read-only, edit, delete, etc.—enforcing security and privacy of sensitive material.

• **Scalability.** If more storage capacity is required, one simply pays the provider incrementally more, rather than having to purchase an additional server. Conversely, if work is cut back and fewer seats are needed, the monthly fee goes down.

• **Operating expenses vs. capital expenses.** Cloud storage and service subscriptions are classified as operating expenses for budget purposes, while in-house servers and software licenses are considered capital assets—which depreciate, generate tax implications and eventually require replacement. Manufacturers utilizing cloud-based services pay for these new IT services much like they do utility bills. This accounting re-categorization allows them to utilize their capital budget for other projects.

Not only do cloud solutions represent a significant shift in how businesses store and access their digital content, but also how they pay for it. In a software-as-a-service (SaaS), or cloud subscription model, all expenses for storage, management, maintenance and support are typically included in a single monthly fee. In effect, the subscriber leases access to the application and data storage space as opposed to owning the software and the in-house hardware necessary to deploy and maintain it.

The cloud offers engineers the benefits of online storage, file sharing and document control without large initial outlays for hardware and licenses associated with on-premise installations. Similarly, cloud-based services typically take up less ongoing IT and operational resources because the cloud service provider takes over the responsibility of operating servers, backing up data and upgrading the technology.

With cloud-based document and quality management, engineering firms can control and regulate business processes and file access, control visibility and manage revisions across multiple departments, protect themselves against data loss and securities breaches, and provide the necessary level of automated communication to those people who are affected by changes. Furthermore, all of these benefits can be realized while eliminating upfront system costs and associated hardware resources.

From this perspective, it's easy to see why cloud-based document management is becoming more and more common among manufacturers today. **DE**

Greg Milliken *is president of M-Files Inc., the developer of M-Files professional document management software and the cloud-based document management service M-Files Cloud Vault. Contact him via de-editors@deskeng.com.*

INFO ➜ M-Files Inc.: M-Files.com

For more information on this topic, visit deskeng.com.

Big, Fast and Hot

The expensive Eurocom
Panther 3.0 mobile workstation
rivals top-of-the-line desktop systems.

BY DAVID COHN

Eurocom has been building some of the world's most powerful portable computers for more than 20 years. After putting the company's impressive D900F Panther system through its paces earlier this year (see *DE* March 2011), we were quite excited when its successor arrived. And once again we were impressed—the Eurocom Panther 3.0 delivered performance on par with the best desktop workstations. But this time around, we had some concerns with the system's cooling and power consumption.

Eurocom bills its new Panther 3.0 as a desktop replacement as well as both a server on-the-go and a blazingly fast high-end gaming machine. Built around an Intel X58 chipset, the Panther 3.0 incorporates components typically found only in true workstations, including a top-of-the-line Intel Xeon 6-core CPU, up to four internal hard drives for 3.25TB of storage and as much as 24GB of memory. But all that power comes at a significant price, both economic and ergonomic. The Eurocom Panther 3.0 is expensive, heavy and noisy, and packs a lot into what may be too small a space.

Well, not that small. The Eurocom Panther 3.0 measures 16.5x11.25x3.0 in., and weighs in at 12.25 lbs., plus another 3.5 lbs. for its enormous (8.25x4.5x 2-in.) 300-watt power supply. The system is configured like a traditional notebook computer, with a nicely finished black plastic and brushed aluminum case, but it's nearly twice as thick as more conventional notebooks. In fact, sitting on a desk, the Panther 3.0 almost looks like it's resting on top of a second, smaller notebook computer.

Raising the lid reveals a beautiful 17.3-in. Full HD display with a native resolution of 1920x1080 (down from 1920x1200 on the D900F Panther), and a full-size keyboard with a separate numeric keypad. There's also a touchpad centered below the keyboard with a pair of buttons and a fingerprint reader.

Our evaluation unit came equipped with a 2-megapixel webcam centered above the LCD, flanked by a pair of mi-crophones. Three speakers are located in the hinge below the LCD, with two more in the top corners above the keyboard. Between these are a series of touch sensors for adjusting speaker volume and toggling the Wi-Fi, Bluetooth and webcam on and off, and hard drive and keyboard status indicators. These controls, as well as the large round power button to the right of keyboard, glow bright blue when active.

Configuration Options Galore

The right side of the case provides jacks for audio line-in, S/PDIF-out, microphone-in and headphone, as well as three USB 2.0 ports and a security lock slot. The left side houses a DVI-out port, cable antenna jack (for the systems' optional TV tuner), RJ-45 LAN jack for the built-in Gigabit Ethernet, HDMI-out port, two USB 3.0 ports, eSATA port, mini-IEEE 1394b (FireWire) port, a 9-in-1 card reader, and the optical device drive bay. There is no provision for a modem, however.

Our evaluation unit came with a Blu-Ray combo drive with 6X Blu-Ray write and 8X DVD+-RW capabilities. The rear of the case contains just the external power connector centered between fan intake vents, while the front includes more air vents as well as a cleverly hidden ExpressCard slot, an infrared receiver (for use with the optional TV tuner) and LED indicators for power and battery.

Like most Eurocom systems, configuring the Panther 3.0 is all about choices. The LCD display is available with or without non-glare coatings, and can be driven by a choice of 16 different GPUs, ranging from the NVIDIA GeForce GTX 460M to the top-of-the-line NVIDIA Quadro 5010M installed in our evaluation unit. That graphics accelerator, equipped with 4GB of GDDR5 memory and 384 compute unified device architecture (CUDA) cores, can also power an external 30-in. monitor at up to 2560x1600—but it also adds $2,050 to the cost of the system.

You can configure the Panther 3.0 with any one of 10 different Intel processors. In our case, Eurocom included a 3.46GHz Intel Xeon X5690 CPU with six cores and 12MB of L3 cache. As mentioned, the system can accommodate up to 24GB of memory in three 204-pin small outline dual in-line memory module (SO-DIMM) sockets, which are located beneath the keyboard. While the manual illustrates how to access these sockets, it's complex enough that most customers will opt to let Eurocom do the work. Our system came equipped with 12GB of memory, using three 4GB DDR3 1600MHz memory modules.

The Panther 3.0 supports up to three standard hard drives (or a fourth if you forego the optical drive), and these are much easier to install. Two mount in a special drive cage within an easily accessible bay on the bottom of the system; the third is located in a separate bay below the battery. Eurocom provided a 250GB Intel solid state hard drive as the primary drive in our evaluation unit, along with a pair of Seagate Momentus 750GB secondary drives. The solid state drive enabled the Panther 3.0 to boot to Windows 7 64-bit in just 50 seconds.

The 78.44WH lithium-ion battery is useful more as an

Design Engineering Mobile Workstations Compared

		Eurocom Panther 3.0 mobile workstation (3.46GHz Intel Xeon 5690 six-core CPU, NVIDIA Quadro 5010M, 12GB RAM)	Eurocom D900F Panther mobile workstation (3.3GHz Intel Xeon X5680 six-core CPU, NVIDIA Quadro FX 3800M, 12GB RAM)		HP Elitebook 8540w mobile workstation with DreamColor display (1.60GHz Intel Core i7 QM720 quadcore CPU, NVIDIA Quadro FX 1800M, 8GB RAM)		Dell Precision M6500 mobile workstation (2.00GHz Intel Core i7 X920 CPU, NVIDIA Quadro FX 380M, 4GB RAM)		HP EliteBook 8530w mobile workstation (2.53GHz Intel Core 2 Duo T9400 CPU, NVIDIA Quadro FX770M, 4GB RAM)		Lenovo ThinkPad W700 mobile workstation (2.53GHz Intel Core 2 Quad Core Q9300 CPU, NVIDIA Quadro FX3700M, 4GB RAM)	Lenovo ThinkPad W700 mobile workstation (2.53GHz Intel Core 2 Quad Core Q9300 CPU, NVIDIA Quadro FX3700M, 4GB RAM)
Price as tested		$8,027	$7,467		$3,657		$4,430		$2,822		$3,524	$4,549
Date tested		8/20/11	1/14/11		12/13/10		4/23/10		12/18/08		10/22/08	10/22/08
Operating System		Windows 7	Windows XP	Windows 7	Windows XP	Windows 7	Windows XP	Windows 7	Windows XP	Windows Vista	Windows XP	Windows Vista
SPECviewperf	higher											
3dsmax-04		74.50	75.97	79.11	43.31	56.10	49.56	52.35	33.38	32.21	34.23	12.12
catia-02		95.31	85.23	82.83	53.08	59.43	64.31	61.72	42.41	39.75	45.01	13.64
ensight-03		101.50	66.02	60.10	41.52	41.05	58.28	47.75	37.42	34.24	43.31	15.03
maya-02		407.09	290.35	236.55	204.12	159.95	283.64	212.05	149.21	108.33	165.87	25.37
proe-04		95.73	84.84	77.09	57.51	57.06	70.91	61.96	42.92	39.33	45.67	10.28
SW-01		198.76	178.01	163.34	109.6	102.93	152.41	132.68	67.98	59.75	90.01	17.19
tcvis-01		91.18	47.36	42.28	29.84	27.15	47.71	39.44	21.42	19.19	28.34	4.48
ugnx-01		86.62	43.18	39.34	27.75	28.10	39.60	33.64	19.85	18.11	30.91	4.18
SPECapc SolidWorks	lower											
Score	seconds	n/a	135.63	n/a	198.16	n/a	175.72	n/a	182.63	n/a	187.27	n/a
Graphics	seconds	n/a	51.94	n/a	67.55	n/a	58.99	n/a	62.16	n/a	60.87	n/a
CPU	seconds	n/a	29.22	n/a	45.42	n/a	37.62	n/a	39.99	n/a	44.40	n/a
I/O	seconds	n/a	58.76	n/a	89.8	n/a	83.48	n/a	83.69	n/a	96.66	n/a
SPECapc SolidWorks	higher											
Score	ratio	n/a	5.66	n/a	4.08	n/a	4.75	n/a	4.75	n/a	4.47	n/a
Graphics	ratio	n/a	3.61	n/a	2.69	n/a	3.09	n/a	3.26	n/a	3.15	n/a
CPU	ratio	n/a	10.19	n/a	7.1	n/a	8.58	n/a	8.07	n/a	7.72	n/a
I/O	ratio	n/a	5.14	n/a	3.53	n/a	3.79	n/a	3.78	n/a	3.65	n/a
Autodesk Render Test	lower											
Time	seconds	53.3	57.16	51.83	188.5	146.81	168.33	180.16	318.4	324.60	162.00	291.60
Battery Test	higher											
Time	hours:min	0:38	n/a	1:17	1:28	1:21	2:06	2:05	3:21	3:00	2:15	1:20

Numbers in blue indicate best recorded results. Numbers in red indicate worst recorded results.

uninterruptable power supply (UPS) rather than for actually powering the system. In fact, you'll want to get back to AC power or shut the system down quickly when running on battery power, as the battery in our evaluation unit lasted just 38 minutes.

In addition to the poor battery life, the Panther 3.0 also proved to be exceptionally noisy and ran quite hot. Even at idle, the fan noise from the Panther 3.0 was louder than any other device in the office and could be heard from more than 20 ft. away. When running benchmark tests, that noise became even more pronounced, and the air leaving the vents on the rear of the system reached 107°F.

Performance at a Price

After the results we recorded when testing the previous Eurocom system, we expected to see stellar performance from the Panther 3.0. Expectations were especially high for the SPECviewperf test, which focuses solely on graphics. We got what we had hoped for: The Panther 3.0 outperformed every mobile system we've ever reviewed, owing largely to its incredible NVIDIA GPU. In fact, the Panther 3.0's viewperf results rivaled those of top-of-the-line desktop workstations.

Unfortunately, because the system came with only Windows 7 Professional installed, we were unable to run the SPECapc SolidWorks benchmark. But judging from the performance we experienced when running SolidWorks and other CAD software, we'd definitely expect the Eurocom Panther 3.0 to do an exceptional job on even the most demanding CAD, CAE and DCC tasks.

The results on our AutoCAD rendering test raised some serious concerns, however. With hyper-threading enabled, the equivalent of 12 CPU cores enabled the Panther 3.0 to complete our rendering in as little as 52 seconds—essentially the same as the previously reviewed Eurocom D900F Panther, which had a slightly slower CPU. But in repeated tests, the system averaged 75.5 seconds, and results varied widely. After conferring with Eurocom, the company concluded that it was a power problem and sent us a second 300-watt power supply, along with a small converter box to join the two power supplies together to power the computer. This did improve performance slightly, but consecutive results continued to be somewhat erratic and trended downward.

While rendering each scene, the Panther's cooling fans became markedly louder. Concluding that there was also a cooling problem, we ran the rendering test manually rather than using a script, waiting several minutes between each rendering until the fans returned to their slowest speed. Sure enough, each rendering now averaged 53.3 seconds, and one iteration took just 49 seconds—a truly impressive result.

That said, all of the Panther's power comes at a staggering price. The standard Panther 3.0 configuration sells for $2,726. Once we tallied up all of the options included in our evaluation unit, however, the system we received would cost $8,215—with the additional power supply and converter adding another $249 to the price—making it the third-most expensive system we've ever reviewed, mobile or desktop.

If that price wasn't daunting enough, the standard Eurocom warranty only covers the system for one year and requires owners to ship the system back to a factory depot for repairs. Extending the warranty for a second year adds $150 to the cost, while adding a third year adds $268. But Eurocom does offer lifetime upgradability, as well as credits of up to 20% if you trade in an older system when purchasing a new Panther.

While the Eurocom Panther 3.0 is meant to replace desktop workstations for power users on the go, it's likely to appeal to a very narrow market. The burden in terms of price, weight, power consumption and CPU cooling are serious considerations. If portable performance is your only criteria, the Eurocom Panther 3.0 packs a lot of power, but it may actually be too much power in too small a space. **DE**

David Cohn *is the technical publishing manager at 4D Technologies. He also does consulting and technical writing from his home in Bellingham, WA, and has been benchmarking PCs since 1984. He's a contributing editor to* Desktop Engineering *and the author of more than a dozen books. Contact him via email at david@dscohn.com or visit his website at DSCohn.com.*

INFO ➜ **Eurocom Corp.:** Eurocom.com

Eurocom Panther 3.0 Mobile Workstation

Price: $8,215 as tested ($2,726 base price)
System requirements:

- **Size:** 16.5x11.25x3.0-in. (WxDxH) notebook
- **Weight:** 12.25 lbs. as tested, plus 3.5-lb. power supply
- **CPU:** 3.46GHz Intel Xeon X5690 6-core w/12MB L3 cache
- **Memory:** 12GB 1600MHz DDR3 SDRAM (24GB max)
- **Graphics:** NVIDIA Quadro 5010M with 4GB memory
- **LCD:** 17.3-in. diagonal (1920x1080)
- **Hard Disk:** 250GB SSD and two 750GB 7200 rpm Seagate Momentus drives
- **Optical:** Matshita Multi Blu-ray Disc write and DVD+/-RW dual layer
- **Audio:** Line-in, SP/DIF-out, microphone, headphone, built-in microphone and speakers
- **Network:** Integrated Gigabit Ethernet (10/100/1000 NIC); Intel 802.11 a/b/g/n wireless LAN; optional integrated Bluetooth 2.0
- **Keyboard:** Integrated 98-key keyboard with numeric keypad
- **Pointing Device:** Integrated 2-button touchpad
- **Other:** Three USB 2.0, two USB 3.0, one mini IEEE 1394 Firewire, eSATA, Express Card slot, 9-in-1 card reader, DVI-out, HDMI-out, 2MP webcam, CATV. Laptop case included

For more information on this topic, visit deskeng.com.

Autodesk for Mac:
Take Two

Autodesk updates and adds to its native Macintosh version in
AutoCAD 2012 for the Mac.

BY DAVID COHN

L ast year, Autodesk finally put
the rumors to rest, releasing
the first native Macintosh
version of its flagship CAD soft-
ware. It had been more than 12
years since the company last pro-
duced a version of AutoCAD spe-
cifically for the Mac—although
ever since Apple switched from
Motorola to Intel processors, lots
of Mac users had been running
versions of AutoCAD using virtual
machine software.

Now, one year later, Autodesk
has released AutoCAD 2012 for
Mac, updating the Mac version
to take advantage of many of the
same new features and functions
introduced in the Windows version of AutoCAD
2012 released earlier this year (see *DE*, July 2011).

AutoCAD 2012 for Mac looks much different than
its Windows counterpart. Instead of the ribbon in-
terface, AutoCAD for Mac puts its tools where Mac users
expect to find them: in the Apple menu bar. In addition,
the Mac version uses a number of palettes docked to the
left and right of the drawing canvas.

By default, the Tool Sets palette appears on the left,
while the Layers and Properties Inspector palettes are
docked on the right. There are also Command Line and
Status Bar palettes across the bottom. All of these palettes
can be dragged around the screen, docked in other loca-
tions, and collapsed in various ways.

We noted that the most important aspects of Auto-
CAD—100% DWG file compatibility and robust cus-
tomization and support for AutoLISP and ObjectARX—
were most definitely hallmarks of the Mac version as well.
What's more, AutoCAD 2011 for Mac introduced other
exclusive features, such as a set of viewport controls lo-

Users are now notified when external references have
changed, and the Layers palette includes Layer Groups,
similar to Layer Filters in the Windows version.

cated in the upper-left corner of each drawing viewport,
and support for multi-touch gestures. Those viewport
controls eventually found their way into the Windows
version of AutoCAD 2012; the gestures remain exclusive
to the Mac.

New Features

Although some aspects of AutoCAD 2012 for Mac remain
exclusive to the Mac, virtually all of the improvements in Au-
toCAD 2012 have made their way into the new Mac release.
For example, the user coordinate system (UCS) icon can be
selected and manipulated, making it much easier to adjust the
UCS, particularly when working in 3D.

AutoCAD 2012 for Mac also includes auto-complete.
As you type at the program's command prompt or dynamic
input field, AutoCAD displays a list of commands, system

AutoCAD LT for Mac 2012

When Autodesk shipped AutoCAD 2011 for Mac, it only addressed the needs of users looking for a native OS X version of its flagship product. This time around, the company has also released a Mac version of AutoCAD LT, for those users who need only 2D drafting and detailing tools. At a cost of $899.99, AutoCAD LT for Mac 2012 shares the same Macintosh interface as AutoCAD 2012 for Mac, while providing a feature set almost identical to that found in the Windows version of AutoCAD LT 2012.
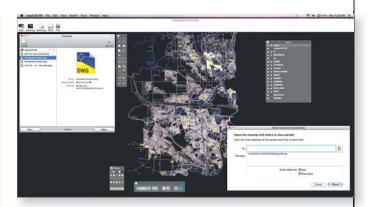
For example, users can create and modify drawing objects, add text and dimensions, and attach external references and image files. But because AutoCAD LT is strictly a 2D program, it lacks the ViewCube and there are no in-canvas viewport controls. There are also no 3D modeling or rendering tools, and no support for AutoLISP or Object ARX—but this is true for both the Windows and Mac versions. So while AutoCAD LT for Mac 2012 also lacks a number of other features found in the Windows version, such as Sheet Sets and eTransmit, the list is not quite as long simply because LT doesn't include all of the features found in the full version of AutoCAD.

Still, AutoCAD LT for Mac 2012 does support gestures and can use Cover Flow to locate drawings. AutoCAD LT for Mac 2012 also takes advantage of the same contextual visors for functions, such as hatch editing, manipulating external references, and adjusting associative arrays as AutoCAD 2012 for Mac. It also has largely the same system requirements.

Unlike the Mac version of AutoCAD or even the Windows version of AutoCAD LT, however, prospective customers cannot download a free trial version of AutoCAD LT for Mac.

Currently, AutoCAD for Mac is available only from Apple's Mac App Store in the US and Canada. There is no subscription option, and AutoCAD LT for Mac does not support network licensing. But it also does not need to be activated. You can begin using the software as soon as it finishes downloading.

If your drawing needs are limited to 2D, AutoCAD LT for Mac may be an ideal native Macintosh solution.

AutoCAD WS for Mac lets you edit DWG files on your Mac and collaborate in real time with other team members, using the web or mobile versions of AutoCAD WS or an actual copy of AutoCAD or AutoCAD LT.

variables, or command aliases matching what you've typed thus far. You can then quickly select from the list.

Other new AutoCAD 2012 features, such as enhanced multifunctional grips, better control over multi-leaders, the new Blend command, improved join functionality, and previews when using the Fillet or Chamfer commands or when working with surfaces have also been included in AutoCAD 2012 for Mac. The Copy command also includes the new Array option, so you can create a linear array of independent objects.
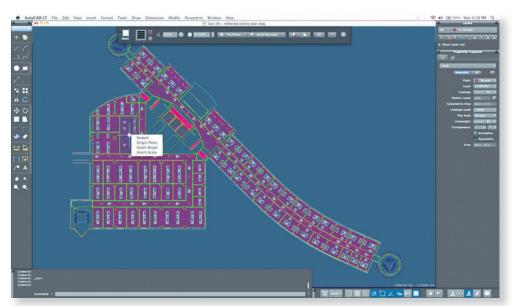
Speaking of arrays, associative arrays—potentially one of the biggest AutoCAD productivity improvements in years—are also included in the new Mac release. Rather than simply making duplicate copies of selected objects arranged in a rectangular or circular pattern, there are now three separate commands for creating rectangular arrays, polar arrays or path arrays. Path arrays distribute items along a path formed by objects such as lines, arcs, polylines or helixes, and can follow 2D or 3D paths. All three types of arrays can also include multiple levels.

Visors to the Rescue

If there were any significant complaints about the initial release of AutoCAD for Mac, they focused more on what was missing. For example, AutoCAD 2011 offered a huge improvement in the way users interacted with hatch patterns and gradient fills, enabling users to preview and adjust these objects before committing changes to the drawing. Much of this capability relied on the ribbon, but because the Mac doesn't have a ribbon interface, AutoCAD 2011 for Mac used a dialog box—similar to the way things worked in AutoCAD 2010 and earlier releases.

Happily, Autodesk has found an excellent alternative to the ribbon: contextual visors. Visors aren't new. AutoCAD 2011 for Mac used a visor when creating or editing multi-

While offering the more limited toolset of AutoCAD LT, AutoCAD LT for Mac 2012 sports the same native Mac interface as AutoCAD 2012 for Mac—and benefits from new features, such as the Hatch Editing contextual visor.

Application Preferences dialog. In the 2011 Mac release, a Content palette made up for some lost functionality because of the absence of the Tool palette, by letting you organize and access frequently used blocks. But you could only add blocks to a drawing by clicking on the block in the Content palette. Now, you can either double-click or drag-and-drop the block from the Content palette to the drawing.

The new Mac release also goes a long way toward filling other gaps that we had noted last year. For example, AutoCAD 2011 for Mac lacked layer filters. To remedy that, AutoCAD 2012 for Mac now offers layer groups within the Layers palette. After creating a layer group, you can add layers to a group by simply dragging and dropping from the main layer list. You can then change the status of all layers in a group with a single click.

AutoCAD 2012 for Mac also now includes a Batch Publish feature nearly identical to the Publish command in AutoCAD, and you can publish to the printer named in each layout page setup—or to single- or multi-sheet PDF files. The Export command also now supports PDF, in addition to ACIS, STL, EPS, DXX, BMP and DWG, and lets you choose what to export (the display, drawing extents or a window). It also lets you override the current page setup. The new Mac release also includes the Drawing Recovery Manager.

Unfortunately, there is still a pretty long list of features and functions totally missing from AutoCAD 2012 for Mac. For example, while you can insert and manipulate dynamic blocks, you still cannot create or edit dynamic blocks in the Block Editor. There's also no Action Recorder, Batch Standards checker, Layer Translator, Sheet Set Manager, AutoPublish, eTransmit or drawing object hyperlink capability.

line text. But in AutoCAD 2012 for Mac, visors are now provided for other operations as well. For example, when creating or modifying hatch patterns or gradient fills, you no longer see a dialog. Rather, the program displays a visor containing controls for selecting the hatch pattern, adjusting its angle and scale, and so on. You can then preview changes made to the hatch or fill before saving those changes to the drawing.

When modifying associative arrays, you see a visor containing tools for changing the base point, editing the source object, replacing items in the array, and resetting the array—similar to those that appear in the ribbon when working in the Windows version. Missing, however, are visor-based controls for adjusting the spacing between arrayed objects or the number of objects in the array; for that, you'll have to rely on grip editing.

There's also a visor that displays when working in a paper space layout so you can add a new viewport, edit the page setup, or print the current layout. Another displays when you select an external reference so you can edit, open or clip the xref and another. You can also now double-click to edit an external reference in place, and AutoCAD 2012 for Mac now notifies you when external references change.

Filling Some Gaps

Some users complained that the icons in the Toolset and Status Bar in AutoCAD 2011 for Mac were too small. Happily, users can now easily adjust these from within the

Good Enough?

AutoCAD 2012 for Mac supports network licensing. Prospective users can download a free 30-day trial version, and existing AutoCAD 2012 users can even activate AutoCAD 2012 for Mac using the same license number and activation code, so multi-seat licenses can be shared. You

AutoCAD 2012 for Mac incorporates most of the new features in AutoCAD 2012, such as associative arrays, and improves upon the previous release with the addition of tools such as contextual visors.

David Cohn *is the technical publishing manager at 4D Technologies and the author of that company's CADLearning eLearning courses for AutoCAD. He also does consulting and technical writing from his home in Bellingham, WA. He's a contributing editor to* Desktop Engineering *and the author of more than a dozen books on AutoCAD, a program that he has used for more than 25 years. Contact him via email at david@ dscohn.com or visit his website at DSCohn.com.*

INFO ➡ Autodesk Inc.: Autodesk.com

AutoCAD 2012 for Mac 2012

Price: $3,995 (full system); $450 (annual subscription)

System requirements:
- **Computer:** Apple Mac Pro 4,1 or later; MacBook Pro 5,1 or later (MacBook Pro 6,1 or later recommended); iMac 8,1 or later (iMac 11,1 or later recommended); Mac mini 3,1 or later (Mac mini 4,1 or later recommended); MacBook Air 2,1 or later; MacBook 5,1 or later (MacBook 7,1 or later recommended)
- **Operating System:** Mac OS X v10.6.4 or later
- **CPU:** 64-bit Intel processor
- **Memory:** 3GB RAM (4GB recommended)
- **Disk space:** 2.5GB free disk space (3GB recommended)
- **Video:** 1280x800 display with true color (1600x1200 with true color recommended)

can even use the License Transfer Utility to import and export a stand-alone license between Windows and Macintosh systems.

Because it represented a complete rewrite of AutoCAD as a native OS X application, AutoCAD 2011 for Mac was essentially a first release. This revision smooths out many of the rough spots. There are still some missing features that could be showstoppers for some, who may elect to run AutoCAD 2012 on their Macs using a virtual machine. But for many, AutoCAD 2012 for Mac provides lots of power—and may be a perfect native Macintosh solution. **DE**

AutoCAD WS Enables Free CAD Sharing

Last year, when Autodesk first announced AutoCAD for Mac, the company also introduced AutoCAD WS, a free program that allowed users to edit and share 2D AutoCAD drawings in real time. Initially released as a free mobile application for the iPad, iPhone and iPod Touch, the company subsequently expanded its capabilities and broadened support to include Android devices and a web-based version accessible using any web browser.

With the release of AutoCAD and AutoCAD LT for Mac 2012, Autodesk has also released AutoCAD WS for Mac, a native Macintosh version of the software, available as a free download from the Mac App Store.

Both the Windows and Macintosh versions of AutoCAD and AutoCAD LT 2012 include tools to upload drawings to AutoCAD WS. After setting up a free online account, you can access your drawings using any version of AutoCAD WS and, if you choose, share them with others. You can set permissions to let others view, edit or download drawings. Once downloaded, a drawing can be opened using AutoCAD or AutoCAD LT.

AutoCAD WS includes tools to pan and zoom, control layers, and draw basic geometry such as lines, circles, polylines and text. You can also move, copy, rotate, scale and erase objects, measure distances and areas, and switch between model space and paper space.

The web-based version includes additional tools not found in the mobile app, while the capabilities of the new Mac version fall somewhere in between. For example, both the web-based and Mac versions let you plot drawings to PDF or DWF, while the Mac and web-based versions include Ortho mode and object snap. The Mac version lacks timeline capabilities, however.

All versions let you co-edit drawings with others at the same time, even with people who don't have a copy of AutoCAD or AutoCAD LT. That's definitely one of the biggest features of this cool little app.

Building Better Products Faster with Workstation Clusters

Virtual clusters using HP Z800 Workstations provide a complete engineering workbench with high-performance cluster computing power.

BY PETER VARHOL

Screen image courtesy of ANSYS

Design engineers are highly sensitive to the pressures of bringing a product to market on time. Not being on schedule means that a manufactured product may not meet its market window and fail.

But schedules are designed for the best of circumstances, when nothing goes wrong. Of course, that's almost never the case. As a result, work doesn't progress as rapidly as planned, and designs aren't the best they can possibly be.

A part of that can be traced back to computing power. If engineers can't run an analysis on their workstations, and have to wait to get time on the enterprise cluster, then they can't properly tune a design. The result is a design that meets specifications, but isn't necessarily the best given the time and money invested.

One answer is the workstation cluster. This cluster shares computing resources with the engineer using the workstation, but not in a way that would slow down the computer for the engineer. Instead, the workstation user has enough resources to create the design, while unused processing power and memory goes toward running simulations on the cluster.

HP Z800 Workstations using virtualization software from Parallels can make this happen. The engineer can make use of memory, processors, processor cores, and storage for design work, while dedicating the rest of those resources to the cluster. The Z800 Workstations can be devoted to two simultaneous tasks — creating the design interactively, and analyzing parts of that design.

Building a Workstation Cluster

Here's how it works. Parallels Workstation Extreme, running on the HP Z800 Workstation package, enables engineers to create a virtual machine that can partition off memory, processor cores, and disk space for separate tasks. Because it uses Intel's VT-d technology for direct I/O, it offers virtual machines with performance comparable to that of native applications.

Engineers can take a group of HP Z800 Workstations, split each one into two virtual machines, and use the second Ethernet port to create a private network, clustering together several of these virtual machines. The result is an inexpensive but powerful

workstation compute cluster that can run analyses, or simulations of parts, for their design while it is in process.

Applications such as finite element analysis and computational fluid dynamics from vendors like ANSYS and Dassault Systèmes are ready to take advantage of these workstation clusters. they can use all available processor cores in the virtual machines. In many cases, these vendors also have licensing models that enable software to easily use all available processor cores for execution.

How does this approach help engineering teams keep to their schedules while delivering great designs? The workstation cluster enables engineers to quickly get feedback on design features, without the necessity of scheduling time on the enterprise cluster. The team can get the results of a simulation within a day or two, and in many cases minutes or hours, and try out different design alternatives to determine the best approach. At the end of the project, the completed design can be simulated on the enterprise cluster at a higher level of detail in order to thoroughly validate it.

In short, engineers have a greater ability to try out design concepts they might not ever get an opportunity to otherwise, and can gradually refine parts of a design based on rapid feedback.

You don't have to settle for good enough. Thanks to the HP Z800 Workstation with two Intel® Xeon® 5600 series processors, up to 192GB of memory, up to 12 processor cores, and dual Ethernet interfaces, even small engineering groups can have a cluster at their disposal. Engineers can keep on schedule, while ensuring that the resulting design is the best possible. Engineering design doesn't get much better than this. **DE**

For more information on an HP and ANSYS Solutions please go to www.hp.com/go/solver or call 800-888-0261.

Designing New Antennas

Researchers help space-challenged RF designers with smaller antennas.

BY JOHN EDWARDS

Like other mobile device antenna developers, Dr. Jennifer Bernhard, a professor of electrical and computer engineering at the University of Illinois at Urbana-Champaign, faced a fundamental dilemma: how to shrink an antenna's physical size without sacrificing major amounts of gain, efficiency or bandwidth.

"Most of the time, when you make an antenna small, you make a bad antenna," she says.

Bernhard's solution was to think flexibly and to use a 3D printing technique to create antennas that fit around a tiny, space-saving hemispherical substrate. Bernhard notes that antennas created with the process are electrically small in relation to the wavelength—usually one-twelfth of a wavelength or less—and feature performance metrics that are an order of magnitude better than those provided by conventional monopole antenna designs.

"It's a great solution to a significant problem," she says.

Addressing the Problem
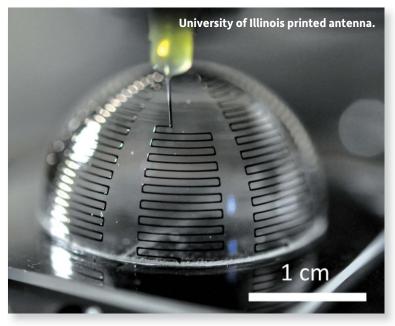
With mobile devices evolving into ever-smaller form factors (SFFs), antenna developers are investigating a growing array of innovative design concepts. They're tapping into 3D printing, exotic materials, fresh fabrication techniques and other innovations that promise to give designers additional choices when specifying antennas for smartphones, tablets and other mobile devices.

Yet as researchers push antenna concepts far beyond existing parameters, they're also confronting long-standing radio frequency (RF)-related design challenges, including the big one: size and shape versus performance.

"The basic job of any antenna is to efficiently send energy into free space," says Dr. Richard Ziolkowski, a professor of electrical and computer engineering and optical sciences at the University of Arizona in Tucson, who is working on his own cutting-edge mobile antenna research. "Yet, with antennas that are tiny compared to the wavelength, most of the signal is simply bounced back to its source."

Overcoming this basic obstacle is fundamental to the work of all researchers exploring new antenna design concepts.

"The goal, of course, is to make the antenna as efficient as possible within the space allowed," says Bernhard. "That's the essential thing."

University of Illinois printed antenna.

1 cm

3D Printing

Bernhard says that she and her team believe 3D printing can allow small antennas to radiate signals effectively, while minimizing the ratio of energy stored to the energy radiated—the "Q" of an electrically small antenna.
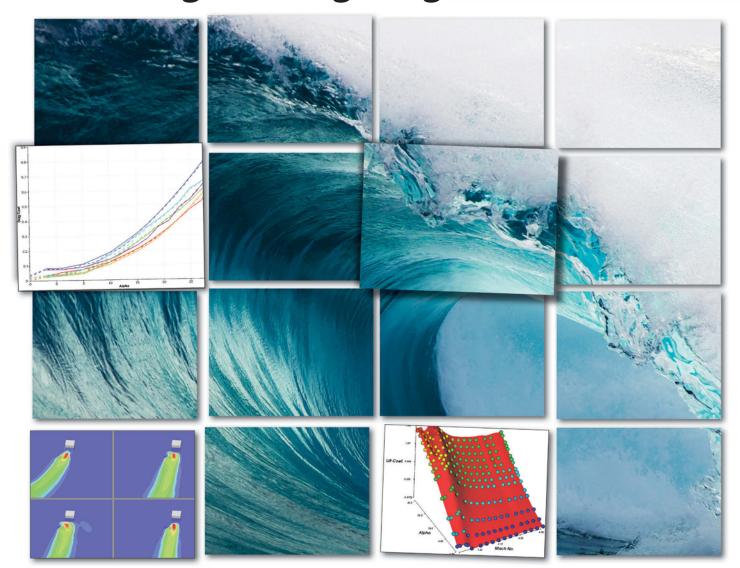
"Printing directly on that hemispherical substrate gives us a highly versatile, single-mode antenna with a Q that very closely approaches the fundamental limit dictated by physics, the Chu limit," she explains.

Building the antenna is relatively simple. Pyrex hemisphere substrates are printed with a conductive ink. Conformal printing can create the meander line antenna on the inside or outside of the substrate for flexibility. Bernhard notes that the approach can be used for an antenna integrated into a mobile device, the inside of a carrying case, or even the wing of an unmanned aerial vehicle. The antenna's operating frequency is determined by the printed conductor cross, as well as the spacing between the meander lines.

Bernhard says that engineers using software design tools to test and evaluate tiny antennas inside SFF platforms actually have an advantage over colleagues working with larger antennas and platforms.

"Once the platforms get larger, there are fewer software products that can do a good job of accurately predicting what the effect of an antenna would be on a larger package, such as an airplane," she says. "Those simulations are complicated, and can take a long time to run."

Metamaterials

Another approach to creating tiny antennas is to build them out of materials that have properties capable of boosting device efficiency despite size and shape constraints. To this end, National Institute of Standards and Technology (NIST) engineers are collaborating with colleagues at the University of Arizona and Boeing to create antennas incorporating metamaterials—materials featuring microscopic structures offering unusual properties.

A metamaterial can make an antenna behave as if it were much larger than it really is, Ziolkowski says: "The antenna structure stores energy and then re-radiates it." The approach eliminates the need to add bulky matching network components to increase efficiency. Better yet, Ziolkowski notes, metamaterials are more "frequency agile," allowing designers to use them on virtually any frequency on the fly.

Metamaterial antennas can radiate as much as 95% of an input signal, while defying conventional design parameters, Ziolkowski says. He notes that the prototype metamaterial antennas are as small as one-fiftieth of a wavelength—and could shrink further.

Ziolkowski says his challenges in developing metamaterial antennas were similar to those faced by antenna researchers investigating other technologies and processes.

"Even with metamaterials, the tradeoffs are always among small size, large bandwidth, high directivity and efficiency," he says. "You're still fighting to get the best results in all of these areas simultaneously and, because of fundamental physics, it remains a challenge."
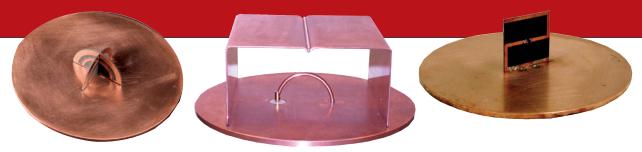
Ziolkowski says that his research team uses ANSYS HFSS, a 3D full-wave electromagnetic field simulation tool, to analyze and tweak their antennas' performance characteristics. He notes that they use the software as a frequency structure simulator to evaluate antenna performance across the spectrum.

Selecting the appropriate antenna design and analysis software, Ziolkowski says, is mostly a matter of personal preference and taking the time to choose a tool with features that are appropriate to the work at hand. He notes that designers have a wide range of software packages from which to choose, including CST Microwave Studio, COMSOL Multiphysics, GEMS and XFdtd, to name a few.

While software vendors take different approaches to antenna design and analysis tasks, Ziolkowski notes that virtually all of the available tools can help designers reach their intended goals.

"All of these software packages look at the performance of the electromagnetic wave and how the antenna is active, coupled with other systems," he says. "With HFSS, for example, you actually incorporate circuit behaviors and portions of the antenna with the circuit itself—called co-design—so that you can actually do both the antenna and the circuit analysis together."

University of Arizona metamaterial antennas.

Liquid Metal Fabrication

At North Carolina State University, researchers are investigating a new antenna fabrication technology that's based on liquid metal injected into elastomeric microchannels. The wirelike antennas can be bent and contorted without snapping because the device's mechanical properties are governed by the elastomer, not the metal.

"Flexibility is a great attribute to have in an antenna, since the antenna's frequency is determined by its shape," says Dr. Michael Dickey, assistant professor of chemical and biomolecular engineering at North Carolina State University. "You can actually tune these antennas just by stretching them."

The antennas are created by injecting an alloy comprised of gallium and indium, which exists in liquid form at room temperature, into hollow channels about the width of a human hair. Once the alloy fills the channel, the alloy's surface oxidizes, forming a "skin." The skin holds the alloy in place, allowing it to retain its liquid properties.

"The alloy essentially remains a liquid," Dickey says. "The alloy assumes the mechanical properties of the material encasing it."

Dickey says a design and simulation tool, like CST Microwave Studio, could be used to integrate a wirelike antenna into a smartphone or similar device.
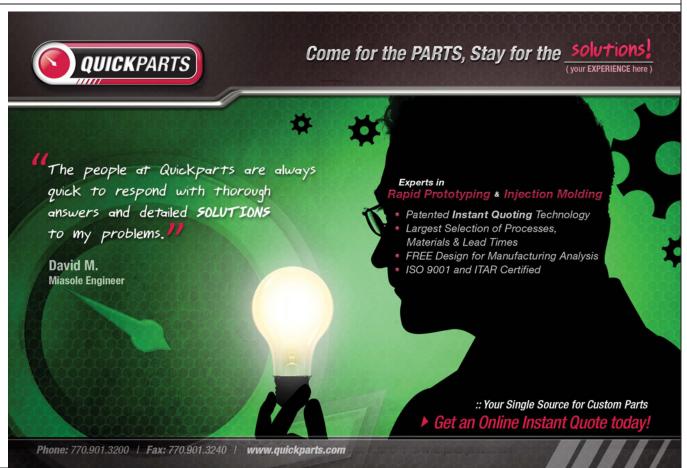
"From a design perspective, the engineer is looking at how informal an antenna can be for a consumer product," he says.

Dickey notes that over the past few years, aesthetic design has become a major consideration in mobile device development—almost a top priority.

"Manufacturers and consumers certainly don't want something like an antenna sticking out of their box," he says. "That's where having good software, and an antenna formed with liquid metal inside a plastic, can be an advantage."

Eliminating Antennas

At the U.S. government's Oak Ridge National Laboratory in Tennessee, researchers are investigating the possibility of doing away with antennas entirely—at least antennas in the

conventional sense of the term. Instead of a conventional metal or dipole antenna, future mobile devices might incorporate an antenna no bigger than a gnat.

Such an antenna might be made possible by a design that allows an electrically charged nano-mechanical oscillator to be tuned to specific electromagnetic waves. The technology detects small electric fields over large frequency ranges while maintaining power efficiency, making it potentially useful in a wide range of mobile communication devices.

"Having to match the antenna length to the wavelength may no longer be necessary," says Dr. Panos Datskos, the technology's co-developer and the lab's nanosystems and structures group leader.

Datskos says that modeling and simulation software played a major role in helping his team develop the technology, as well as measure results. The team chose COMSOL Multiphysics, although a variety of other products could have also been used.

"It has magnetic modules, and that's what we used to simulate the electrical distribution, the magnetic interactions—where the charges should be—and to create a more sensitive antenna," he explains.

The Future

Bernhard notes that innovative antennas, combined with effective design tools, provide a potent one-two punch that will help designers cope with growing pressures to develop smaller and more feature-packed smartphones and other radio-equipped devices. But she predicts that design know-how and talent will continue to be important factors in mobile device development for many more years.

"Technology alone won't get you to an optimum design," she observes. "Most tools, for instance, have some kind of optimizing function in them, but it's important be able to set up the optimizer properly. It's very easy to get answers that are incorrect due to the inexperience of the user."

Ziolkowski agrees. "You start with the technology and the tools," he says. "The rest is up to the designer." DE

John Edwards *is a technology writer based in Gilbert, AZ. His work has appeared in* IEEE Signal Processing, Electronic Design *and other publications. Contact him at jedwards@johnedwardsmedia.com.*

INFO ➜ ANSYS: ANSYS.com

➜ **CST:** CST.com

➜ **COMSOL:** Comsol.com

➜ **GEMS:** 2comu.com

➜ **Remcom:** Remcom.com

For more information on this topic, visit deskeng.com.

Modeling Fuel Cells

Simulation has become crucial to designing and implementing the next generations of fuel cells.

BY MARK CLARKSON
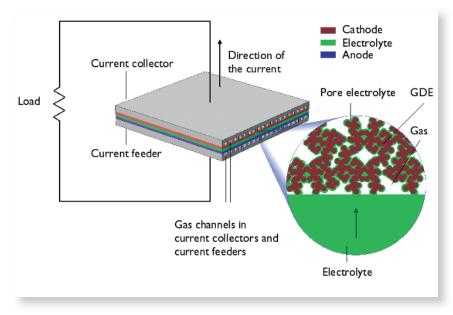
At first glance, fuel cells aren't very complicated. Hydrogen gas (H_2) is fed in under pressure through the anode to a catalyst, usually platinum, which splits it into positively charged hydrogen ions and free electrons. The anode conducts the electrons to an outside circuit to power a motor, etc., before they return to the cathode side of the fuel cell.

On the cathode side, oxygen (O_2) gas is fed under pressure to the catalyst, where the molecules are split into negatively charged oxygen atoms. These combine with the returning free electrons and positively charged hydrogen ions—drawn from the other half of the cell through a one-way membrane—to form water (H_2O.)

While the chemical reaction behind a fuel cell is simplicity itself, modeling one at any significant level of detail is a challenge. The anode and cathode use intricate channels to deliver their respective gases to the catalyst. The catalyst itself is a porous, composite material with an extremely convoluted or "rough" surface, to maximize the exposed surface area. To achieve reasonable voltages, multiple fuel cells are stacked together, connected with bipolar plates.

Any little bit can be surprisingly complicated to simulate. Take the basic matter of electrode thickness, for example.

"You have a tradeoff among conductivity, mass transport and current density distribution," explains Ed Fontes, CTO of COMSOL. "If you make a thick electrode, you get a more uniform current density distribution, so the bipolar plate design is less critical. On the other hand, you get more ohmic

A unit cell in a fuel cell (anode, cathode, electrolyte and current collectors and feeders). The image shows that the gas channels impede the transfer of current, while the parts of the current collector/feeder that are in electronic contact with the electrode impede the transport of gas to the electrode (optimization problem in design). GDE = gas diffusion electrode. *Image courtesy Mark Clarkson.*

losses and losses in mass transport. It's a highly non-linear problem to optimize the thickness of the electrode depending on the operating conditions, catalysts and optimal properties. Even 1D models can be quite complex."

Fully-coupled, Multi-scaled

What do you want to simulate? Are you doing transient analysis simulation to test various hypotheses about what's causing your real-world test results? Are you concerned with the smoothness of the gas flow? Are you modeling the electric current? Heat transfer?

There is, as they say, an app for that:

computational fluid dynamics (CFD), system modeling software, and so forth.

That's good, because a fuel cell design must balance the momentum of the fluids as they flow through the system, the energy transfer (heat, etc.); the transport of the chemical species (hydrogen, oxygen, water, electrons, etc.), and the electrical currents. That's hard enough, but these are all closely coupled.

"When you consume oxygen," says Fontes, "you're transferring electric current, so the electric current balances are coupled with mass transport and fluid flow. The losses in the system produce heat, so you have heat transfer."

Mathematical Battery Models

Dr. John McPhee of the University of Waterloo heads a research team sponsored by Toyota and Maplesoft to develop models of hybrid electric vehicles and sub-systems. Their work is on batteries rather than fuel cells, but still uses simulations of physics and chemistry to aid engineers.

Models developed by McPhee's team are math-based, rather than numeric.
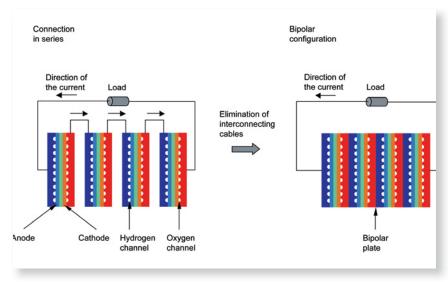
"With simulations based on finite elements or finite differences or CFD," says McPhee, "the math is embedded in the code. You can't see the equations; all you see are nice animations or plots or reams of data."

That's exactly what Toyota didn't want, he says.

"We created a pretty good circuit-based model of a battery," he recalls. "It matched experiments and wasn't too computationally intensive. The Toyota guys looked at it and said, 'Where's the chemistry?' So we went back to the drawing board. We now model the actual chemical kinetics, and the rate of chemical reactions within different battery chemistries."

McPhee is quick to point out that you don't have to know the math to work with the models. "The whole point of this three-way collaboration is for us to develop tools for MapleSim that Toyota can then use. We can't overwhelm the front line engineers with equations.

"An engineer can simply click-and-drag a NiMH battery into an electric vehicle simulation. They don't have to look at the math, unless they want to. The difference is that in MapleSim, you can."

The electrical connection principle behind a fuel cell stack. *Image courtesy Mark Clarkson.*

The performance of the electrodes and catalyst depends on the water content, so you'll need to take evaporation and condensation into account as well.

Simulating all this is doable on an engineering workstation with software such as COMSOL, but only a bit at a time. The cutting edge is to model it all simultaneously and fully coupled.

"To run just CFD isn't very interesting," says Fontes, "so you try to couple CFD and fluid flow with heat transfer, with chemical species transport, and with electrochemical processes [such as] electric and ionic currents."
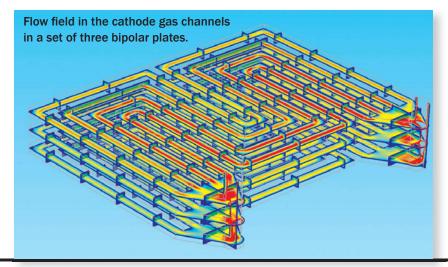
Not only do you want to have different coupled phenomena, Fontes adds, "you also want to have different scales—micro scale and macro scale. This is multi-scale modeling. That's the absolute cutting edge in this field. You can move all the way to molecular simulations for the catalyst. There are national labs doing that."
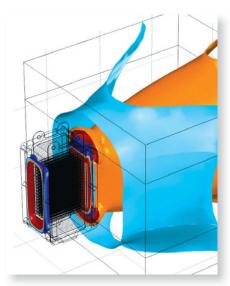
You won't be doing this on your desktop anytime soon. On the other hand, simulations you can run on your desktop today took a supercomputer just a decade or so ago.

Coming of Age

"At Ford, we rely heavily on modeling," says Dr. Dawn Bernardi, Ford research engineer. "Models are an integral part of [fuel cell and battery] development."

It wasn't always thus, she says.

Flow field in the cathode gas channels in a set of three bipolar plates.

Flow velocity iso-surfaces of cooling air obtained from a CFD simulation of an air-cooled stack. *Image courtesy Siegel Schleimer Ingénieurs-conseils, Luxembourg*

Different aspects of the flow velocity field (from the ambient into the cooling channels in the fuel cell stack).
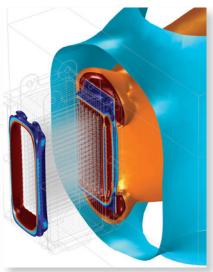
"We didn't get respect in the old days," says Bernardi. "People thought we were just 'fitting things in.' They'd say, 'You could fit an elephant with a 200-term polynomial!' They didn't realize that our equations were based on true phenomena, like diffusion and reaction and heat transfer. They actually did mean something about the system."

Today, she says, people are finally warming up to the concept because the models actually do work.

"We're able to rely on these models, because they reproduce all our life test results," Bernardi says. "Now we can do a year's worth of experiments, and then extrapolate out to 10 to 15 years of life. The challenge is simulating the degradative performance. Although our batteries are designed to outlive the vehicle, we still need to watch them reach their end of useful life."

Fuel cells don't last forever. They suffer corrosion, thermal cracking and other structure changes, not to mention abuse and accidents. Watching them age and die on computer is scientific research now, but the kind of simulation you can run on your workstation today was scientific research 10 or 15 years ago.

In another five to 10 years, this kind of information will be available in materials databases—just like the ones that inform today's software about electrical conductivity and tensile strength.

The problem with fuel cells is that they've been the "next big thing" for, well, decades. Are they ever going to actually break through?

"I got this question during my defense of my PhD [in 1995]," says Fontes. "I said, 'This particular one I'm looking at, the expectations are too high. It's never going to be as useable as people think.'"

That doesn't really matter though, he says. The same methods and software apply to any type of fuel cell, to batteries, to solar cells, and even to corrosion. **DE**

Contributing Editor **Mark Clarkson** *is DE's expert in visualization, computer animation, and graphics. His newest book is "Photoshop Elements by Example." Visit him on the web at markclarkson.com or send e-mail about this article to de-editors@deskeng.com.*

||||||||||||||||||||||||||||||||||||||||||||||||||||||||

INFO ➜ COMSOL: COMSOL.com

➜ **Maplesoft:** Maplesoft.com

For more information on this topic, visit deskeng.com.

A **New** Point of View

FieldView 13 CFD software makes its debut from Intelligent Light.

BY VINCE ADAMS

The size of computational fluid dynamics (CFD) models are increasing exponentially as manufacturers recognize the value of detailed transient flow data for temperature, flow resistance and other critical product data. However, one engineer told FieldView Product Manager Matt Godo that he was only able to post-process 1% of his CFD data because of the volume of information and the available time to interpret it. Herein lies the mission driving the developers at Intelligent Light in Rutherford, NJ.

Intelligent Light has been providing products and services for CFD analysts for 27 years. FieldView 13 is the culmination of five years of development—based on customer feedback as well as internal experience from its consulting organization. The release was guided by the following four principles:
1. Revolutionize CFD data management.
2. Provide unprecedented speed and performance.
3. Enable higher productivity sessions and workflows.
4. Don't require existing users to significantly adapt or adjust the way they have been working.
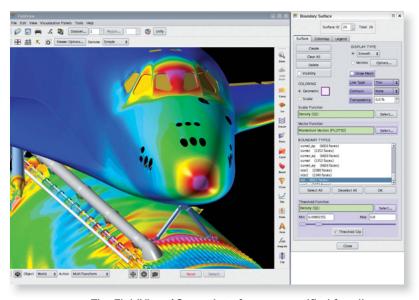
The FieldView 13 user interface, now unified for all platforms (Windows, Linux and Mac).

Setting New Standards

FieldView 13 contains a new graphics processing engine. It uses a many-core, multi-CPU design leveraging GPU acceleration for parallel rendering, multi-threading and hardware optimization to take advantage of all available cores. This release raises the bar for CFD data processing and the benefits scale with the power of the local or client system. The bigger and better the machine, the faster and more powerful the improvements in data visualization and analysis—with interactive operations running at least 20 times faster than the previous release of FieldView.

By default, FieldView displays high-quality renderings of overlapping transparent objects, with anti-aliasing at interactive frames rates. FieldView is designed to deliver the best performance on systems with the latest and most up-to-date graphics cards. It can also be configured to run on systems with more modest graphics resources. High quality rendering extends to batch or automated operation—images generated interactively or in batch are pixel-matched.

Another new feature that improves post-processing speed,

contributing to better interactive experiences for users, is Sweep Caching. When a user completes a streamline or particle path animation or a transient surface sweep, the post-processing objects for each step are stored in memory. Subsequent animations are played from memory for greatly increased frame rates. Free of the limitations of a movie, users can zoom, rotate and/or translate the dataset as it animates.

Breaking Down Barriers

Graphics processing speed can provide incredible improvements in efficiency, but it doesn't address one of the critical problems facing analysts as datasets increase. With high-performance computing (HPC) clusters and the anticipation of cloud computing, engineers are no longer solving their large, problematic models on their local computer, which is often a laptop in mobile, global design environments. Traditional CFD post-processing requires that results sets in the double-digit gigabyte or even terabyte range be copied across networks before the data can be viewed.

"Our current high-water mark is a transient model covering thousands of time steps, where each step contains over 2 billion cells," says Godo. "Even larger and more ambitious modeling efforts will be presented before the end of the year, fully relying on features delivered with FieldView 13."

A direct numerical simulation of a liquid fuel jet with 6 billion gridpoints is conducted to elucidate the multi-scale, turbulent physics of liquid fuel spray atomization. (*FieldView image courtesy of Dr. Matsuo, Japan Aerospace Exploration Agency, or JAXA*)

FieldView 13 addresses this by using reduced representations of the results of interest called Extract Database, or XDB, files. Users can create compact subsets of all the results data of interest on the server system so that a minimal amount of data gets pulled over the network for processing. The XDB file can be as much as 100 times smaller than the raw results dataset. Even if the raw data exceeds physical or realistic transfer or FTP limits, the XDB file will be portable. Then, once the dataset has been opened, it can be animated, sliced, probed and manipulated—with no loss in numerical accuracy, but at greatly enhanced speeds.
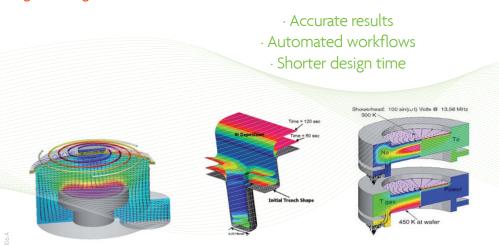
A single XDB file can be created for an entire model, or one can be made for each component or domain in a model. Separate, identically structured XDB files can be created for each step in a parameter sweep, or each time step in a transient analysis. These multiple XDB files can be processed simultaneously in Field-View 13 so that animations or plots across them can be generated. XDB files can also be created using different time steps in a transient model, so that fewer images of slower-moving segments of an event are captured, but resolution on the critical segments can be retained.

XDB files can also be created using batch operations that work remotely on the server, so that the analyst is freed up to work on other tasks. FieldView 13 calls this True Batch Processing. In addition to the productivity enhancements, True Batch also enhances data security. The database never leaves the server on which it was created, and the file system isn't manually accessed (with all the risk that entails). The automation tools in FieldView 13 come into play here.
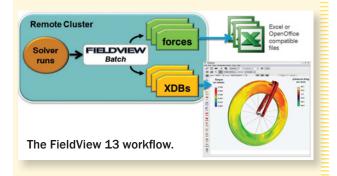
Unsteady Flow over Bicycle Wheel
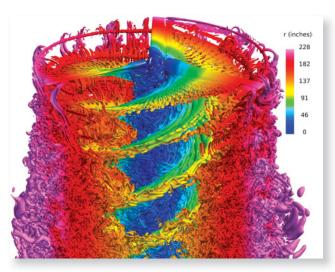
The power of FieldView 13 is best understood in the context of an actual case study. The bicycle racing wheel market is extremely competitive, with manufacturers banking up to $500,000 in development, including simulation and wind tunnel testing, for incremental improvements in stability, aerodynamics and performance (*Source: NYvelocity.com, "Zipp Lead Engineer Josh Poertner," Aug. 27, 2009*). The importance of CFD is growing rapidly for this industry—along with the sizes of the models being solved.

One manufacturer examines multiple geometry iterations each week for aerodynamic performance, using a transient solution to capture unstable loads a high-performance rider might experience. Using a 32-core, parallel processing HPC cluster, the raw results file size is 2.8GB per time step. Simulations with up to several hundred time steps are common.

An FVX script was developed to batch process results on the HPC using only eight cores while running four jobs concurrently. The script created Microsoft Excel-compatible files of the time-dependent forces and XDB files of the various components in the assembly, such as the wheel, fork and frame. The total XDB file size was about 60MB per step, which equates to a 46:1 reduction in processed file size for results visualization and analysis.

Furthermore, it took less time to read and animate the XDB files for all components across all time steps than it took to simply call a single time step's results, using standard interactive visualization in the CFD solver's native post-processor.

The FieldView 13 workflow.

High-resolution CFD simulation of V-22 Osprey rotor. Post-processed with FieldView on dual-CPU/8-core workstation. Multi-grid OVERFLOW simulation contains 14k grids. 668M grid points. (*FieldView 13 image used with permission: Chaderjian, N. M. and Buning, P.G., "High Resolution Navier-Stokes Simulation of Rotor Wakes."* Proceedings of the American Helicopter Society 67th Annual Forum, *Virginia Beach, VA. May 3-5, 2011.*)

Improving Productivity with Automation

The breakthrough in CFD data management through batch creation of XDB files was built on existing automation methodologies in FieldView. The most basic automation option is Restarts. A Restart stores the information needed by FieldView to recreate an image or plot composed by a user when it is saved.

Moving up one level of complexity, FieldView 13 has implemented a scripting capability to automate simple repetitive tasks. At the top of the automation hierarchy is FieldView FVX, a natural language programming tool that provides access to most of the functions within FieldView. The developers at Intelligent Light recognized that users didn't need to learn new or abbreviated syntax, so the calls within FVX capture words that are used in the FieldView graphical user interface. FVX can also be used to run existing Restarts or Scripts as the basis for more advanced automation.

Godo confirms that the significant gains in performance and productivity were achieved without disrupting the familiar environment and workflows in prior versions of FieldView. Version 13 is 100% backward compatible, he says, including scripts. DE

Vince Adams, *an account manager for LMS, is a simulation educator, consultant and speaker. He has authored three books on FEA and numerous magazine articles. Contact him via de-editors@deskeng.com.*

INFO ➡ **Intelligent Light:** iLight.com

Don't Settle for Good Enough with Workstation Clusters

Simulation via workstations clusters is an easy, affordable way to speed multiple design iterations.

BY PETER VARHOL

Until recently, design simulation has been limited to the realm of experts with access to HPC resources placed behind glass walls that are serviced by large chillers and immense amounts of electricity. Today, while you may still need experts for large, complex multi-modal simulations, the simulation of parts and small assemblies can be performed by a broader set of users. Independent software vendors like Autodesk, SolidWorks, ANSYS, Siemens PLM, Altair and others are making simulation-based design more accessible by providing step-by-step help for the novice. Now these users can create, test and modify design ideas in a virtual lab and reduce the number of expensive prototypes.

In tandem with the advances in software is an equally important set of advancements in hardware. Today, a workstation with two Intel® Xeon® processors looks and feels like a supercomputer just three years ago. When you cluster them together in the formation of a distributed cluster, they can deliver results at speeds you would not have thought possible.

The Formula for Building a Workstation Cluster

The hardware

• Build a workstation capable of doing design and simulation. Start with a workstation configured with two Intel® Xeon® processors, at least 24GB of memory a 2TB SATA hard disk drive, a 300GB solid-state drive and a professional graphic card.
• Have either a 1Ge or 10Ge network switch available to act as the distributed workstation cluster compute fabric.
• If you are planning on using a 10Ge switch, make sure your workstation has one 10Ge NIC.

The software

• Ensure Parallels Workstation Extreme (PWE) is installed on each workstation.
• Have a copy Windows and Linux (if you run your simulation applications in a non-Windows OS) operating systems. If you are using LINUX, PWE may bundle a copy of Adaptive Computing's workload management product, MOAB. If you are using Windows, you will need a single copy of Windows HPC Server.
Now you can build the distributed workstation cluster.

How Clusters Work

What makes a distributed Intel® Xeon® based workstation clusters possible? The lead item is hardware technology. Intel® Virtualization Technology for Directed I/O (Intel® VT-d for short) enables users to create hard partitions between resources of a workstation. That means CAD users will perform their operation at near native performance and simulation users will also have access to dedicated resources and near-native processor performance.

What does it mean for product development? According to Bob Deragisch, Manager, Enterprise Systems, at Parker Hannifin, it means that engineers "can expedite design decisions from applications like ANSYS, CFDesign, and LS-Dyna."

Engineering organizations will be able to:

1. Extend the benefits of high-performance computing to small or remote design sites.

2. Per Intel's white paper "Optimizing Engineering Productivity with Workstation Centered Computing:" provide performance equal to that of a local data center while reducing power utilization by up to 10%.

In short, a workstation in a distributed cluster can provide engineers greater ability to try designs they might not get to otherwise, and can gradually refine parts of a design based on rapid feedback on its strengths and limitations.

You don't have to settle for good enough. Thanks to the Intel-based workstations with up to 192GB of memory, up to 12 processor cores, even small engineering groups can have a cluster at their disposal. Engineers can keep on schedule, while ensuring that designs are the best possible. **DE**

INFO ➡ **Intel Corp:** intel.com/go/workstation

Hot New AM Materials

Customer requirements up the ante for new materials to address biocompatibility, strength, transparency, static shock and more.

BY SUSAN SMITH

What prototypes are made of comprises a very important part of manufacturing. The demand for new materials is unrelenting, and keeps additive manufacturing (AM) companies in constant competition to meet the need. In the past year alone, several new materials have been released, addressing such requirements as transparency, biocompatibility, static-dissipative capabilities and advanced composite materials. New research is also being done with 3D printing and tissue engineering.

Combining Temperature Resistance and Toughness

Bruce Bradshaw, director of marketing at Objet Geometries, says the company has debuted four different materials in the past year:

Objet's claim to fame is that it has developed Connex technology, which allows you to combine two materials. Although that's not in itself a new development, it is the basis for one of Objet's new materials: Acrylonitrile Butadiene Styrene (ABS)-like Digital Material (RGD5160-DM).

It helps allay a concern in manufacturing engineering plastics: the tradeoff between temperature resistance and toughness.

"You either get very high temperature resistance and the product is very brittle, or you get toughness and little or no resistance to high temperature," says Bradshaw. "Because we can take two materials and combine them together, our ABS-like material gives us the best of both worlds. By blending two materials together, we get the material properties of a high temperature material, in which ABS-like can go up to [203°F] after post-processing in an oven—and we get very high toughness."

ABS-like is used for form and function and break testing. One example is from an automotive company in Ger-

many that leather coats its steering wheels: The process requires a high temperature to leather coat, so high-temperature molding applications are needed. Another application example would be Kohler running hot water through its faucet designs.

Bradshaw noted there are 4,000 ABS materials in the manufacturing world. Objet offers 65 different materials with its Connex technology: 14 base materials and 51 digital materials. Newly introduced materials include:
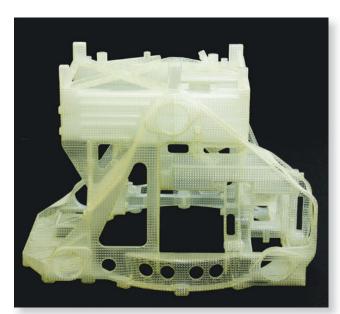
- **MED610:** The biocompatible material is designed for applications requiring prolonged skin contact of more than 30 days and short-term mucosal membrane contact of up to 24 hours. Examples might be Gillette testing toothbrush designs, or any company needing to check the customized fit of devices such as surgical orthopedic guides or dental delivery trays. MED610 is made in a clear material. It has five medical approvals of the standard ISO 10993-1: Cytotoxicity, Genotoxicity, Delayed Type Hypersensitivity, Irritation and USP Plastic Class VI.

- **VeroWhitePlus:** This material is more durable than its predecessor, VeroWhite. Plus is used in the prototyping world.

- **VeroClear:** This material is used for anything that requires transparency, such as lenses, plexiglass or lamps—and also in medical manufacturing for things like bottles.

Objet also offers three different groupings of printers: desktop printers, Eden family and Connex family. VeroWhitePlus works on all the printers; VeroClear and MED610 work on Eden and Connex; and ABS-like works on Connex only.

Investment Casting Materials

3D Systems' new products—Accura CastPro, Accura PEAK and VisiJet Clear—hit the market with their varying range of capabilities: from use in investment casting patterns to designing prototypes for windtunnel models. These materials are used to build models and prototypes for functional testing, proof of concept, dental, architecture and much more.

For use with stereolithography systems (SLA), Accura

ABOVE: 3D Systems' Accura CastPro is designed for the production of investment casting patterns for the QuickCast process. *Photo courtesy of 3D Systems.*

LEFT: Objet's new VeroClear can be used for anything that requires transparency. *Photo courtesy of Objet..*

CastPro is designed for the production of investment casting patterns for the QuickCast process. The material is dimensionally stable even in humid and moist environments, making it a good candidate for highly accurate master patterns. Its low viscosity ensures drainage of very thin walls, and minimizes the amount of residual material in a pattern that needs burnout. It also has low coefficient of thermal expansion, which ensures successful burnout of casting patterns, and a high productivity rate.

Accura PEAK is designed to withstand aggressive environments with rigid and stable parts. It is stiff and resistant to moisture, and can withstand temperatures over 300°F. It is a good choice for parts that must be accurate and stable over time, such as jigs, fixtures, master patterns/models, windtunnel models and higher-temperature applications.

VisiJet Clear was added to the VisiJet materials line for the ProJet 6000 Professional 3D printers using SLA part technology. VisiJet Clear complements the line by adding UV curable acrylic plastic with a polycarbonate look and feel—ideal for applications that require transparency. Models will allow visualization of internal structures or the study of fluid flow, and can be used for prototyping optical components such as lenses and lamps. The output for VisiJet Clear is strong and durable, according to the company.

Stephen D. Hanna, director of global sales and mar-

Parts made with ZCorp's zp150 cost just $3 per cubic inch of material, or about $20 for a typical handheld part.
Photo courtesy of Z Corp.

keting, rapid manufacturing materials, says VisiJet Clear meets the standards of USP Class VI: "This opens the door to a variety of medically related applications, and greatly expands the utility of the PJ6000 to these customers, allowing for the potential for design and development of medical instruments, dental models, medical drill guides, surgical study models, etc."

High Feature Detail and Soft Parts

Aslo new in the lineup of AM materials is the Iflex 500 from EnvisionTEC, for use on the Perfactory and ULTRA systems. EnvisionTEC is known for its photopolymers for manufacturing flexible soft industrial parts, and in that vein, the Iflex 500 produces soft, rubber-like parts with great tear resistance and UV stability that last for a long period of time. When properly light post-cured, the part can be used for underwater applications.

EnvisionTEC LS600 is an ABS-like photopolymer with high feature detail used on EnvisionTEC's 3D printers. LS600 is a material that produces parts with high impact resistance similar to that of thermoplastics. It can be used in prototypes for connector and electronic covers, automotive housings and dashboard assemblies, as well as consumer sporting goods.

Static Dissipative Material

Another contribution to the list of new materials this year is ABS-ESD7 from Stratasys. Fred Fischer, director of business development, notes that this electrostatic discharge (ESD) material is used by Fortus AM systems.

ABS-ESD7 is for applications where a static charge can damage products, impair performance or cause an explosion. It pre-

---

vents a buildup of static electricity. In addition, the ESD material eliminates the attraction and buildup of particulate, such as dust or powders, which are often attracted by static electricity and can degrade product performance. ABS-ESD7 is also valuable for designing products that avoid attracting atomized liquid, such as medicine inhalers, which must deliver the entire drug dose and not leave mist clinging to the inhaler's internal surfaces.

In general, all materials, whether plastics or metal, have a conductivity range. By definition, most plastics are insulators (they are not very conductive) whereas most metals are conductive. "There's a range of materials' conductivity, so if you look in the middle, there is anti-static, and static dissipative," notes Fischer.

The material's static dissipative properties means the surface and interior of a part built have very slight conductivity. Products that are sensitive to static shock—printed circuit boards or microprocessors, for example—can benefit from the static dissipative properties.

"We as humans can be the generators of static electricity, so when our feet rub against the ground, we are acting as a conduc-

Products that are sensitive to static shock, such as printed circuit boards or microprocessors, can benefit from the static dissipative properties of Stratasys' ABS-ESD7. *Photo courtesy of Stratasys.*
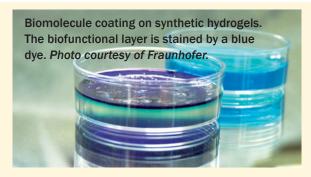
BLOOD VESSELS OF THE FUTURE

Biomolecule coating on synthetic hydrogels. The biofunctional layer is stained by a blue dye. *Photo courtesy of Fraunhofer.*

Tissue engineering researchers worldwide have been working on creating artificial tissue, and even entire organs, in the lab. Along those lines, a team at the Fraunhofer Gesellschaft in Germany has been experimenting with making artificial blood vessels on a 3D printer for use with transplants of lab-created organs. The challenge for tissue engineers has always been supplying artificial tissue with nutrients that must get to and from the organ via capillary vessels. The team believes they have solved that problem using 3D printing and a technique called multiphoton polymerization, developed in polymer science.

"The individual techniques are already functioning, and they are presently working in the test phase; the prototype for the combined system is being built," says Dr. Gunter Tovar, who heads the BioRap project at the Fraunhofer Institute for Interfacial Engineering and Biotechnology IGB in Stuttgart.

Tissue engineering is charged with printing a part of the vascular system, something as small and complex as a blood vessel that has tiny branches and spaces. The scientists at Fraunhofer Institute combine 3D printing technology with two-photon polymerization, which involves shining intense laser beams onto the material to stimulate the molecules.

The material then becomes an elastic solid, which allows researchers to create highly precise and flexible structures that can interact with human body tissue.

To keep the synthetic tubes from getting rejected by the living organism, researchers coat their walls with modified biomolecules. These biomolecules are part of the composition of the "inks" used in the blood vessel printer, combined with synthetic polymers.

"We are establishing a basis for applying rapid prototyping to elastic and organic biomaterials," Tovar says, noting that the 3D printer they use is a new prototype built by the Fraunhofer Gesellschaft.

The same technology may one day make it possible to create artificial tissue and complex organs.

tor and building up static charge," Fischer says. "What these static dissipative plastics do is absorb the static shock and dissipate it into the mass of plastic, so it becomes like a mini lightning rod. You can build products out of static dissipative thermal plastic material, and then the plastic is generally the most conductive and most likely object to absorb and dissipate that static shock. This way, you are protecting the components that are housed inside of a component.

"Historically, the previous plastics that we have were limited to non-static sensitive applications or products," Fischer points out. "This material is allowing us and our technology to be used in a broader spectrum of applications out on the manufacturing floor, in non-static sensitive products—and now in static sensitive products as well. We at Stratasys have some static controlled areas on our production floor, and just about every manufacturer I've walked through, either they or their suppliers have areas on their manufacturing floors that are static controlled areas."

ABS-ESD7 will run on Fortus 400mc and Fortus 900mc FDM systems.

Composite Material Offers Prototyping Performance

ZCorp's latest material is zp150. "Zp150 is the most advanced composite material system in our company's history, and far surpasses the performance of our previous materials," says Joe Titlow, vice president of product management.

Zp150 maintains ZCorp's claim of being the most affordable 3D printing system by costing just $3 per cubic inch of material, or about $20 for a typical handheld part. Parts made with zp150 can be finished differently—and for strong functional performance, can be infiltrated with Z-Max resin.

This material works with the entire line of ZPrinters from ZCorp, from the $14,900 ZPrinter 150 to the top-of-the-line ZPrinter 650. **DE**

Susan Smith *has been working as an editor and writer in the technology industry for more than 18 years. As an editor, she has been responsible for the launch of a number of technology trade publications, both in print and online. Currently, she is a contributing editor for* Desktop Engineering. *Contact her via de-editors@deskeng.com.*

INFO ➡ **3D Systems:** 3Dsystems.com

➡ **BioRap project:** IGB.Fraunhofer.de/en.html

➡ **EnvisionTEC:** EnvisionTEC.com

➡ **Objet Geometries:** Objet.com

➡ **Stratasys:** Stratasys.com

➡ **ZCorp:** ZCorp.com

For more information on this topic, visit deskeng.com.

Editor's Picks

by Anthony J. Lockwood

Each week, Tony Lockwood combs through dozens of new products to bring you the ones he thinks will help you do your job better, smarter and faster. Here are Lockwood's most recent musings about the products that have really grabbed his attention.
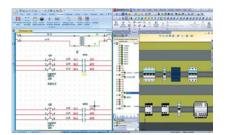
Print 3D Models with 14 Materials

More than 50 printing material options available for new Objet260 Connex 3D printer.

Design engineers have and will always see a lot of their work at all different stages—on screen and paper. But a growing number of leading-edge engineers also get to see and hold in their hands their evolving ideas. The tactile quality of holding your notions with nature's own hand is one of the beauties of 3D printing. And the idea of 3D printing just got a lot more beauty to it with the introduction of the Objet260 Connex 3D printer from Objet Geometries. With it, you can print a single 3D model made of up to 14 individual materials, according to the company. The materials can have different textures, colors and properties,

MORE ➜ deskeng.com/articles/aabbmn.htm

E-CAD Integrates with SolidWorks

Real-time processing/bidirectional updates between schematics and 3D views.

As its name implies, elecworks for SolidWorks from Trace Software integrates with SolidWorks. It's bidirectional integration, so that changes in either application propagate to the other automatically and instantly. That's important for a number of reasons, not the least of which is that you can see and fix problems right away in 3D instead of discovering them later during the autopsy of a charred prototype. And, yes, 3D. Schematics, of course, are in 2D, but elecworks lets you visualize individual wires and components in a 3D panel or machine assembly. It auto-routes in 3D SolidWorks files directly, and components are linked with the schematic BOM.

MORE ➜ deskeng.com/articles/aabbnh.htm
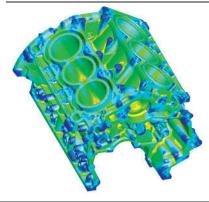
Stratasys Introduces Cross-Over 3D Printer

Fortus 250mc provides 3D prototyping and direct digital manufacturing in one unit.

The new Fortus 250mc Production 3D Printer is designed to deliver concept and functional prototypes as well as direct digital manufactured parts. What this thing is said to do is combine the ease of use of the Stratasys Dimension line of 3D prototyping systems with the control software used in its Fortus line. That software is key.

The software, called Insight, automatically prepares your STL part file for 3D printing so, like using a Dimension, you don't have to think much about it. But Insight also lets you do a Ross Perot: You can get under its hood and set control parameters for the Fortus additive build process.

MORE ➜ deskeng.com/articles/aabbpj.htm

Flow Science Releases FLOW-3D Version 10.0

Newest edition of CFD software offers FSI and thermal stress evolution models.

Flow Science released version 10 of FLOW-3D computational fluid dynamics (CFD) software recently, and it sounds pretty interesting.

For those of you who are not familiar with the application, the company describes FLOW-3D as a CFD system to use in your design phase and to improve production processes. It's engineered as an all-inclusive package, meaning that you do not need special modules for meshing or post-processing. Its integrated user interface is said to tie everything from problem setup to post-processing together. It's available in a distributed-memory version, and there are specialized versions as well.

MORE ➜ deskeng.com/articles/aabbsa.htm

1 MSC Announces Marc Simulation Software

MSC Software (mscsoftware.com) has announced the new Marc simulation solution, which includes a new user interface, and improved CAD interoperability and meshing features. It is designed to make it easier for engineers to create finite element analysis (FEA) models. According to the company, it enables smarter setup of nonlinear contact problems. For multiphysics analyses like heat transfer and electrostatics, it enables engineers to use a "perfect" glue approach for contact between bodies. Several new electromagnetic tetrahedral and triangular elements are also available.

Femap 10.3 Released

Siemens PLM Software (siemens.com/plm) has released a new version of Femap pre- and post-processor for engineering finite element analysis (FEA) running on Windows. Enhancements in Femap 10.3 are designed to make it easier for end users to work with complex geometry models. A new user interface feature further extends support for functionality in the NX Nastran FEA application. Features include a new automatic geometry preparation feature that reduces the time necessary to prepare complex geometry for FEM preprocessing (meshing); an enhanced user interface feature that supports interactive model set-up for the aeroelastic analysis capability of NX Nastran; and the Femap Postprocessing Toolbox has been extended to include the Free Body Tool.

2 View XVL Files on the iPad

Lattice Technology's (lattice3d.com) iXVL View application for the iPad will read previously captured/converted XVL files, displaying 3D models that can be rotated, panned and zoomed using the fingertip controls common in iPad applications. XVL's compression will allow large CAD files to be viewed on all iPad devices. The free software is available from the Apple App Store.

Delcam Launches FeatureCAM 2012

Delcom (delcam.tv/fc2012/lz/) has launched FeatureCAM 2012. The new release includes numerous keyboard shortcuts. It is also easier to select multiple items from the operations tree and to move all the features in a group.Toolpaths are displayed more clearly, selection of toolpaths has been simplified and an "undo" option added. The release also includes back boring and spiral roughing strategies.

NI Unleases the Watchdog

National Instruments' (ni.com) Watchdog Agent Prognostics Toolkit for NI LabVIEW system design software expands the capabilities of LabVIEW in machine prognostics and health management (PHM) applications. It is intended to give engineers a ready-to-run prognostics solution that can increase engineering efficiency for developing any PHM application.

Omega Introduces Thermal Imager

The OSXL160 thermal imager from Omega (omega.com) has a temperature range of -20 to 1,200°C (-4 to 2,192°F). Special features include voice annotation, analysis software, FLASH memory storage, USB connectivity to a PC, and a colored thermal image with sounds and alarms.

3 Eurocom Expands Notebook Display Options

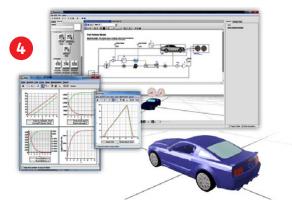Eurocom (eurocom.com) has increased the 15.6 and 17.3 LCD display choices for its notebooks, and now offers the displays as replacement options for users who want to upgrade. Display options include matte, glossy, and up to 95% NTSC. Eurocom offers display choices from several manufacturers: AU Optronics (AUO), LG Electronics and HannStar.

Take 17TB of Storage with You

NextComputing's (nextcomputing.com) Radius EX and Vigor EX portable workstations now support up to 17 fixed SAS/SATA 2.5-inch hard drives or 12 hand-removable SAS/SATA drives. With drives available up to 1TB each, users will be able to create a mobile workstation with up to 17TB of internal storage.

4 Maplesoft Releases Component Library

Maplesoft (maplesoft.com) has released MapleSim Driveline Component Library, a physical modeling and simulation tool built on a foundation of symbolic computation technology. It handles all of the mathematics involved in the development of engineering models, including multi-domain systems, plant modeling and control design. By using MapleSim with the MapleSim Driveline Component Library, engineers can mix physical models and empirical data. The compo-

nents of the library use an acausal approach to model development where the user can connect the components in the required configuration without worrying about issues like torque direction and load-flow. These issues are resolved automatically by the underlying solver.

New 3D Printers Announced
EnvisionTEC's (envisiontec.de) new Perfactory 3D printers have a projector resolution of 1920 x 1200 pixels. At the heart of all EnvisionTEC systems lies Texas Instruments' DLP technology (dlp.com), which, with its Digital Micro-Mirror Device (DMD), can expose each entire layer in one pass. This technology is intended to increase build speeds without sacrificing part accuracy or detail.

5 Add Rendering Decals
SimLab Soft (simlab-soft.com) has released SimLab Composer 2.4, which introduces a decal system that enables the user to add decals to single or multiple faces. Created decals are viewable in generated 3D PDF files. Version 2.4 focuses on material management with new capabilities and a revised GUI that assists the user with the rendering process by combining identical materials, and changing properties of multiple materials simultaneously.

3D Scanning from Afar
FARO Technologies' (faro.com) Focus3D laser scanner includes the new Multi-Sensor hardware feature, which incorporates a compass and height sensor to complement the device's existing dual-axis compensator. The new sensors define height and orientation against a set point for each scan. Users can now perform and download scans remotely with the new WLAN remote control included in each Focus3D. The device may also be equipped with one of three application-specific adaptors. **DE**

Advertising Index /// Publisher does not assume any liability for index errors or omissions.

GET SMART
with Embedded Software

Unlike a lot of media campaigns, IBM's "Smarter Planet" ads feature a basic truth behind them: A high proportion of the products and devices we interact with are getting smarter. The proposition is that these will enable a decent standard of living for the soon-to-be 7 billion people on Earth to live sustainably within the resources of the planet.

Of course, technology only enables. The real drivers are people's needs, cultures and organizational capabilities. The Internet hosts new types of organizations, accelerating the trend to inter-connect us. So far, it links people to content and other people. In the future, though, it will link people to things—and even link things to things.

Smart devices are products and equipment that are a combination of mechanical, electrical, semiconductor and software

> **Smart products change the design process and the role of engineers.**

elements. It is the embedded software that controls these devices and enables the increasingly complex interactions between people and devices. But a consequence of this is that the design process to develop smart devices cannot simply break the smart product into independent parallel designs by engineers in each discipline for each element. The elements interact.

A function to meet a customer requirement, such as triggering an actuator, might be realized in mechanical, electrical or software controls. The design trade-offs are complex, and extend beyond the traditional engineering function. The unit volume, the unit price, the power consumption, the time to market and the flexibility to change the product capability all interact.

There is also a time dimension to the design trade-offs. As volume increases, implementing a function in a different discipline might reduce cost. Downloading new software to an existing product in the field, creating new functionality, might generate post-sales revenue and maintain customer loyalty.

One group that will be affected first by these changes are design engineers. Smart products change the design process

and the role of the engineers. Fortunately, engineers are highly skilled and adaptable people! But they will need smarter tools.

Looking Back to Look Ahead

When I was an engineering student, more than 40 years ago, undergraduate courses were for specific disciplines: electrical, mechanical, electronics. We did not expect to need sales or commercial skills. Our success would be measured in terms of product innovation—introducing new products to the market. Selling them profitably would be somebody else's problem.

But when we were unleashed onto the industry, we quickly found out that making sure our innovative product met a market demand turned out to be our problem, too. Many of us acquired commercial skills on the job.

Since I studied, the narrow course I took has transformed to a comprehensive grounding in engineering, management and economic principles. Company-based projects are a significant feature of the course, enabling students to apply what they have learned to real industrial problems. The head of the Institute of Manufacturing at Cambridge University, Dr. Mike Gregory, sets out the objective as "combining intellectual rigour with real-world relevance with an integrated community of academics, students and industrialists with a shared passion for modern manufacturing."

Real-world problems turn out to be much more complex than excellence in design in one's own discipline. Engineers have developed techniques to handle complexity and interactions, such as multi-disciplinary design teams. A systems engineering approach has become common. The International Council for Systems Engineering (INCOSE) defines systems engineering as "an interdisciplinary approach and means to enable the realization of successful systems. It focuses on defining customer needs and required functionality early in the development cycle, documenting requirements, then proceeding with design synthesis and system validation while considering the complete problem."

Software Development Becomes Key

Smart products mean that all product development engineers will have to engage more effectively with one particular dis-

Cases in Point

Better visualization of virtual products means that future potential users can understand the look and feel of the physical product from a realistic simulation. For example, BMW Research has already shown how this technique would assist a technician to dismantle an engine. Images of the parts to remove are registered over the user's view of the physical world. These "augmented reality" techniques could validate requirements specifications and assist making better estimates of the demand for the product. Better customer feedback is likely to improve products, but encourage late requirements changes. To cope with this, we will need a better design process.

Dassault Systèmes demonstrated technologies for augmented reality at the recent Application Innovation Summit in Paris. Its vision is that product design and the software tools it supplies to assist engineers will change fundamentally. The company expects use of these tools to extend to most skilled workers in a manufacturing enterprise.

At present, support for embedded software development is poorly integrated with mechanical and electrical engineering design and data management tools. However, we are beginning to see a more integrated approach from the bottom up.

For example, Microvisk Technologies' handheld medical device provides home tests for patients using anti-coagulant treatments. It demonstrates the kind of problem engineers will face. This device measures viscosity using microelectromechanical sensors (MEMS) on disposable strips incorporating a small cantilever. A finger prick of blood is fed to the microchip by a micro capillary channel. Tissue factor is introduced to begin a reaction known as the Clotting Cascade. The change from a free-flowing to a gel-like substance is detected; the cantilever deforms and the result of the test displayed.

This is an extremely challenging and complex device to design. The industry is tightly regulated. Material science, mechanical, microelectronics and software elements all interact. A simulation-led design process has allowed Microvisk's design team to reduce physical prototyping and development time by integrating several disciplines' analysis and simulation. Microvisk used COMSOL Multiphysics to simulate the MEMS' performance mechanically, thermally and electrostatically.

This example is just one step in what we expect to be a gradual change to smarter software tools that support a changing design process.

cipline: software development. Traditionally, engineers have always been able to program computers a bit. But they considered this as supporting their discipline, not an integral part of the product they are designing. Embedded software as an element of a product must be developed with the same professional engineering techniques. Engineers from all disciplines will have to understand these constraints, and how product functionality realized in software interplays with their own discipline.

Smart products add flexibility to products, but complexity to the design process. It can be very difficult to track which requirements are satisfied where. For smart products, a successful design process will need further emphasis on requirements definition; modeling the design at successive stages of refinement; simulation and analysis of a virtual product; and test and verification that the design requirements are met as we replace the virtual product design with physical realization.

Many engineering software tools still only work within a single discipline. What is needed are tools that comprehensively support all the different disciplines, and allow designs to be verified both at the detail of the problem and as a complete solution to the requirements. **DE**
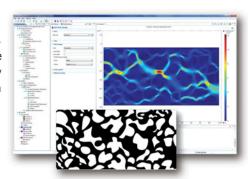

Mike Evans *is research director of Cambashi Ltd, based in Cambridge, UK. Contact him via de-editors@deskeng.com.*
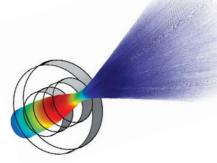
Image to Material Conversion

COMSOL Inc. has unveiled COMSOL Multiphysics version 4.2a, which includes a number of new capabilities. Image-to-material conversion can shorten time to solution for image-based simulation by allowing users to bypass both geometry creation and computational-demanding meshing of microscopic details. They can use image data to represent 2D material distributions or to identify regions with different materials by their color or gray scale. Images used in this way can have many origins such as scanning electron microscope (SEM), computed tomography (CT), or magnetic resonance imaging (MRI).

Particle Tracing
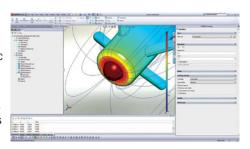
The Particle Tracing Module extends the functionality of the COMSOL Multiphysics environment for computing the trajectory of particles in a fluid or electromagnetic field, including particle-field interactions. Any COMSOL add-on module combines with the Particle Tracing Module, giving users access to additional modeling tools and fields to drive the particle motion. Applications include flow visualization, mixing, spraying, particle separation, mass spectrometry, ion optics, beam physics, ion energy distribution functions, acoustic streaming, and ray tracing.

LiveLink for Creo Parametric

The new expansion of the COMSOL LiveLink CAD interoperability products now embrace Creo Parametric design software from PTC. By establishing an associative connection between the two applications, any changes to a feature in the 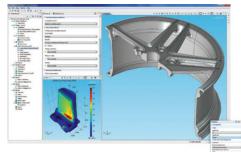Creo Parametric CAD model automatically update the geometry in COMSOL Multiphysics while retaining physics settings. The LiveLink for Creo Parametric includes all the capabilities of the COMSOL CAD Import Module and enables import and defeaturing of CAD files from major CAD packages.

New Default Kernel

The Parasolid geometry kernel from Siemens PLM Software is now the default geometry kernel for users of the CAD Import Module and the LiveLink products for CAD. Parasolid enables the handling of more advanced geometry objects for users of LiveLink products, which include versions for AutoCAD, Inventor, Creo Parametric, Pro/ENGINEER, SolidWorks, and SpaceClaim.

➜ **DO YOU HAVE A COOL TOOL?** *Desktop Engineering* editors are always on the hunt for great tools to share with our readers. Send us your ideas at de-editors@deskeng.com.

HIGHLIGHTS

COMSOL Multiphysics Version 4.2a is a significant upgrade, which also features:

- Digital Elevation Map (DEM) import of topographical surface data for applications such as fluid flow, structural, or electromagnetics.

- Post-processing of simulation results using interactive slice and isosurface plots with slider control.

- A new user interface for large parametric sweeps and quick visualization of response surface plots.

- A new electromechanics physics interface in the MEMS Module combines solid mechanics and electrostatics with a moving mesh to model the deformation of electrostatically actuated structures.

For detailed information on COMSOL Multiphysics version 4.2a please visit *www.comsol.com/4.2a*

NOW I THINK IN 3D.
AND MY DESIGNS JUST KEEP GETTING BETTER.

With our Dimension® 3D Printer, I know my model will represent my idea exactly. And that makes it easier to improve my design with each iteration. Our Dimension is right here in the office, and that helps us get our products to market faster. The Dimension models we create are made in ABS, so they're tough and durable. Leases start at only $275 a month, or you can buy a printer for under $15,000, so they're also really affordable. Overall, Dimension gives me an amazing sense of freedom—and creativity.

Find out more at www.dimensionprinting.com/de4

dimension®
3D Printers

STRATASYS®