November 2010 / deskeng.com

DE Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

The HPC Revolution

Parallel Processing Power

P. 16

SPECIAL SECTION
FOCUS ON Rapid Tech

What's New at omega.com

Please Call to Receive Your FREE Dilbert™ CONTROLCAT New Horizons® in Temperature, Pressure, Flow and More!

or Visit Our Web Site: omega.com/literature

HOT OFF THE PRESS!

Over 150 Full Color Pages of State-of-the Art Products!

- Automation and Control
- Calibrators: Temperature & Pressure
- Chart Recorders
- Data Acquisition Systems
- Data Loggers
- Flowmeters
- Heaters: Flexible, Immersion, & Strip
- Infrared Handheld Instruments
- Load Cells
- pH and Conductivity Instrumentation
- Pressure Transducers
- Photoelectric and Proximity Sensors
- RTD Probes and Thermometers
- Sanitary Temperature and Pressure Sensors, Transducers
- Signal Conditioners
- Strain Gages
- Thermocouples
- Terminal Blocks
- Transmitters
- USB Probes and DAS Systems
- Wind Tunnels
- Wireless Sensors and Meters

Includes 105* DILBERT *Cartoons!

Visit omega.com/dilbert for your daily dose of DILBERT!

For Sales and Service, Call TOLL FREE

1-800-82-66342

1-800-TC-OMEGA®

Shop Online at

omega.com®

OMEGA®

by Steve Robbins

The Design Revolution

This issue is titled "The HPC Revolution," but the revolution happening in the design engineering team's world is much bigger than just computational power. Don't get me wrong, the speed of our computing systems are enabling everything we do, but the tools we use are software based, and engineering processes are changing as our tools get more intuitive and powerful.

The design revolution is changing the way we work. The decisions made at the very beginning, starting with the conceptual design, are influencing every aspect of the design process down to the manufacturing floor. Design engineers are now able to simulate their models on the plant's automation equipment, validating that the designs will be able to be manufactured.

The design revolution taking place is based on changing from serial to parallel processes.

The Parallel Process

As the design team works together to build CAD models, run analysis and simulations, configure systems, and collaborate with their team members, the design becomes increasingly complex. The design revolution taking place is based on changing from serial to parallel processes, and on the organization of large amounts of engineering data that is generated. This data needs to be accessible to team members—the right data to the right engineer—instantaneously, wherever they are working. With high-performance computing and communication networks becoming increasingly available, concurrent processes enable quick and rapid change. Quick change compresses the design and manufacturing cycles. The design revolution will not just organize data, it will help us comprehend it. It will let us know what we need to do next.

One of the challenges this presents is changing the way we think. Most of us have been trained to think serially. This was natural when a design was passed from the mechanical engineer, to the electrical engineer, to the manufacturing engineer. But in the new paradigm, models will be revised as analysis and simulations occur simultaneously and the design is being virtually manufactured to reduce time to market. CAE applications will not just create a visual post-processing report. They will suggest design changes with the data to back them up. Simulation lifecycle management software will mine analysis data for factors that affect all aspects of the design and manufacturing process.

The simulation of the mechanical, electronic, and software systems shortens the design process cycle and creates a system where the components are designed and specified before the manufacturing process. The software code developed for a mechatronic simulation is reused for the completed design. Mechatronics promotes electronic design automation, creating specifications for electronic components at the beginning of the design cycle. It is a holistic approach to virtual prototyping, and we can expect increasingly larger, more complex models virtually prototyped as we continue to make gains in software development and its use of multi-core computing.

The Revolution Will Be Accessible

The revolution is not just for large enterprises. Powerful workstations are democratizing CAE and simulation software. Software developers are incorporating analysis into their CAD software. Workstations can access some of the most powerful computers on earth. Microsoft has a vision of connecting all scientific and engineering databases with search capabilities. This access to engineering data will be as profound as the first time you searched the Web, and all it will take is your workstation.

Engineering jobs are not going to get any easier, but they sure are going to get more interesting. As Carl Bass, CEO of Autodesk once said in an interview: "There isn't a problem in the world that a great designer can't solve." Welcome to the design revolution. **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.*

DE'S NEW LOOK

You've likely noticed this issue of *Desktop Engineering* looks a bit different. We've updated the design to respect your busy schedule by providing quick hits of relevant information, but also respect your intelligence by providing in-depth articles that help you do your job. We think this new design strikes that balance through small things, such as adding page numbers to the articles on the cover so you can flip right to them, to larger things like adding new departments focusing on the areas you want to read about most. There are dozens of such changes, and we'll continue to refine the new design based on your feedback. Tell us what you think at de-editors@deskeng.com.

From napkin sketch to prototype in just weeks. Let's see what we can create in SolidWorks®.

You and your fellow designers are going to help Jeremy create some never-before-built products. You'll submit projects, share design ideas and vote on key design decisions. If you think you're ready, let's go design. **Watch. Share. Vote.** LetsGoDesign.tv

LET'S GO DESIGN
AN INTERACTIVE WEB SERIES
LetsGoDesign.tv

presented by SolidWorks

DE November 2010 VOLUME 16/ISSUE 3

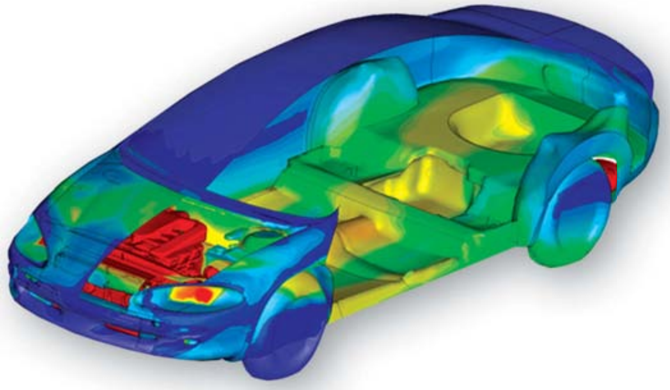

COVER STORY

ON THE COVER: Engineers break through bottlenecks as hardware and software combine to boost speed and efficiency.
Images courtesy of iStock.

HIGH-PERFORMANCE COMPUTING

SIMULATION & ANALYSIS

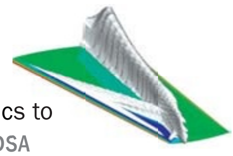

FOCUS ON RAPID TECH

November 2010 VOLUME 16/ISSUE 3

DEPARTMENTS

deskeng.com

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

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

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

***DE* ON TWITTER**
Follow us @ DE_editor@twitter.com

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

DE Desktop Engineering®

EDITORIAL
Steve Robbins | Executive Editor
Jamie J. Gooch | Managing Editor
Anthony J. Lockwood | Editor at Large
Heather Pittinger | Copy Editor

CONTRIBUTING EDITORS
Mark Clarkson, David S. Cohn, Al Dean, Mike Hudspeth, Susan Smith, Peter Varhol, Pamela J. Waterman, Kenneth Wong

PUBLISHER
Thomas Conlon

Advertising Sales
603-563-1631 • Fax 603-563-8192
Jeanne DuVal | Account Manager (x274)

ART & PRODUCTION
Darlene Sweeney | Director (x257)

A LEVEL 5 COMMUNICATIONS PUBLICATION
Steve Robbins | Chief Executive Officer
Thomas Conlon | President

ADVERTISING, BUSINESS, & EDITORIAL OFFICES
***Desktop Engineering*® magazine**
Level 5 Communications, Inc.
1283D Main St., PO Box 1039 • Dublin, NH 03444
603-563-1631 • Fax 603-563-8192
E-mail: DE-Editors@deskeng.com
www.deskeng.com

SUBSCRIBER CUSTOMER SERVICE
***Desktop Engineering*® magazine**
PO Box 677 • Northbrook, IL 60065-0677
847-559-7581 • Fax 847-564-9453
E-mail: den@omeda.com

***Desktop Engineering*®** (ISSN 1085-0422) is published monthly by Level 5 Communications, Inc., 1283D Main Street, P.O. Box 1039, Dublin, NH 03444, 603-563-1631. Periodicals postage paid at Dublin, NH, and at additional mailing offices. Desktop Engineering® is distributed free to qualified U.S. subscribers. Subscription rates for non-qualified: U.S. $108 one year; Canada and Mexico $126 one year; all other countries $195 one year.

LIST RENTALS: For information on list rentals, contact Statlistics, Danbury, CT: 203-778-8700.

POSTMASTER: Send all address changes to ***Desktop Engineering***, P.O. Box 677, Northbrook, IL 60065-0677.

Address all editorial correspondence to the Editor, Desktop Engineering. Opinions expressed by the authors are not necessarily those of *Desktop Engineering*. Unaccepted manuscripts will be returned if accompanied by a self-addressed envelope with sufficient first-class postage. Not responsible for lost manuscripts or photos.

BPA WORLDWIDE

by Kenneth Wong

AMD Launches Ultra-High End FirePro V9800

With 4GB frame buffer memory, ATI FirePro V9800 has twice the memory capacity of its predecessor, the V8800.

In September, just as Siggraph 2010 got under way, NVIDIA unleashed a series of Fermi-class Quadro cards. A month later, it was AMD's turn. Whereas its 2GB ATI FirePro V8800, released in April, was described as "high end," AMD characterizes its new 4GB V9800 as "ultra-high end."

Initially, AMD planned to sell the V9800 for $2,499, but shortly before its release, the company decided its listed price should be $3,499. That's significantly more than its predecessor, the V8800, which is listed for $1,499. However, in benchmark test results AMD provided, the V9800 appears to offer similar performance to the V8800's. In fact, in CINEBENCH 11.5 (based on usage of MAXON's Cinema 4D animation software), the V8800 seems to outperform the V9800 by a small margin.

Comparing Cards

At $3,499, AMD's V9800 costs more than its rival NVIDIA's Quadro 5000 ($2,249, 2.5GB memory), but less than Quadro 6000 ($4,999, 6GB memory). In single precision processing, the V9800 is said to reach 2.72 teraflops. Its predecessor reached 2.64 teraflops.

The V9800 supports up to six display ports through AMD's Eyefinity setup, allowing you to run programs on multiple monitors. It's powered by 1,600 stream processors, with support for DirectX 11, OpenGL 4.0 and Shader-Model 5.

Even though it has become common knowledge that AMD eventually plans to phase out its ATI brand (just the brand, not the product line), the V9800 still carries the familiar ATI logo.

In last month's *DE*, David Cohn reviewed AMD's ATI FirePro V8800, as well as the V3800, V5800 and V4800. For his take on the lineup, including benchmarks, visit deskeng.com/articles/aaaynr.htm. This month, he reviews NVIDIAs cards on page 28. **DE**

INFO ➜ American Micro Devices (AMD): amd.com

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

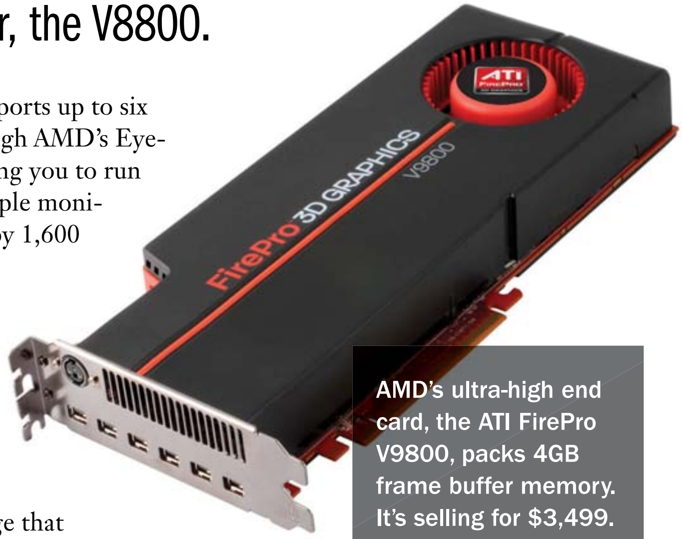

AMD's ultra-high end card, the ATI FirePro V9800, packs 4GB frame buffer memory. It's selling for $3,499.

V9800 Specifications

- GPU stream processors: 1600
- Memory interface: 256-bit
- Memory size/type: 4GB GDDR5
- Mini DisplayPort: 6
- Max Res.: 2560x1600 @ 60Hz
- Stereo (3-pin DIN): Yes
- DirectX: 11.0; OpenGL: 4.0; Shader Model: 5.0
- Microsoft Windows 7, XP, Vista, Linux 6 (32-bit or 64-bit)
- Max Power: <225 W
- Slots: 2
- Form Factor: Full Ht./Full Length

by Kenneth Wong

3DVIA Mobile for iPad

While Autodesk woos Mac fans with AutoCAD for Mac, Dassault Systèmes curries favor from the same crowd by releasing of 3DVIA Mobile HD.

Described as an "interactive 3D model viewer for the iPad," the new 3DVIA Mobile HD application lets you use an iPad to browse, view and play with 3D models hosted at 3DVIA.com, Dassault Systèmes' 3D content sharing platform.

Dassault's 3DVIA.com portal lets you upload models saved in 3DXML, DAE (COLLADA), 3DS, OBJ, IGES, STEP, and KMZ formats. You don't need an iPad to sign up for 3DVIA.com. However, if you do, the iPad-targeted 3DVIA Mobile HD application lets you zoom, pan and rotate your models using fingertips on the portable device. You may also export quick renderings of your models and publish them to Facebook from 3DVIA Mobile. iPhone users can use the application in a similar fashion, but the larger screen on iPad makes 3D presentations more effective.

Dassault Systèmes' 3DVIA Mobile HD for iPad might be a precursor to other lightweight viewing and markup applications for the Apple-dominated portable market.

Models hosted in 3DVIA.com are usually in lightweight, simplified geometry; therefore, they're better suited for online viewing then the more complex, history-based parametric CAD models used in manufacturing.

Dassault operates a site similar to 3DVIA.com called 3D Content Central. With a greater focus on mechanical engineering, 3D Content Central is more closely associated with its mid-range MCAD package SolidWorks. The company is also getting ready to release a version of its 2D drafting program, DraftSight, which will run on Mac OS as a native program. **DE**

INFO → Dassault Systems: 3dvia.com

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

by Kenneth Wong

Kepler and Maxwell on the Roadmap for NVIDIA

When they hit the market within the next three years, these two GPU architectures will take performance to new heights.

On Sept. 21, NVIDIA's cofounder, president and CEO, Jen-Hsun Huang, took the stage at the NVIDIA GPU Technology Conference (GTC 2010, San Jose Convention Center, CA), donning his trademark black T-shirt. The company's current GPU architecture, dubbed Fermi, represented a technological leap, he reminded us: "Fermi is six times faster than the previous generation."

Packing as many as 896 cores on a single board (Quadro Plex 7000), it could justifiably be considered a supercomputer on a chip. Under Fermi, NVIDIA GPU performance reached 1.5 double-precision gigaflops (DP GFLOPS) per watt, according to the company. Now, Huang revealed his ambitious roadmap for GPU, plotting the timeline between 2011 and 2013 with two new GPU architectures, dubbed Kepler and Maxwell. Kepler, scheduled to debut in 2011, is expected to push GPU performance to 4-6 DP GFLOPS per watt; Maxwell, targeted for 2013, will push it up to 14-16 DP GFLOPS per watt.

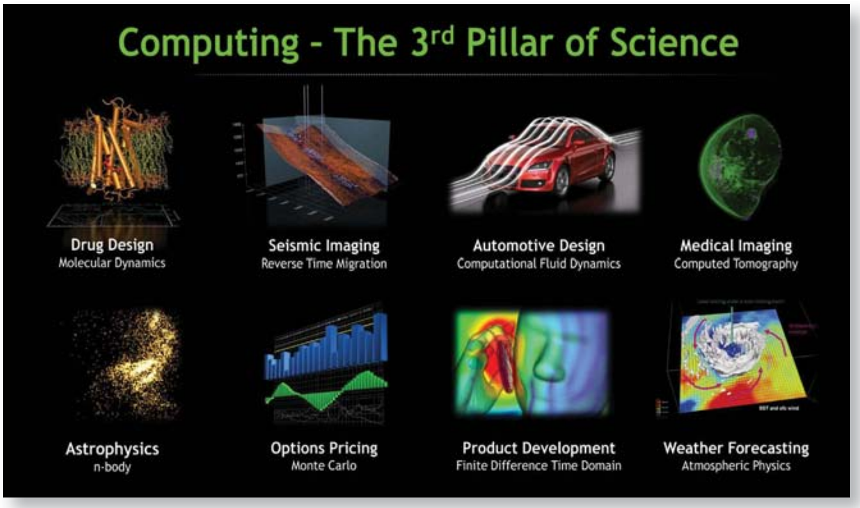

NVIDIA targets the HPC market with Tesla GPU clusters for science and research.

Focus on HPC

With high-performance computing (HPC) becoming prevalent in research and science ("the third pillar of science," as NVIDIA would like to call it), NVIDIA is also pushing GPU into the HPC market. The company targets consumers and gamers with its GeForce brand, mobile apps with Tegra and professionals with Quadro. For the HPC market, it develops and offers Tesla GPUs.

"Scientists, researchers, and engineers around the world want to have access to HPC for their research. With CUDA, we've now made it more affordable," said Huang.

NVIDIA has struck up a partnership with The Portland Group (PGI), a supplier of HPC compilers, to develop a CUDA compiler for industry-standard, general-purpose 64- and 32-bit x86 architectures. PGI plans to demonstrate it at the Supercomputing Conference 2010 this month.

CUDA, NVIDIA's programmable environment for GPUs, made it possible for many to develop applications that can take advantage of GPUs' parallel processing power. That means GPU is no longer confined to the graphics market, which had been dominated by game developers, animators and filmmakers. It can now reach into mechanical simulation, medical research, financial analysis and other fields. This puts GPU on a collision course with CPU—and by extension, its market leader, Intel, which is also targeting the HPC market with its multi-core CPUs currently available, and many-core CPUs that are still in development.

NVIDIA Tesla-enabled HPC systems are now available from IBM, T-Platforms and Cray, among others. Application developers supporting GPU-based computation also appear to be growing. According to NVIDIA, in 2009 the CUDA developer kit was downloaded approximately 293,000 times. This year, at press time, the number had jumped to 668,000. Recent additions to the list of GPU-enabled software brands include MatLab's Parallel Computing Toolbox, ANSYS Mechanical R13, and Autodesk 3ds Max. DE

INFO → NVIDIA: nvidia.com

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

by Kenneth Wong

3ds Max Adds iray as a Rendering Option

According to its manufacturer, mental images, iray is 'the world's first interactive and physically correct, photorealistic rendering solution.'

Beginning this fall, 3ds Max users on subscription will get an update that allows them to use mental images' GPU-accelerated iray rendering engine to visualize their scenes. (mental images is a wholly owned subsidiary of NVIDIA.)

In the announcement jointly issued by Autodesk and NVIDIA, the companies said, "While iray produces identical images on either CPUs or GPUs, 3ds Max users will enjoy up to 6X faster results over dual quad-core CPUs when using a GPU such as the new NVIDIA Quadro 5000 or Tesla C2050."

At the NVIDIA GPU Technology Conference, Ken Pimentel, director of visual communication solutions, media and entertainment, Autodesk, and Michael Kaplan, vice president of strategic development, mental images, joined NVIDIA CEO Jen-Hsun Huang to demonstrate the feature in 3ds Max.

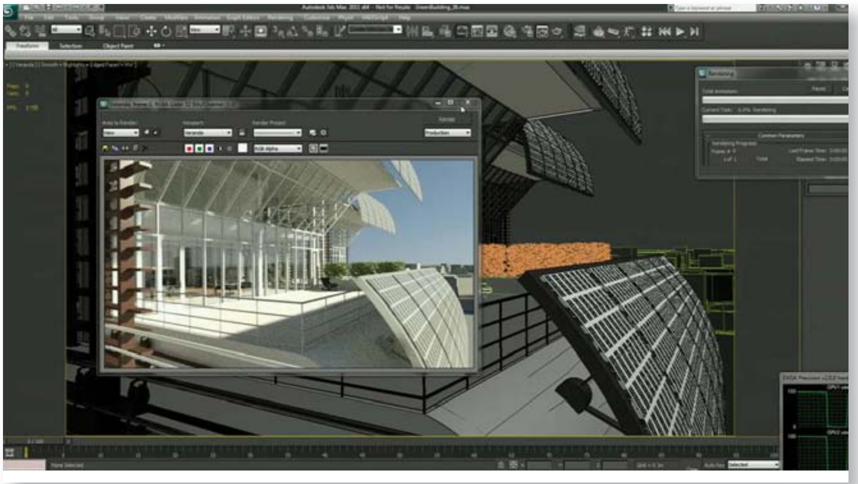

Autodesk 3ds Max users will soon have the option to use GPU-accelerated iray as their rendering engine.

Accelerated Computation

"Because rendering systems use interpolation to accelerate things, they have this command panel that looks like you're flying a 747. There are just so many dials and settings," said Pimentel. "What's great about iray is that it introduces another level of usability: It's point-and-shoot."

iray takes advantage of the GPU to accelerate computation—often producing the results faster than a CPU alone could. Because iray calculates light bounces according to physical laws, it requires no interpolation (or "fake reality," as some might call it). Thus, the input interface for rendering on iray is also comparatively simpler, according to Pimentel.

Browsing Renderings

Pimentel, Kaplan and Huang also demonstrated a browser-based solution, which lets presenters use a standard laptop to access a web-hosted 3ds Max scene (hosted by PEER1) from a web browser and render it with nearly real-time speed as they walked around the scene. The only data a user sent was its virtual position, as represented by the mouse pointer. All rendering took place in the cloud, on Fermi-class GPUs hosted remotely, explained Kaplan.

"(Cloud-hosted GPUs) are all running exactly the same iray software that comes with 3ds Max," he said. "We can guarantee that the image that you get from (cloud-hosted iray renderer) is exactly the same, pixel for pixel, as what you would get from 3ds Max."

Previously, rendering and animation was not intended to imitate reality. For the most part, they represented an artist's or an animator's representation of reality. Certain details in the result, be it a still image or an animation sequence, were often approximated or simplified because computing resources available weren't sufficient to reproduce every photon's path. The rise of multi-core CPUs brought some relief, but NVIDIA hopes more will turn to GPU cores to accomplish the same tasks at what it says is greater speeds. DE

INFO ➜ NVIDIA: nvidia.com

Autodesk: autodesk.com

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

Software Resale Issue Resurfaced

Courts are changing their stance on whether software is purchased or licensed.

A year ago, when U.S. District Judge Richard A. Jones rendered his judgment in Vernor vs. Autodesk, he sided with eBay merchant Tim Vernor, declaring "Mr. Vernor has prevailed in his claim that Autodesk's copyright does not prevent him from reselling his AutoCAD packages." (Source: "Software Resale on eBay: Court Sided with Reseller," Oct. 13, 2009).

Near the end of his 26-page statement, Jones also cautioned, "[Courts] across the nation have issued rulings that adopt and reject the equivalent of the parties' positions here. Vernor, 555 F. Supp. 2d at 1174 [citing rulings from courts in California, New York and Utah.] Since that opinion, at least two more courts have joined the cacophony. The MDY court (which decided the MDY Industries vs. Blizzard Entertainment case) favored the copyright holder, whereas the court in UMG Recordings (UMG Recordings, Inc. vs. Augusto) favored the consumer."

Jones' observations proved prophetic. The 9th Circuit U.S. Court of Appeals recently reversed his ruling. "Autodesk retained title to the software and imposed significant transfer restrictions." In the summary for the unanimous three-judge panel, 9th Circuit Judge Consuelo Callahan wrote, "It stated that the license is nontransferable, the software could not be transferred or leased without Autodesk's written consent …" (Source: "Autodesk Wins Copyright Appeal," Reuters, Sept. 10, 2010).

Most software merchants like Autodesk contend that their software is "licensed," not "sold;" therefore, they argue, software buyers (or licensees) cannot resell, trade or transfer their software—like they could with, say, used books or furniture. DE

INFO ➜ Autodesk: autodesk.com

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

ZWSoft Lays Groundwork for ZW3D

Product will be ZWSoft's all-in-one offering, signaling a consolidated focus.

In September, after two years of confidential negotiations, Chinese software maker ZWSoft bought VX Corp. With the completion of the transaction, VX Corp. becomes a wholly owned subsidiary of ZWSoft.

"The VX CAD/CAM brand will go away. It'll be transitioned into ZW3D," clarified Bob Fischer, vice president of sales and marketing for VX Corp.

In the announcement, ZWSoft states, "This new all-in-one CAD/CAM solution will extend enterprise-level mechanical design and engineering capabilities to the desktop at a fraction of the cost of other comparable systems."

When ZWSoft releases ZW3D, VX customers will be able to purchase it from their resellers, Fischer explained.

ZWSoft attracted attention with its AutoCAD clone ZWCAD, primarily a 2D drafting package. As a less-expensive alternative to AutoCAD and AutoCAD LT, ZWCAD begins to test the market long dominated by Autodesk. With the pending release of ZW3D, ZWSoft is about to butt heads with other 3D CAD products, including Autodesk Inventor.

VX CAD/CAM is known for sophisticated surfacing functions. The repackaged product, ZW3D, will continue to offer the same functions, confirmed Fischer.

For VX, the Chinese CAD market is uncharted waters, difficult to navigate without the help of a local force like ZWSoft. For ZWSoft, repackaging an established, robust 3D CAD product is a better alternative to spending years of R&D developing its own.

Most CAD products are based on a 3D modeling kernel, licensed from either ACIS or Parasolid. VX CAD/CAM is an exception to the rule; it's based on its own (and thus exclusive) VX kernel.

Just as this issue goes to press, ZWSoft struck an agreement with Siemens PLM Software to license the latter's D-Cubed software, to be incorporated into ZWCAD Professional. D-Cubed offers design and drafting capabilities, DWG file import/export and application programming interfaces. DE

INFO ➜ ZWSoft: zwcad.org

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

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 KennethongSF, or send e-mail to DE-Editors@deskeng.com.*

Industrial Sector Surges Ahead

1 **ThomasNet's** newest Industry Market Barometer (IMB) shows North America's industrial/manufacturing sector surging forward, accelerating the momentum of its recovery with the promise of further expansion to come. Several indicators demonstrate a sector that's "caught on fire" with more company growth, fewer declines, a resumption of hiring, a winding down of lay-offs, and new investments.

"Looking at three Industry Market Barometers over the last 18 months, we've seen a winning formula unfold. Those companies that reported deploying new strategies, such as new product development and online marketing, have propelled to growth mode. Buoyed by the results of their initiatives, these successful companies are moving to the next level, fueling industry's recovery," says Eileen Markowitz, president of ThomasNet.

According to findings reported in the ninth annual Industrial Indicator Survey conducted by **GlobalSpec**, the number of organizations reporting expectations of increased revenue in the current year doubled from last year (from 24% to 51%), reflecting optimism about the recovering economy within the U.S. industrial sector. In addition, 66% stated company revenue is on target or above for 2010.

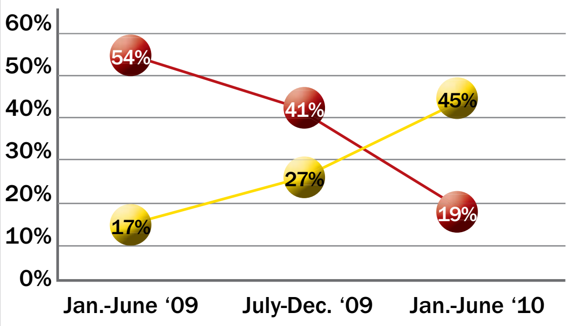

The industrial sector is showing more growth (yellow) than declines (red), according to research from ThomasNet.

Think3 Acquired by Versata Enterprises

2 **Think3, Inc.**, a provider of design and modeling software, will join the family of Versata affiliated software companies. Think3 provides technology that links three separate design areas: the concept, its development, and the finished product.

Think3 will continue to operate as a stand-alone corporation within the Versata family of software businesses.

Scott Brighton has been named CEO of Think3. Brighton most recently served as president of Trilogy Enterprises.

"Our goal is to make our entire customer base, which includes over 31,000 customers worldwide, completely successful with Think3," says Brighton.

Chris Smith is the newly appointed COO of Think3.

Bits from Bytes Acquired by 3D Systems

3 **Bits from Bytes,** which manufactures the BfB 3000 and RapMan 3D printers, has been bought by 3D Systems. The acquisition will see the entire Bits from Bytes management team remain while operations will continue from its Clevedon, UK, headquarters.

Manufacturing affordable kits and printers ranging between $1,300 to $4,000, Bits from Bytes gained attention for opening 3D printing to a range of sectors where previously the price was too prohibitive.

Employing 13 people, the company was set up by Ian Adkins and Iain Major in December 2008 and accounted for 17% of all 3D printer sales worldwide, ranking it second in total shipments in 2009, according to *Wohlers Report*.

AutoCAD WS Extends Design into the Cloud

4 **Autodesk, Inc.** has announced the availability of AutoCAD WS, a new free web application that uses cloud computing technology to enable AutoCAD software users to view, edit and share their AutoCAD designs and DWG files through web browsers and mobile devices. The AutoCAD WS mobile application for iOS is also now available as a free app in the App Store for iPad, iPhone and iPod touch.

Users can upload and manage designs in their online workspace from their AutoCAD desktop software through a new free plug-in, or as part of the Subscription Advantage Pack for AutoCAD 2011. It will also be integrated into the AutoCAD for Mac software. DE


100% modo image from Bruce Long, Giannini Creative, Chicago
Form follows fun
Luxology

The HPC Revolution

Top tech companies push for HPC-driven engineering.

BY KENNETH WONG

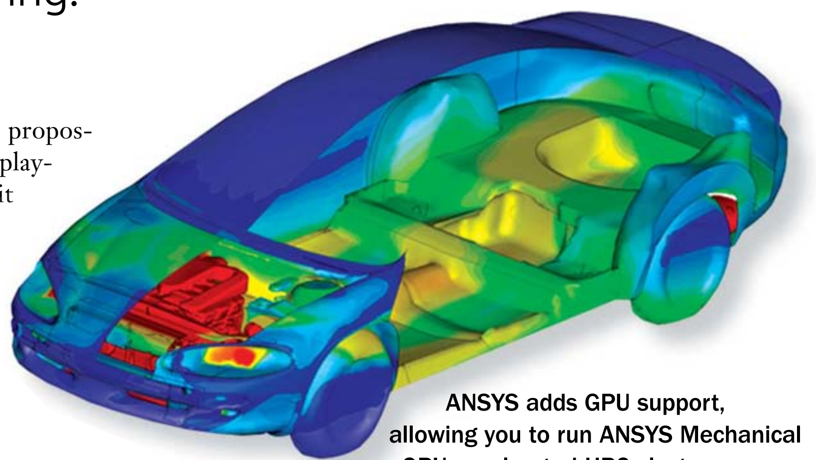

ANSYS adds GPU support, allowing you to run ANSYS Mechanical on GPU-accelerated HPC clusters.

Some might consider what Microsoft is proposing at its technical computing portal as playing God in pixels. Others might view it as technology imitating life. "Modeling the World" (modelingtheworld.com) showcases an initiative to enable scientists, engineers, and researchers to digitally reproduce complex bio-eco-mechanical systems that affect our lives—weather patterns, traffic flows, disease spread, to name but a few—so they can be used for predictions.

In his personal note posted to the site, Bob Muglia, president of Microsoft Server and Tools Business, writes, "Our goal is to unleash the power of pervasive, accurate, real-time modeling to help people and organizations achieve their objectives and realize their potential … One day soon, complicated tasks like building a sophisticated computer model that would typically take a team of advanced software programmers months to build and days to run, will be accomplished in a single afternoon by a scientist, engineer or analyst working at the PC on their desktop."

It's not a coincidence that the site went live in May, a few months before the unveiling of Microsoft HPC Server 2008 R2, the company's software for supercomputing. Modeling the world, the company believes, will require:

- Pushing technical computing to the cloud.
- Simplifying parallel development.
- New technical computing tools and applications.

HPC in the Cloud

For most designers and engineers, modeling the world may be confined to rendering, simulation, and analysis. The latest generation of workstations equipped with multicore CPUs and GPUs can easily churn out photorealistic automotive models and architectural scenes produced in a high-end 3D modeling software within a few seconds or minutes, but in the near future, you may be able to take advantage of remote HPC systems through a browser to speed up the process close to real-time.

At the GPU Technology Conference 2010, NVIDIA cofounder, president, and CEO Jen-Hsun Huang was joined by Ken Pimentel, director of visual communication solutions, media and entertainment, Autodesk; and Michael Kaplan, VP of strategic development, mental images.

Pimental and Kaplan were there to break the news about an online rendering option, slated to become available to 3ds Max users in the near future (no definitive date yet). Powered by iray rendering engine from mental images (a subsidiary of the GPU leader NVIDIA), the new feature lets you create photorealistic renderings from a browser window, tapping into the computing horsepower of a GPU cluster hosted offsite.

For those who routinely perform computational fluid dynamics (CFD), finite element analysis, and other engineering simulations, a similar web-hosted option is available from Dezineforce, a UK-headquartered company.

Dezineforce gives you access to analysis software from well-known vendors, including ANSYS, MSC Software, Oasys, and Dassault Systèmes. You can buy monthly subscription packages, essentially paying for the number of CPU core hours used on Dezineforce's HPC platform. Prices start from £5/$7.50 per core hour, excluding the analysis software license cost.

Modeled on SaaS or software as a service, the new HaaS or hardware as a service model is catching on among other engineering software makers. Altair, for instance, complements its PBS Works software suite with on-demand HPC services. If your in-house computing resources are inadequate for your needs, you tap into Altair's HPC resources, then pay for your usage per hour, per node.

According to Microsoft, the company will soon release an update to Windows HPC Server that allows customers to provision and manage HPC nodes in its Windows Azure cloud platform from within on-premises server clusters, in order to utilize computing power on-demand.

HPC in a Box

On-demand HPC, or cloud-hosted HPC, is a viable option for those who need it periodically. For those who must perform intense computation tasks regularly, HPC in a box may be a better option. Dezineforce also sells HPC in a box.

Measuring roughly 28x30x47 in., Dezineforce's HPC Simulation Appliance is powered by Intel Xeon X5670 processors, configurable in 48 to 384GB of memory; 0.5 to 10TB of storage; and 24, 48, 96, or 192 cores.

For years, GPU has been relegated to the realm of digital content creation, game development, and filmmaking. But in the last few years, NVIDIA began moving into biomedical imaging, financial analysis, and engineering simulation. The releases of several new GPU-based HPC systems (IBM BladeCenter, T-Platforms TB2, and Cray XE6), all timed to coincide with the GPU Technology Conference, showed NVIDIA's Tesla cards are making inroads in HPC.

NVIDIA's foray into HPC is partly made possible by the launch of CUDA, the company's programmable environment for parallel processing on GPUs. Much in the same way you can take advantage of multicore CPUs for parallel processing, you can now break down a large computational operation into smaller chunks, executable on many GPU cores.

A Hybrid Processor

As the competition between CPU and GPU heats up, AMD, which provides both types of processors, is planning on a hybrid model, an APU (accelerated processor unit) that combines both a CPU and a GPU on a single die.

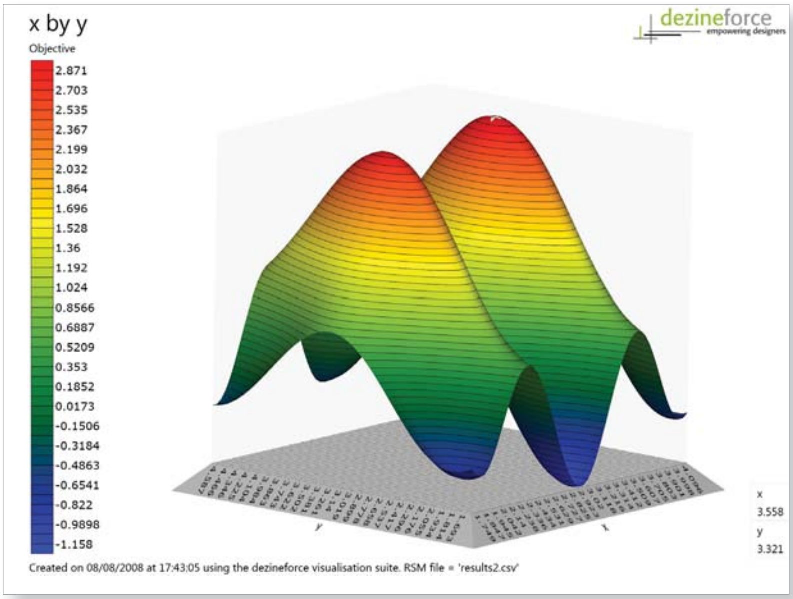

Dezineforce delivers computing resources on demand, giving subscribers remote access to HPC systems to run analysis and simulation programs.

"APUs combine high-performance serial and parallel processing cores with other special-purpose hardware accelerators, enabling breakthroughs in visual computing, security, performance-per-watt and device form factor," the company declared.

"What we don't see so much is an entire application running on one architecture or the other [on CPU or GPU]," explains Patricia Harrell, AMD's director of HPC. "An HPC application most likely will make use of both multi-core CPUs and GPUs."

AMD plans to release Fusion brand APUs initially into the desktop and laptop market before migrating them into the server market. At the present, the APU is still in development. (The product has a dedicated home page at AMD site, but it greets you with, "Coming soon!") The company currently offers its Opteron CPUs and FirePro GPUs for original equipment makers.

Core-Based Licensing Falls Apart with GPU

The emergence of GPU-driven HPC systems may also persuade simulation and analysis software developers to re-examine their licensing model, traditionally tied to the number of cores.

In late 2010, ANSYS plans to add GPU-powered acceleration to its product line, beginning with ANSYS Mechanical R13. The hundreds of computing cores embedded in a single GPU (as opposed to CPU cores, which currently come in 2 to 32 cores in a single machine) make it impractical to license GPU-acceleration on a per-core basis.

"The number of cores is so much higher on a GPU than on a traditional processor—and the performance per core is so different from traditional processors—that it calls for a different

TODAY'S NETBOOK, YESTERDAY'S HPC

By 1970s' standards, an average desktop or notebook computer available from Best Buy or Fry's today would be a supercomputer. The computing titan of the time, Cray 1, was installed at Los Alamos National Laboratory in 1976 for $8.8 million. It boasted a speed of 80MHz, executing 160 million floating-point operations per second (160 megaflops) with 8MB of memory. By contrast, today's HP Mini 110 (a netbook) in the most basic configuration is powered by a 1.66GHz Intel Atom processor, with 1GB memory. The price is roughly $275.

approach on the GPU (e.g., the core-based licensing makes no sense on a GPU)," says Barbara Hutchings, ANSYS' director of strategic partnership. "To support our new capability, we will be making the GPU available if you have purchased our HPC Pack license. That is, the customer can buy a single license that enables both traditional processor cores and the GPU. This allows customers to choose the hardware configuration that makes sense for their application, without licensing being a barrier to the use of GPU acceleration."

Similarly, Altair also offers its PBS Works suite for both CPU and GPU. Robert Walsh, Altair's director of business development for PBS GridWorks, says, "Today we don't even charge you for GPU management. So, for example, you have a machine with a quad-core CPU and a 128-core GPU, you're only paying for using the licenses of PBS Works you run on the CPU."

Some software users feel handicapped by the traditional per-core licensing model, which limits them to run the analysis on a finite number of cores even though they may have invested in hardware with a far greater number of cores. "Many simulation vendors are now offering highly scalable licensing," notes ANSYS' Hutchings, pointing out that software makers are adapting.

Data Deluge

With so much content being created and shared through public forums and social media, the democratization of technology seems to have produced an unforeseen side effect: what *The Economist* calls "The Data Deluge" (Feb. 27-March 5, 2010). Citing the article during the "Pushing through the Inflection Point with Technical Computing," keynote address at the High Performance Computing Conference in September, Bill Hilf, Microsoft's general manager of technical computing, said: "We're outpacing our ability to store data. We're actually creating more bytes than we have the capacity to store. The challenge really is, what do you keep, what do you throw away? Many analysts predict that over the next five years, we'll produce more data than we have ever produced in the history of human kind."

To analyze, understand, and extract precious nuggets of wisdom from the sea of data, Microsoft expects many will turn to parallel processing and HPC, and specifically to Microsoft Visual Studio, an integrated software development platform, to identify and take advantage of parallelization opportunities.

HPC is beyond Microsoft's traditional territory, operating systems. To navigate these uncharted waters, the company needs to bring hardware developers onboard. A Microsoft press statement explained, "Both NVIDIA Parallel Nsight 1.5 and Intel Parallel Studio 2011 make Visual Studio an integrated development environment for parallelism. NVIDIA provides source code for debugging GPU applications along with tracing and profiling for both CPU and GPU, all on a single correlated timeline. Intel Parallel Studio adds several tools to improve parallel correctness checking, race condition detection, and parallel code analysis and tuning." DE

Kenneth Wong *has been a regular contributor to the CAD industry press since 2000, first an editor, later as a columnist and freelance writer. He has closely followed the migration from 2D to 3D, the growth of PLM, and the impact of globalization on manufacturing. He is* Desktop Engineering's *resident blogger. Email him at Kennethwongsf@earthlink.net or follow him on Twitter@KennethwongSF.*

WHAT WILL YOU DO WITH HPC?

When computing resources are limited (say, to a single workstation), you may be forced to simplify a problem to reduce the workload. In conducting analysis, you may be forced to decrease the mesh resolution, thereby compromising the accuracy of your results; in rendering a scene, you may be required to reduce surface reflectivity. Therefore, greater access to HPC could mean fewer compromises.

Whereas single-core computing usually confines you to simulating one design scenario at a time, you may be able to run several variations of the same scenario simultaneously using the multi-core/many core computing paradigm. This lets you simulate, for example, airflow patterns around a racecar's hood, designed with several different curvatures, all at the same time.

For small businesses that may not be ready to invest in HPC hardware, cloud-hosted HPC offers a viable alternative, but enabling the function is easier if incorporated into the design, rendering, or analysis software itself. In computer-aided machining (CAM) and tooling, software developers like CoreTech (makers of Moldex3d) and HSMWorks (makers of HSMWorks, a SolidWorks CAM plug-in) have built HPC-support into the software interface itself.

Autodesk has a number of technology previews hosted at Autodesk Labs. Project Centaur, a web-hosted design optimization technology, lets you specify materials, constraints, load conditions, and variable parameters, then find the best design option through computation. Similarly, Project Neon let you upload an AutoCAD file with camera angles and saved views, then use remote computing resources to render the scene.

Because technologies in Autodesk Labs are not yet part of commercial products, it's unclear how the company plans to license these features or charge for accessing them. But they serve as working prototypes of how you will be able to take advantage of remote HPC in the future.

Get Your Head in the Cloud

It may not be ready for prime time, but get ready to kick the tires with your engineering software.

BY PETER VARHOL

Many who have heard anything about cloud computing believe that it has no impact on engineers. In reality, we use cloud computing almost every day, for tasks mundane and highly technical. And in many cases, it can be an appropriate solution to engineering computing needs.

Maybe you think you aren't interested in cloud computing, but chances are you use cloud applications already. If you use Gmail, Hotmail, Yahoo Mail or any of the Google applications, you are already well into the cloud. Store your photos on Picasa, Snapfish, Flickr or Shutterfly? You're also using the cloud, albeit probably for personal rather than professional purposes.

The question is whether the cloud can meet the needs of engineers in design and analysis work. Can cloud computing help you with your job? You may already be using Gmail and Google Apps at work, for example, but for supporting roles. Can the same approach you use with other software work on the most important software you use?

Cloud Categories

There are a number of different technologies that fall under the umbrella of cloud computing. At the most basic level, cloud computing involves using applications that run across the Web—in "the cloud," so to speak. Web email services, Salesforce.com, Google Apps and a host of other applications fall into this category.

But a second category that has emerged over the last couple of years is capacity on demand, best characterized by Amazon's EC2 cloud. These clouds are simply empty servers that users can "rent" for a set period of time, for a price that varies depending on the type of server and quality of service. Typically, it costs just a few cents an hour. Users create a virtual machine of their application and upload it to the server, where it executes.

A third category is also capacity on demand, but with significant support software, including operating systems, application frameworks and databases. Users run their applications without having to include all of the support software and services into the virtual machine. Microsoft's Azure and Google's App Engine fall into this category.

A variation of this category is for a commercial software vendor to operate its own cloud with its own application, inviting engineers to rent time on their servers if they can't afford to purchase a license outright, or if their needs are only occasional.

Here's where it gets interesting for design engineers. As more vendors look at how they might rent software to customers with only occasional needs or with limited budgets, it may make sense for these vendors to operate their own data centers running their software. While this sort of software rental is in its infancy, it has significant potential for growth, if in fact they can successfully deliver design features and performance over the Web.

Large and geographically distributed design shops might also set up their own private cloud for internal use. These shops have multiple groups with similar needs, and it may make sense to consolidate engineering software under a single umbrella, and make it available to users over the internal network.

There's some attraction to using design and analysis software in the cloud. One significant advantage is versioning: Engineering customers are always on the latest version of the software, and vendors manage upgrades on their own servers. Another is the pay-as-you-go model, which may appeal to the corporate financial types who prefer writing a series of small checks rather than one large one. Lastly, for distributed teams, it provides an easy way to share files and combine design efforts.

Running in the Clouds

Even if your favorite design and analysis software were installed and available on cloud servers, you probably won't just drop your existing licenses and sign up for the same software delivered on the cloud. There are some downsides, and they are significant enough to bring pause to anyone eying the cloud as a solution to existing problems:

The first is performance. You would think that servers on the cloud might be more powerful than those systems available at your desk. That may be true, but remember that multiple users are sharing those servers. Plus, there is the inevitable lag caused by the user interface. Engineering software is often visually detailed, and that means that a lot of bits have to go over the Web to reach the engineer's desktop. Some vendors are responding by "re-architecting" their software in a client/server configuration to run computations in the cloud and display results directly on the engineering workstation. Others are providing the ability to do some types of design through a browser.

Then there's a matter of security. You and your organization probably want to protect your files and data from unauthorized access. With that data normally residing outside the organization, and stored on systems that are administered by another company, there may be legitimate worries of having designs lost or stolen.

There may also be cultural issues involved. Most organizations want a certain amount of control over their intellectual property, and that typically means that designs don't go outside the company. There may be a greater sense of control and protection by having design data on the box in the next room rather than in an indeterminate location some distance away.

Betting on Improvements

There are vendors who are working to address these limitations and make the cloud a better place for engineering applications. Autodesk has been developing browser-based applications that are currently in development, including Autodesk Project Twitch, Autodesk Project Butterfly, Autodesk Project Dragonfly and Autodesk Project Showroom. For example, Project Butterfly is an Autodesk labs project that allows users to edit and collaborate on AutoCAD drawings using a web browser, while Project Dragonfly does the same thing with

Standard On-Demand Instances	Linux/UNIX Usage	Windows Usage
Small (Default)	$0.085 per hour	$0.12 per hour
Large	$0.34 per hour	$0.48 per hour
Extra Large	$0.68 per hour	$0.96 per hour
Micro On-Demand Instances		
Micro	$0.02 per hour	$0.03 per hour
High-Memory On-Demand Instances		
Extra Large	$0.50 per hour	$0.62 per hour
Double Extra Large	$1.00 per hour	$1.24 per hour
Quadruple Extra Large	$2.00 per hour	$2.48 per hour
High-CPU On-Demand Instances		
Medium	$0.17 per hour	$0.29 per hour
Extra Large	$0.68 per hour	$1.16 per hour
Cluster Compute Instances		
Quadruple Extra Large	$1.60 per hour	N/A*

** Windows is not currently available for Cluster Compute Instances.*

Current Amazon EC2 cloud pricing enables users to rent even server clusters at far less than purchasing comparable systems.

floor planning software. A shipping product, AutoCAD WS, is a mobile version of AutoCAD for iPad and iPhone.

SolidWorks is also dipping its toe into the cloud. At the company's conference last February, it previewed a virtualized version of the new SolidWorks release running on a netbook computer, and made statements of its commitment to the cloud.

Special-purpose middleware is also making it possible to run applications that have been tuned to work with that software. Middleware from companies such as Acceleware and ScaleMP sit between the operating system and the application, and use both general-purpose and industry-specific algorithms to take parts of the application execution and break them into parallel components. Those components can then execute on the underlying Windows operating system in threads that run in parallel on separate cores. Because you can rent as many processors and cores as you like on a cloud, this makes it possible to take advantage of far more cores than may be available in-house.

On the software side, MathWorks MATLAB language provides a keyword that enables parallel execution of defined parts of the code. MATLAB also enables the programmer to define the execution environment, even defining a specific cloud such as the Amazon EC2. By configuring that location for execution, the programmer can ensure that the application is optimized for the cloud.

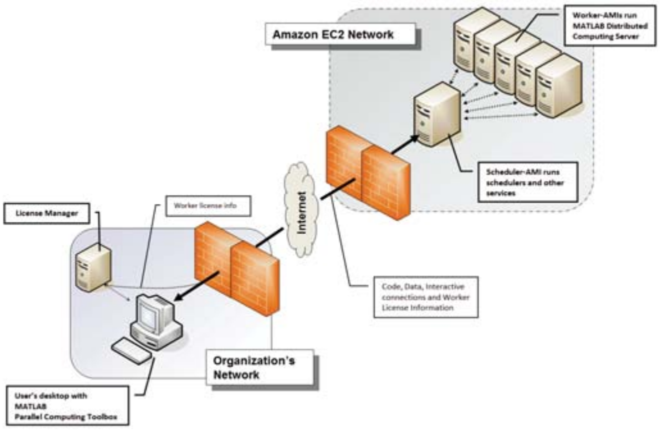

MathWorks' MATLAB language lets engineers dynamically deploy a computation to the cloud for execution.

Managing engineering workloads in the cloud is also a significant challenge. Adaptive Computing is attempting to make the cloud more amenable to engineers requiring HPC services. Its Moab Cluster Suite provides an engine working in conjunction with cloud infrastructure that integrates scheduling, managing, monitoring and reporting of cluster workloads. Its companion product, HPC Suite, changes a node's operating system on the fly in response to workload needs.

While you may not be ready for the cloud today, it is likely to save you money and offer better convenience and collaboration features in the future. The time when cloud computing serves all of our engineering design needs is a long way off, but there can be real benefits to getting started now. **DE**

GETTING STARTED

Running engineering software in the cloud may be the future, but the future is not yet here. There are too many open questions to fully move to a cloud model for most engineering groups.

Still, there are ways you can start using the cloud right now, and good reasons for doing so. First, you can usually get more computational power out of renting servers in the cloud than you can from your desktop, and the cost of doing so is low. Second, renting software is ultimately what vendors want, so that they can even out their revenues—and they will make it compelling to do so. But you need to prepare and build experience in cloud computing. Here's how:

1. Find out what cloud offering your software vendors have or are developing. Negotiate a test drive so that you have a good feel for their capabilities and plans.

2. Determine what you use in the way of computational power for your engineering software, and price out equivalent capacity on Amazon EC2 or an equivalent cloud. If you run your own computation jobs in-house, start packaging up some of them to run on EC2.

3. Look at the capabilities currently available in the cloud, and how those capabilities can supplement the work that you're doing today—either by offloading some existing computing, or doing supplementary design or analysis work.

4. Make plans to move some work to the cloud today. The experience will prepare your team and your IT support to use the cloud for large projects in the future.

Contributing Editor **Peter Varhol** *covers the HPC and IT beat for DE. His expertise is software development, math systems, and systems management. You can reach him at DE-Editors@deskeng.com.*

INFO ➜ **Acceleware:** acceleware.com

➜ **Adaptive Computing:** adaptivecomputing.com

➜ **Amazon EC2:** aws.amazon.com/ec2

➜ **Autodesk:** autodesk.com

➜ **MathWorks:** mathworks.com

➜ **ScaleMP:** scalemp.com

➜ **SolidWorks:** solidworks.com

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

The Siemens answer:
PLM Software to build the right product — and build the product right.

Escalating costs. Tighter schedules. Increased complexity. The challenge of turning more ideas into successful products has never been greater. You need a product lifecycle management solution that improves everything from design to manufacturing – and beyond. Find out how Siemens PLM Software can help your company stand out, no matter how tough the challenge. **www.siemens.com/plm.**

Answers for industry.

SIEMENS

HP Z200 Brings Power *to the* People

HP's entry-level CAD workstation is big on performance and easy on your wallet.

BY DAVID COHN

Earlier this year, HP introduced the Z200, the newest member of its Z-series workstations. Targeted at entry-level CAD users, prices start at $769, making the Z200 one of the most affordable systems out there certified for use with CAD applications.

The HP Z200 we received came housed in a newly designed case that bears considerable resemblance to the BMW Designworks styling we first saw in the flagship HP Z800 we reviewed earlier this year (*DE* January 2010). While its lineage is apparent, the Z200's configuration is more conventional, with three 5.25-in. drive bays with exposed front panel access. The top-most bay contains a tray-loading 16X SATA dual-layer SuperMulti LightScribe DVD+/-RW optical drive, while the lower bay houses a 22-in-1 media card reader. These bays can be rotated 90° to convert from a minitower to a desktop configuration and a Blu-ray writer is available as an option.

Below these bays, vertical fins conceal the air intake and internal speaker while a power button, two USB ports, headphone and microphone jacks, and an IEEE 1394a FireWire connector occupy a narrow vertical panel down the right side. The rear panel adds six more USB 2.0 slots, PS/2 keyboard and mouse connectors, an RJ-45 jack for the integrated Intel 82578 Gigabit LAN, and audio-in, audio-out, and microphone jacks. There are also DVI-I and DisplayPort connectors for the integrated Intel graphics.

Plenty of Expansion Options

Removing the side panel on the tool-less chassis reveals the clean interior we've come to expect from HP, although a far cry from the modular organization of the flagship Z800. Below the external drive bays are three 3.5-in. bays, one of which contained a Western Digital 500GB 7,200 rpm SATA drive. Drives of up to 1.5TB are available, and HP also offers solid-state drives.

CPU options for the Z200 range from a 2.56GHz Celeron up to a 3.6GHz dual-core i5 or a 3.06GHz Xeon quad-core processor. Our evaluation unit came with a 3.47GHz Intel i5-670 dual-core CPU with 4MB of L2 cache mounted in the single CPU socket. One of Intel's newest 32 nanometer CPUs,

CAD Workstations Compared

		HP Z200 Workstation (one 3.47GHz Intel i5-670 dual core CPUs, NVIDIA Quadro FX 1800, 4GB RAM)		**Lenovo D20** workstation (two 2.66GHz Intel Xeon X5550 quad core CPUs, NVIDIA Quadro FX 4800, 8GB RAM)		**Dell Precision T3500** workstation (one 2.27GHz Intel Xeon E5520 quad core CPU, NVIDIA Quadro FX 3800, 4GB RAM)		**Lenovo S20** workstation (one 2.27GHz Intel Xeon E5520 quad core CPU, NVIDIA Quadro FX 3800, 4GB RAM)		**HP Z800** workstation (two 3.2GHz Intel Xeon X5580 quad core CPUs, NVIDIA Quadro FX 4800, 12GB RAM)		**HP xw8600** workstation (two 3.4GHz Intel Xeon X5492 quad core CPUs, NVIDIA Quadro FX 4800, 4GB RAM)		**Levono ThinkStation S10** workstation (2.66GHz Intel Core 2 Q6700 quad core CPU, NVIDIA Quadro FX 4600, 2GB RAM)
Price as tested		$2,089		$5,943		$2,544		$3,885		$10,604		$9,307		$2,589
Date tested		8/7/10		1/11/10		7/30/09		7/29/09		4/24/09		12/22/08		6/30/08
Operating System		Windows XP	Windows Vista	Windows XP	Windows Vista	Windows XP	Windows Vista	Windows XP	Windows Vista	Windows XP 64	Windows Vista 64	Windows XP	Windows Vista	Windows XP
SPECviewperf	higher													
3dsmax-04		**60.87**	60.65	50.38	51.21	39.91	42.75	48.43	52.59	50.55	51.51	52.24	54.61	**37.88**
catia-02		**68.13**	66.87	61.79	62.01	51.85	53.33	60.40	60.61	62.10	61.66	63.17	62.48	**48.25**
ensight-03		53.85	53.06	55.26	53.51	47.26	47.84	51.74	**55.33**	53.99	53.62	54.44	50.82	**43.33**
maya-02		238.59	208.40	**250.41**	223.73	220.79	199.04	232.92	207.87	231.80	209.74	234.50	193.15	**191.10**
proe-04		**68.03**	65.74	64.83	63.66	55.67	55.54	61.56	64.49	63.59	61.48	52.73	57.15	**48.86**
SW-01		138.22	137.48	144.17	**145.19**	123.28	120.57	136.81	139.54	135.24	128.08	109.91	119.29	**90.90**
tcvis-01		35.60	34.81	**40.55**	39.51	28.71	28.07	29.17	38.76	28.93	28.29	29.84	27.58	**24.46**
ugnx-01		30.91	31.23	**34.93**	34.52	33.40	32.27	33.41	33.19	33.34	32.38	34.17	31.14	**27.04**
SPECapc SolidWorks	lower													
Score	seconds	148.72	n/a	141.59	n/a	178.39	n/a	**140.42**[1]	n/a	145.17[1]	n/a	164.71	n/a	**188.01**
Graphics	seconds	56.83	n/a	41.48	n/a	62.99	n/a	47.33[1]	n/a	**41.31**[1]	n/a	54.18	n/a	**60.13**
CPU	seconds	32.81	n/a	33.00	n/a	36.68	n/a	**31.01**[1]	n/a	32.68[1]	n/a	**44.36**	n/a	41.48
I/O	seconds	**63.10**	n/a	67.73	n/a	83.35	n/a	65.86[1]	n/a	71.94[1]	n/a	69.96	n/a	**90.19**
SPECapc SolidWorks	higher													
Score	ratio	5.27	n/a	6.28	n/a	4.66	n/a	5.91[1]	n/a	**6.38**[1]	n/a	4.84	n/a	**4.56**
Graphics	ratio	3.23	n/a	4.68	n/a	2.92	n/a	3.92[1]	n/a	**4.85**[1]	n/a	3.55	n/a	**3.15**
CPU	ratio	9.83	n/a	9.78	n/a	8.80	n/a	**10.41**[1]	n/a	9.87[1]	n/a	**7.27**	n/a	7.72
I/O	ratio	**5.02**	n/a	4.67	n/a	3.80	n/a	4.81[1]	n/a	4.40[1]	n/a	4.52	n/a	**3.51**
Autodesk Render Test	lower													
Time		137.4[1]	135.2[1]	64.00[1]	63.60[1]	118.2[1]	125.0[1]	99.0[1]	117.6[1]	59.00[1]	**52.00**[1]	64.40	67.60	**153.20**

Numbers in **blue** indicate best recorded results. Numbers in **red** indicate worst recorded results. Results are shown separately for portable and desktop workstations.

Our HP Z200 evaluation unit turned in some of the best test results ever, making this system a performance leader.

the 670 has a maximum Thermal Design Power of 73 watts, a maximum turbo-boost frequency of 3.73GHz, and it supports HyperThreading. Our system also included 4GB of 1,333MHz DDR-3 memory, installed as a pair of 2GB DIMMs, leaving two additional memory sockets available for future expansion. The Z200 supports up to 16GB of RAM. Power was provided via a 320 watt 89% efficient power supply.

The motherboard provides seven slots: one PCIe Gen2 x16 graphics slot, one PCIe Gen1 x1 slot, one PCIe Gen1 x4 slot (x1 electrically), one PCIe Gen1 x16 slot (x4 electrically), and three full-length PCI slots. The graphics slot was filled with an NVIDIA Quadro FX 1800 graphics accelerator with 768MB of memory. There are also up to five more internal USB ports available. Despite fans on the graphics card, CPU, power supply, and rear panel, the Z200 was virtually silent.

Excellent Performance

While the Intel i5-670 is a pretty fast processor, it's only a dual-core CPU, so we were not expecting to see any records fall. Boy, were we surprised. On the SPECopc Viewperf graphics benchmark, the HP Z200 equipped with the NVIDIA Quadro FX 1800 graphics card turned in some of the fastest results we've ever recorded on three of the datasets and very respectable results on the others.

On the SPECapc SolidWorks benchmark, which is more of a real-world test (and breaks out graphics, CPU, and I/O performance separately from the overall score), the Z200 results were a bit lower, reflecting its more modest components. But even on this test, this affordable HP workstation yielded great results, including the best I/O performance we've seen to date.

The AutoCAD rendering test results were also excellent, although definitely well behind those of the hyper-threaded dual-quad-core equipped systems we've recently reviewed. Because AutoCAD's mental ray rendering engine is multi-threaded, this test clearly shows the benefits of multiple cores. The Z200 averaged 135 seconds to complete the test rendering, just 10 seconds behind the Dell T3500 despite having half the number of CPU cores.

HP rounds out the Z200 with its excellent 104-key USB keyboard and a two-button optical mouse and backs the system with a limited 3-year warranty on parts and labor.

While Z200 prices start at just $729, at that price you get a Celeron CPU, 2GB of memory, integrated graphics, and a relatively small hard drive. This spring, HP also released a small form factor version of the Z200, with prices starting even lower at $699. Adding options obviously boosts the price, but even at $2,089 as tested, our HP Z200 evaluation unit delivers plenty of power without breaking the bank. It may just be the best balance of engineering workstation power and performance available today. DE

David Cohn *is a computer consultant and technical writer based in Bellingham, WA and has been benchmarking PCs since 1984. He's a contributing editor to* Desktop Engineering, *the former editor-in-chief of* Engineering Automation Report *and* CADCAMNet, *and the author of over a dozen books. You can contact him via email at david@dscohn.com or visit his website at www.dscohn.com.*

INFO → HP: hp.com

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

HP Workstation Z200 At-a-Glance

- **Price:** $2,089 as tested ($729 base price)
- **Size:** 7.0x17.9x17.6 in. (WxDxH) tower
- **Weight:** 27 pounds
- **CPU:** Intel Core i5 670 3.47GHz dual-core w/4MB L2 cache
- **Memory:** 4GB (16GB max) DDR3 1,333MHz
- **Graphics:** NVIDIA Quadro FX 1800 w/768MB
- **Hard Disk:** WD 500GB 7,200 rpm SATA
- **Floppy:** none
- **Optical:** 16X SATA dual-layer SuperMulti LightScribe DVD+/-RW
- **Audio:** High-definition integrated Realtek ALC262 audio
- **Network:** integrated Intel 82578 Gigabit LAN
- **Modem:** none
- **Drive bays:** three external 5.25-in. bays, three internal
- 3.5-in. bays
- **Ports (front):** two USB 2.0, one IEEE 1394a (FireWire), one microphone in, one headphone out
- **Ports (rear):** six USB 2.0, one audio out, one audio-in, one microphone in, two PS/2, one RJ-45 to integrated LAN, DVI-I, DisplayPort
- **Ports (internal):** five USB 2.0
- **Keyboard:** 104-key HP keyboard
- **Pointing device:** two-button optical HP scroll mouse

Simulation-Based Design Evaluation and Verification

Virtual prototypes improve design and reduce time to market.

BY PETER VARHOL

New product strategies typically involve designing the product on a computer, building one or more prototypes, testing those prototypes, and going back and adjusting the design based on the test results. This cycle can occur multiple times until the design has been proven through testing to meet requirements, and then prepared for manufacture.

While this approach usually results in a well-designed product, it can be expensive and time-consuming, with multiple cycles of design, prototype, and test before a design is final.

Today, however, there is an alternative. Engineering workstations and design and simulation software make it possible to prototype and test virtually, without ever building a physical product. The result is a quality product in a shortened design cycle that meets your market window.

How much shorter? A 2006 study from the Aberdeen Group reports that engineering design groups that use digital prototyping build half the number of physical prototypes as the average manufacturer, get to market 58 days faster than average.

Hardware and Software Show the Way

Software from design, analysis, and simulation vendors lets engineers take designs and subject them to both static and dynamic tests that provide data on characteristics such as strength and fitness for purpose. Rather than estimating, or overdesigning, engineers can make more precise evaluations of design characteristics before a single prototype is built. The result is fewer prototypes, and fewer design-test-redesign-retest cycles.

What makes it possible to design and test from your desktop workstation is the performance of today's engineering workstations. Using industry-standard Intel processors, engineers have more computational power at their fingertips than was even available on most supercomputers only a few years ago.

Inexpensive workstations and workgroup clusters are based on the 64-bit Intel® Xeon® 5600-series processor, with six cores and two threads per core. These workstations and work group clusters deliver the compute capacity of high performance computers that were only available in the data center or on a supercomputer just a few years ago.

For high-end designs and detailed simulations, server clusters can be employed. Because many high-end servers use Intel

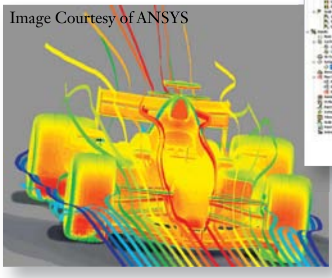
Image Courtesy of ANSYS

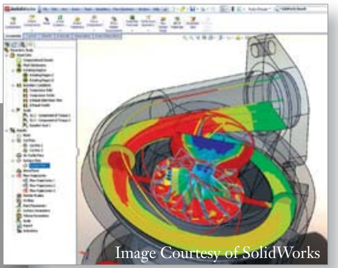
Image Courtesy of SolidWorks

Xeon processors, the chip maker has an interest in ensuring that vendor hardware scales up well, and works with both components and software. Developed with hardware and software vendors, Intel Cluster Ready lets engineering groups match engineering design applications to today's leading platforms and components. This includes servers from Appro, Dell, SGI, and Super Micro, among others.

Software is also key to producing designs that are just about ready for manufacture. Software from vendors such as ANSYS, Autodesk, and SolidWorks can enable engineers to perform both static and dynamic analysis, testing strength, tolerances, cost, and other characteristics critical to the success of the design. High-end analysis and simulation software, running on fast and inexpensive engineering workstations, can substitute for early prototypes in many projects.

New Ways of Accomplishing Product Design

In many cases, engineers can work on a design at their desktops, using feedback from simulation and test results to refine the design until it has a high probability of meeting requirements. A single set of prototypes may well be sufficient to demonstrate the merits of the design and send it on to manufacturing.

Virtual prototyping using high-performance personal workstations won't completely eliminate the building and testing of physical prototypes. It will still be necessary to build at least one set of prototypes in order to confirm the results of simulation and virtual test. But it should eliminate one or more prototype cycles, and enable engineers to deliver great products sooner. DE

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

➜ intel.com/go/hpc

➜ intel.com/go/cluster

More Graphics Power

NVIDIA's latest collection of workstation graphics accelerators have CUDA support.

BY DAVID COHN

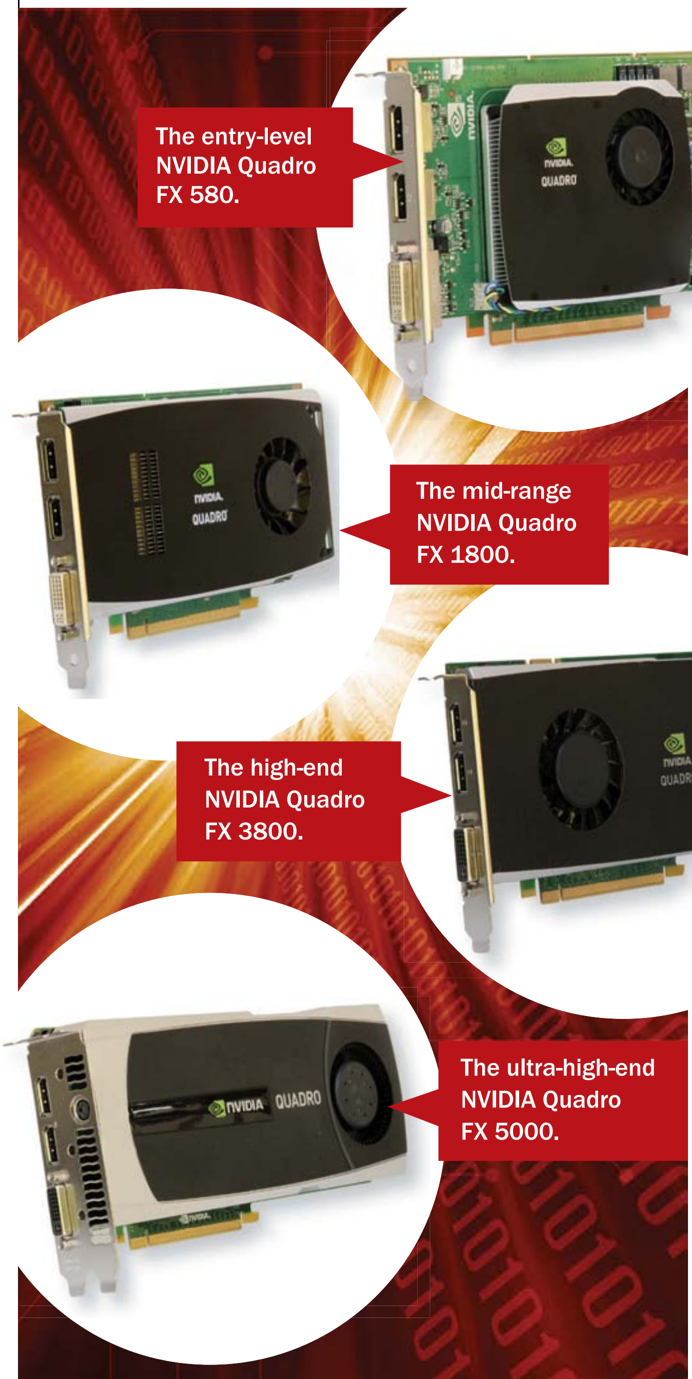

It's been two years since we last looked at the latest workstation-class graphics accelerators from NVIDIA—the Quadro FX series (see *DE* July 2008). Since then, the company has updated its entire lineup, including a trio of high-end boards based on the company's new Fermi architecture. We've included one of these new boards, the NVIDIA Quadro FX 5000, in this roundup. Its new Quadro graphics processing unit (GPU) combines parallel processing with high-performance graphics functions.

According to NVIDIA, the new second generation GigaThread Engine in the FX 5000 allows up to 10x faster context switching operations over previous generation GPUs. The GigaThread Engine also provides dual Copy Engines, which allows simultaneous transfer of data into and out of the GPU in addition to either 3D processing or compute. Fermi is also the first GPU architecture to support ECC memory, which enables the graphics card to detect and correct some types of errors in hardware as data is accessed. The new Quadro architecture also incorporates a Scalable Geometry Engine that allows multiple graphics processor to create twice the number of raw polygons in a single clock cycle, and hardware tessellation for generating detailed geometry without sacrificing performance.

We also received three other boards: the FX 3800, FX 1800, and FX 580. Although released more than a year ago, these boards represent NVIDIA's current offerings in the entry and mid-range price/performance points.

All four of the boards we tested use a PCI Express 2.0 x16 bus interface. The three older boards support DirectX 10 and Shader Model 4.0 while the new FX 5000 supports DirectX 11 and Shader Model 5.0. OpenGL support varies from 3.0 on the FX 580 and FX 1800, to 3.1 on the FX 3800 and version 4.0 on the new Quadro FX 5000. But all four of these NVIDIA graphics accelerators feature a full 30-bit display pipeline, as well as full-scene anti-aliasing, hardware 3D window clipping, and HDMI support via an HDMI adaptor.

SPEC Viewperf benchmark results for NVIDIA Quadro FX Series

	NVIDIA Quadro FX 5000 NEW!	NVIDIA Quadro FX 4800	NVIDIA Quadro FX 3800 NEW!	NVIDIA Quadro FX 3700	NVIDIA Quadro FX 1800 NEW!	NVIDIA Quadro FX 1700	NVIDIA Quadro FX 580 NEW!	NVIDIA Quadro FX 570
Manufacturer's price	$2,249	$1,999	$1,199	$1,599	$699	$699	$170	$199
Average street price	$1,765	$1,620	$818	$899	$405	$585	$150	$182
SPECviewperf 11.0 (from NVIDIA)								
catia-03	29.21	20.14	20.06	18.32	18.96	N/A	14.84	N/A
ensight-04	44.75	15.66	15.59	11.00	13.12	7.49	8.33	6.05
lightwave-01	45.29	45.33	44.93	44.79	43.81	21.06	40.83	17.82
maya-03	60.05	42.32	30.45	16.29	27.33	N/A	17.89	N/A
proe-05	6.11	5.94	5.97	5.96	5.94	N/A	5.93	N/A
sw-02	43.08	36.71	36.46	31.51	32.62	10.23	26.66	8.26
tcvis-02	30.62	17.68	17.61	15.45	15.71	6.40	10.49	3.81
snx-01	45.45	18.29	18.05	14.98	16.79	8.61	11.55	6.10
SPECviewperf 10.0 (HP xw6600)								
3dsmax-04	62.83	62.74	62.78	62.75	62.48	41.62	59.66	29.25
catia-02	78.25	70.34	70.08	66.88	68.40	48.72	62.60	40.54
ensight-03	112.78	59.47	56.80	52.90	54.36	31.02	45.37	23.46
maya-02	435.65	243.73	237.36	214.47	224.23	135.23	182.04	73.43
proe-04	57.07	55.48	56.69	55.11	55.29	41.61	53.93	31.96
sw-01	103.63	102.98	102.70	102.89	101.63	60.26	87.21	44.04
tcvis-01	80.89	40.47	40.04	34.95	35.61	18.28	23.71	9.09
ugnx-01	89.33	37.80	37.59	31.98	34.62	19.83	23.22	8.90
SPECIFICATIONS								
Bus architecture	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16
Extra power req'd	Yes	Yes	Yes	Yes	No	No	No	No
Form factor	4.38"x9.75"	4.36"x10.5"	4.38"x9.0"	4.38"x9.0"	4.38"x7.8"	4.38"x6.6"	4.38"x6.75"	4.38"x6.6"
Slots used	2	2	1	1	1	1	1	1
Max Power (watts)	152W	150W	108W	78W	59W	42W	40W	38W
PCIe version	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0
Length	3/4-length	3/4-length	2/3-length	2/3-length	2/3-length	half-length	half-length	half-length
Processors	352	192	192	112	64	32	32	16
Memory configuration	2.5GB (GDDR5)	1.5GB (GDDR3)	1GB (GDDR3)	512MB (GDDR3)	768MB (GDDR3)	512MB (DDR2)	512MB (GDDR3)	256MB (DDR2)
Memory interface	320-bit	384-bit	256-bit	256-bit	192-bit	128-bit	128-bit	128-bit
Memory bandwidth	120 GB/sec	76.8 GB/sec	51.2 GB/sec	51.2 GB/sec	38.4 GB/sec	12.8 GB/sec	25.6 GB/sec	12.8 GB/sec
Number of Dual-Link DVI Outputs	1	1	1	2	1	2	1	2
Number of Display Port Outputs	2	2	2	0	2	0	2	0
Stereo 3D Connector (3-pin)	Yes	Yes	Yes (1)	Yes	No	No	No	No
Framelock/Genlock	Yes (2)	Yes (2)	No	No	No	No	No	No
SLI	Yes	Yes	Yes	Yes	No	No	No	No
OpenGL version	4.0	3.0	3.1	2.1	3.0	2.1	3.0	2.1
DirectX/Shader Model	11/5.0	10.0/4.0	10.0/4.0	10.0/4.0	10.0/4.0	10.0/4.0	10.0/4.0	10.0/4.0
"Maximum Resolution Support (@ 60 Hz)"	2560x1600	2560x1600	2560x1600	2560x1600	2560x1600	2560x1600	2560x1600	2560x1600

Notes:
1. Requires optional 3-pin Mini-DIN bracket
2. Requires optional G-Sync option card

NVIDIA was the first graphics card vendor to create a unified graphics architecture, in which processing power can be dynamically allocated to compute, geometry, shading, and pixel processing to optimize GPU performance. The entry-level FX 580 features 32 such processors while the GPU in the new Quadro FX 5000 provides 352 Compute Unified Device Architecture (CUDA) processors.

The new NVIDIA Quadro FX boards also include NVIDIA's Application Configuration Engine (ACE), which automatically adjusts graphics settings for optimized application performance. All four boards also provide a single dual-link DVI-I output and a pair of DP connectors. Only two of the three digital outputs can be active at a time; the DVI port can also support a single analog monitor.

The latest boards also support the new NVIDIA CUDA GPU-computing software architecture. Some types of computing tasks have parts that are better suited for CPU computing while other parts are better suited for being run on the GPU. CUDA provides a standard C language interface so that programmers can synchronize hybrid programs running on both CPUs and GPUs.

The New Quadro FX Family

At the entry level, the NVIDIA Quadro FX 580 provides 512MB of GDDR3 memory. The extra memory and more powerful GPU, with 32 NVIDIA CUDA processors yield double the bandwidth of the board it replaces. The FX 580 requires a single PCIe x16 slot and consumes 40 watts of power.

The mid-range is the sweet spot. The NVIDIA Quadro FX 1800 has doubled the number of GPU processors—from 32 in the older FX 1700 to 64 in the FX 1800—and a 192-bit memory interface. With 768MB of GDDR3 memory, the memory bandwidth triples to 38.4 GB/second, resulting in graphics performance anywhere from two to three times that of the older board. Yet the FX 1800 requires a single PCIe x16 slot and consumes just 59 watts.

At the high end, NVIDIA replaced the FX 3700 with the Quadro FX 3800. This board comes with 1GB of GDDR3 memory—double that of the older board—while retaining the same 256-bit memory interface. This results in the same 51.2 GB/second memory bandwidth as the older board, but with 192 parallel processors—80 more than its predecessor—the NVIDIA FX 3800 easily outperforms the older board. This board also supports stereoscopic 3D output (with an optional 3-pin Mini-DIN bracket) as well as NVIDIA's Scalable Link Interface (SLI). SLI technology enables users to leverage multiple NVIDIA Quadro FX graphics boards to dynamically scale graphics performance. The Quadro FX 3800 requires just a single PCIe X16 slot, but it consumes 108 watts and requires an auxiliary connection to the system power supply.

At the ultra high end, the FX 5000 features 352 CUDA parallel processing cores and 2.5GB of GDDR5 memory to deliver up to eight times the performance of the previous generation Quadro FX 4800 when running computationally intensive applications such as ray tracing, video processing, and computational fluid dynamics. The FX 5000 provides two DP connections and a Dual-link DVI connector. It also supports stereoscopic 3D output, and has the 3-pin stereo connector built-in. The FX 5000 also supports NVIDIA's SLI technology, and NVIDIA also offers an optional G-Sync card for frame lock and genlock. All of this power results in a bigger board consuming 152 watts, requiring an auxiliary power connection and the space of the adjacent expansion slot.

Pricing

- **NVIDIA Quadro FX 580:** $170 (MSRP), $150 (average street)
- **NVIDIA Quadro FX 1800:** $699 (MSRP), $405 (average street)
- **NVIDIA Quadro FX 3800:** $1,199 (MSRP), $818 (average street)
- **NVIDIA Quadro FX 5000:** $2,249 (MSRP), $1,765 (average street)

Benchmarking the Boards

We tested the four NVIDIA Quadro FX boards using the same HP xw6600 workstation equipped with a pair of 3GHz Quad-Core Xeon E5450 processors, so all of our results are directly comparable. Since NVIDIA updated its driver software for the newest boards, we also retested the previous generation of NVIDIA graphics cards using the updated driver, to get an accurate comparison of the performance improvements. This is the same system we recently used to test the latest ATI FirePro boards (see *DE* October 2010), so we can also compare the relative performance of boards from both manufacturers.

All tests were performed using both version 10 of the SPEC Viewperf benchmark (spec.org) and version 11, which was recently released, both at a resolution of 1280x1024.

Based on our own independent results, the new generation of NVIDIA Quadro FX boards do indeed surpass the performance of the older boards. In fact, in most instances, the newer board at the lower price point beat the performance of the next step up from the previous generation.

With these new Quadro FX boards, the competition between NVIDIA and AMD remains as fierce as ever, giving you, the user, lots of excellent choices at every price point. DE

David Cohn *is a computer consultant and technical writer based 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. You can contact him via email at david@dscohn.com or visit his website at dscohn.com.*

INFO → NVIDIA Corporation, nvidia.com

→ Advanced Micro Devices, Inc. (AMD), amd.com

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

Optimizing **OpenFOAM** Productivity in Clusters

This study concludes that, when configured correctly, users can see cost savings in addition to higher productivity by using HPC clusters.

BY GILAD SHAINER, TONG LIU, PAK LUI, JEFFERY LAYTON AND SCOT SCHULTZ

Computational fluid dynamics (CFD) enables the study of the dynamics of things that flow, typically fluids, by generating numerical solutions to a system of partial differential equations that describe fluid flow. The primary purpose of CFD is to better understand qualitative and quantitative physical phenomena in the flow, which can be used to improve engineering design.

CFD simulations are typically carried out on high-performance compute clusters (HPCC) or other large computational systems because they require a compute resource that can handle complex problems with large computational requirements, large memory requirements, and possibly large storage requirements. Because CFD codes are scalable, HPCCs typically offer the best price/performance solution for running simulations.

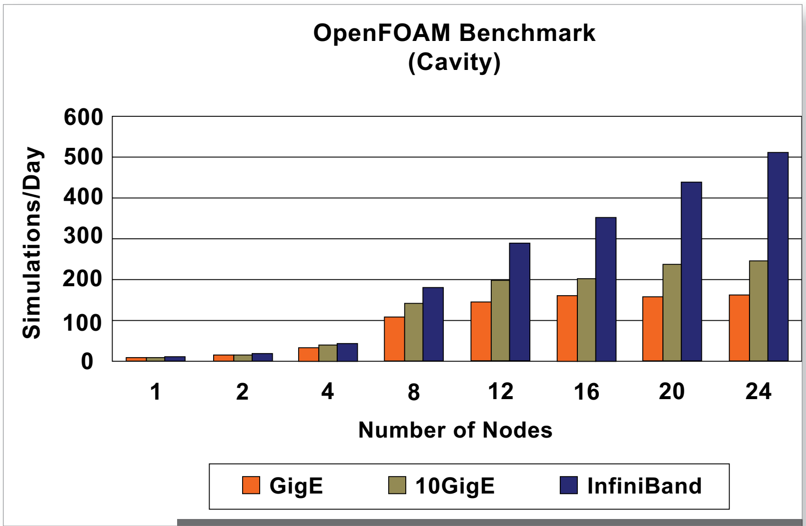

FIGURE 1: Productivity of OpenFOAM with InfiniBand vs. 10GigE and GigE.

Connecting Compute Nodes

A basic HPCC system consists of off-the-shelf servers, a high-speed interconnect and storage. The interconnect has a great influence on the total cluster performance and scalability. A slow interconnect will cause delays in data transfers between servers and between servers and storage, causing poor utilization of the compute resources and slow execution of simulations. An interconnect that requires a large number of CPU cycles as part of the networking process ("onloading") will decrease the compute resources available to the applications—and therefore will slow down and limit the numbers of simulations that can be executed on a cluster. This will limit the cluster scalability because when the number of CPUs increases (the number of cores), a corresponding increase in CPU cycles will be required to handle networking.

The InfiniBand high-speed interconnect provides fast communication between compute nodes to constantly feed the CPU cores with data and help eliminate idle times. Moreover, InfiniBand was designed to be fully offloaded, meaning all the communications are being handled within the interconnect with no involvement from the CPU. This further enhances the ability for the code to scale up, with close to linear performance when more compute resources are required.

OpenFoam CFD Applications

Open Field Operation and Manipulation, or OpenFOAM, is an open-source CFD application for simulating fluid flow problems. The broad physical modeling capabilities of OpenFOAM have been used by the aerospace, automotive, biomedical, energy and processing industries.

The original development of OpenFOAM began in the late 1980s at the Imperial College, London, motivated by a desire to find a more powerful and flexible general simulation platform than the de-facto standard at the time, Fortran. Since then, it has evolved by exploiting the latest features of the C++ language, having been effectively re-written several times over. The predecessor, FOAM, was sold by Nabla Ltd. before being released to the public under the general public license in 2004. It is now developed primarily by OpenCFD Ltd., with assistance from an active user community.

OpenFOAM is among the first major scientific packages written in C++ that provides relatively simple, top-level, human-readable descriptions of partial differential equations. It uses polyhedral cells. This functionality is a natural consequence of the hierarchical description of simulation objects, and it is the first general-purpose CFD package to be released under an open-source license.

At the core of any CFD calculation is a computational grid, used to divide the solution domain into thousands or millions of elements used for solving the partial differential equations. OpenFOAM uses a finite volume approach to solve systems of partial differential equations on any 3D unstructured mesh of polyhedral cells. In addition, domain decomposition parallelism is fundamental to developing solvers, and OpenFOAM has the ability to decompose the domain for good scalability.

For performance and scalability analysis, OpenFOAM provides a set of solvers and utilities representing typical industry usage for performing pre- and post-processing tasks. It also has a set of libraries of physical models. Mesh generation is simplified by representing a cell as a list of faces, and each face as a list of vertices. OpenFOAM applications handle unstructured meshes of mixed polyhedra with any number of faces: hexahedra, tetrahedral or even degenerate cells. The scalability analysis of OpenFOAM was done using the lid-driven cavity flow benchmark case.

Using OpenFOAM to Analyze Scalability

Using InfiniBand, we were able to achieve higher productivity in every job size (number of cores) compared to Gigabit Ethernet (GigE) and to 10 Gigabit Ethernet (10GigE). In particular, at 24 server nodes (192 cores), we achieved 219% higher performance vs. GigE and 109% higher than 10GigE. The lid-driven cavity flow benchmark timings are shown in Figure 1 as the number of jobs run per day, as a measure of productivity.

The ability to execute more than twice the number of simulations per day solely because of the interconnect demonstrates the importance of choosing the right elements when building HPC clusters. With the same server configuration, a user with an InfiniBand-connected cluster will be able to reduce the engineering design time by more than half.

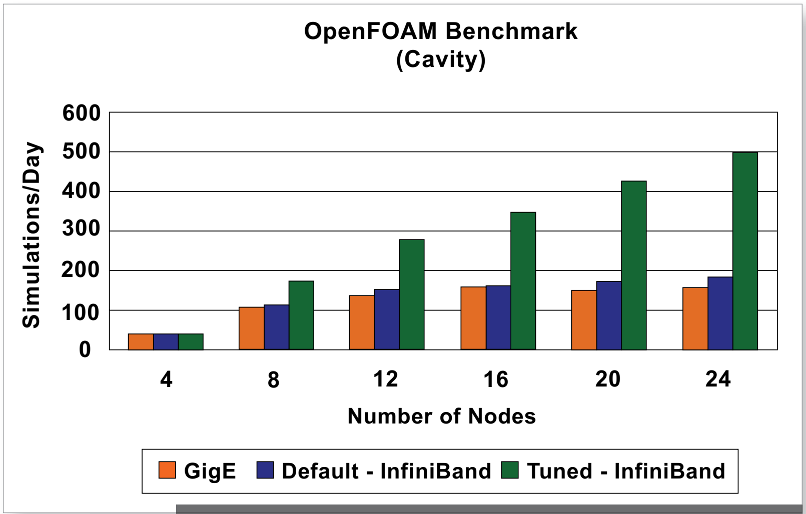

FIGURE 2: Productivity of OpenFOAM with optimized InfiniBand vs. the out-of-the-box settings.

We also noticed that OpenFOAM did not scale with Gigabit Ethernet beyond 16 servers (128 cores) for this problem, which means that any server added to a 16-node cluster with Gigabit Ethernet will not add any performance or productivity to OpenFOAM tasks.

In recent CFD and FEA application profiling analysis, the dependency on highest throughput and lowest latency for optimal scalability was shown. Many of the CFD codes generate large Message Passing Interface (MPI) messages, in the range of 16KB to 64KB, for data transfer, and use fast small messages for synchronizations. When cluster increases in server count, or node count, the number of messages both large and small increase as well; this increases the burden on the cluster interconnect. This is also where Ethernet becomes the bottleneck and limits any performance gain, or productivity. The profiling results of lid-driven cavity flow are shown in Figure 3.

The most used MPI functions tend to be MPI_All-

SCALABILITY ANALYSIS SPECS

The scalability analysis was performed at the HPC Advisory Council compute center, using the following components:

- Dell PowerEdge SC 1435 24-node cluster with Quad-Core AMD Opteron 2382 (“Shanghai”) CPUs
- Mellanox InfiniBand ConnectX HCAs
- InfiniBand

Each node had 16GB memory, DDR2 800MHz. The operating system was RHEL5U3, InfiniBand drivers OFED 1.4.1, OpenMPI-1.3.3 and OpenFOAM 1.6. For testing purposes, the input dataset of lid-driven cavity flow—mesh of 1000 x 1000 cells, icoFoam solver for laminar, 2D, 1000 steps—was used.

CONFIGURING OPENFOAM

The results shown in Figure 1 are not out-of-the-box results that can be achieved with OpenFOAM. The default OpenFOAM binary is definitely not optimized for InfiniBand-connected clusters. For some reason, when InfiniBand is present, the default OpenFOAM binary will still use the Ethernet network for partial data transfers, and does not take full advantage of the InfiniBand networking capabilities. Choosing a precompiled Open MPI doesn't solve the issue, and recompiling OpenFOAM properly is necessary. There are two ways to compile OpenFOAM. Please note that Option 2 takes longer to create—more than four hours:

Option 1:

- Modify OpenFOAM-1.6/etc/bashrc to use MPI entry rather OPENMPI (WM_MPLIB:=MPI).
- Change MPI entry within settings.sh to system OpenMPI (export MPI_HOME=/usr/mpi/gcc/openmpi-1.3.3).
- Add the following to wmake/rules/linux64Gcc/mplib (PFLAGS = -DOMPI_SKIP_MPICXX; PINC = -I$(MPI_ARCH_PATH)/include; PLIBS = -L$(MPI_ARCH_PATH)/lib64 –lmpi).

Option 2:

- Keep the default OPENMPI entry in bashrc.
- Modify default Open MPI compiler option in ThirdParty-1.6/Allwmake (refer to Open MPI website for full compiling options).

With the proper optimization, InfiniBand-delivered productivity will be improved by 172%, as shown in Figure 2, and will deliver the performance advantages shown in Figure 1.

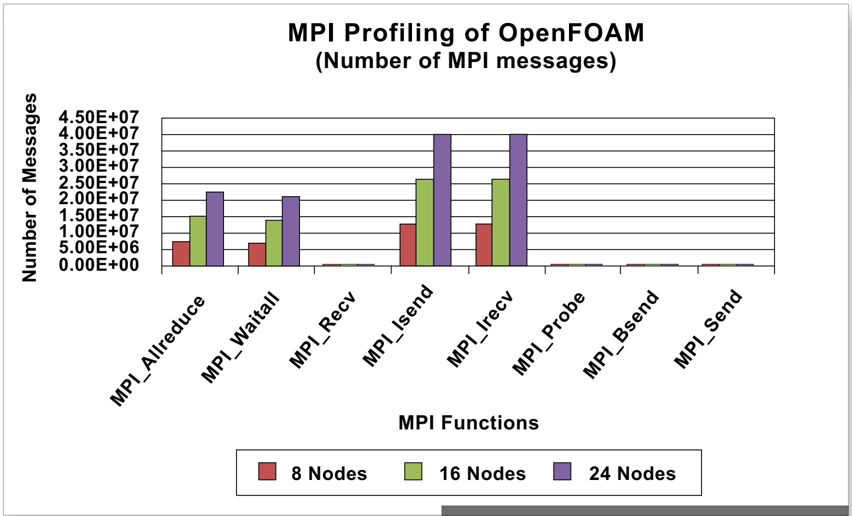

FIGURE 3: OpenFOAM MPI profiling.

reduce, MPI_Waitall, MPI_Isend and MPI_recv. The number of MPI functions increase with cluster size, as expected. The majority of the MPI timing overhead is within the MPI_Allredcue collective operation. The recent technology development between Mellanox Technologies and Oak Ridge National Laboratory of MPI collectives offloads can increase OpenFOAM productivity through better, efficient execution of MPI_All reduce and MPI_Waitall—and more importantly, by enabling asynchronous progress. In other words, by overlapping of the MPI compute cycles and the MPI communications cycles.

OpenFOAM provides an open-source CFD simulation code, and has demonstrated a high degree of parallelism and scalability. This enables it to take full advantage of multi-core HPC clusters. InfiniBand shows greater efficiency and scalability with cluster size, providing twice the productivity. .

The study shows that economical integration of the CPUs and the InfiniBand network can save up to $8,400 a year in power to achieve the same number of application jobs vs. GigE, and up to $6,400 to achieve the same number of application jobs vs. 10GigE. This is based on 24-node cluster configuration. As cluster size increases, more power can be saved. For the user, this translates into a reduction of operating expenses. DE

Gilad Shainer *is a high-performance computing evangelist who focuses on HPC, high-speed interconnects, leading-edge technologies and performance characterizations.* **Tong Liu** *is the manager of application performance at Mellanox.* **Pak Lui** *works at Mellanox Technologies as a senior HPC applications performance engineer.* **Jeffrey Layton** *is the enterprise technologist for HPC within Dell.* **Scot Schultz** *is a board member of OpenFabrics Alliance, representing AMD.*

INFO ➜ AMD: AMD.com

➜ **Dell:** Dell.com

➜ **HPC Advisory Council:** HPCAdvisoryCouncil.com/pdf/OpenFOAM_Performance_Analysis.pdf

➜ **InfiniBand Trade Association:** InfiniBandTA.org

➜ **Mellanox Technologies:** Mellanox.com

➜ **Oak Ridge National Laboratory:** ORNL.gov

➜ **OpenCFD Ltd.:** OpenFOAM.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.

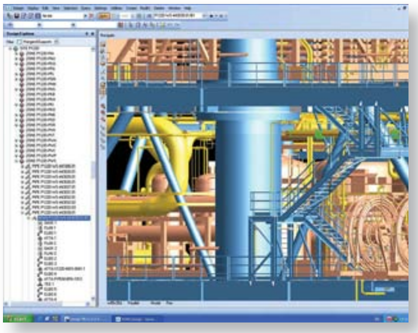

AVEVA Plant 12.0.SP6 Released

AVEVA plant supports Windows 7; delivers performance and functional enhancements.

Industrial sites are like little cities. They have layer upon layer of complexity that, when well-managed, makes the complexity seem ordinary. The trick, of course, is designing, building/revising, then engineering the whole thing well. Not an easy to execute or manage job.

AVEVA has just released version 12.0.SP6 of AVEVA Plant, its a set of integrated applications for plant engineering, design, and design management. Support for Windows 7 is the headline for this new version, but there are functional extensions, a spiffier interface and general performance optimizations.

MORE → deskeng.com/articles/aaaxzm.htm

BOXX Debuts 4880 XXtreme Workstation

Claims it's the "world's fastest" workstation; includes new NVIDIA Quadro 5000.

You can get yourself a new workstation, then you find it is providing just marginal improvements. This was common when we were first crossing over to 64-bit computing. Then it was the applications that were not leveraging multi-threading. Today, the software is cramming stuff into the hardware.

The folks at BOXX Technologies have made it their mission to get the most of your hardware and software. And just the other day, it introduced its new 3D BOXX 4880 XXtreme workstation. BOXX calls the 3D 4880 XXtreme a "visual supercomputer." I call it a screamer.

MORE → deskeng.com/articles/aaayas.htm

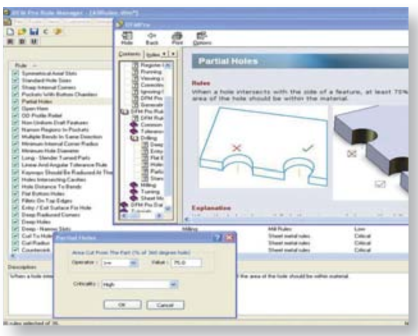

Geometric releases DFMPro version 2.2 for Pro/ENGINEER

Includes support for Wildfire 5 and Windows 7.

When I was 30, ratcheting back on the beer gave me a quick loss of five pounds. Now, I'm lucky if I get a pound. If I had only been consistent in applying the rules of healthy living from the beginning, I would not be paying so dearly today.

Kind of sounds like the new product design process, no? If you had just designed your parts from the start with a consistent set of rules that made late changes to the design less likely, you could eliminate all that time and money fixing them later. That's why you should consider something like Geometric's DFMPro design for manufacturability (DFM) tool.

MORE → deskeng.com/articles/aaaybp.htm

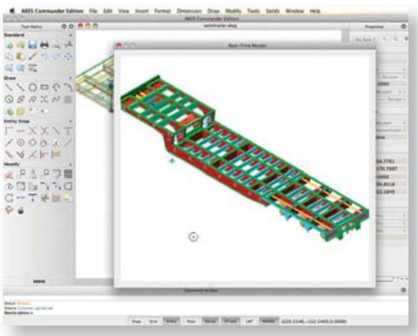

Multi-Platform CAD System Now Runs on Mac OS X

Application said to be the first pro CAD solution for Mac, Windows, and Linux.

Back when Apple first released dolled-up UNIX in the form of OS X, a chill ran up the spine of Mac mavens. The hope was that MCAD, CFD, FEA, and other engineering software would join LabVIEW and a few others with versions for the Mac. And indeed, that has happened. Too slowly. And MCAD mostly has been a laggard. Until now.

Graebert, a developer with a long history, has announced that Edition 1.0 of its low-cost ARES Commander 2D/3D CAD software for Mac OS X. There are three cool things to know about ARES for Mac OS X. Read all about them in my online message.

MORE → deskeng.com/articles/aaaycn.htm

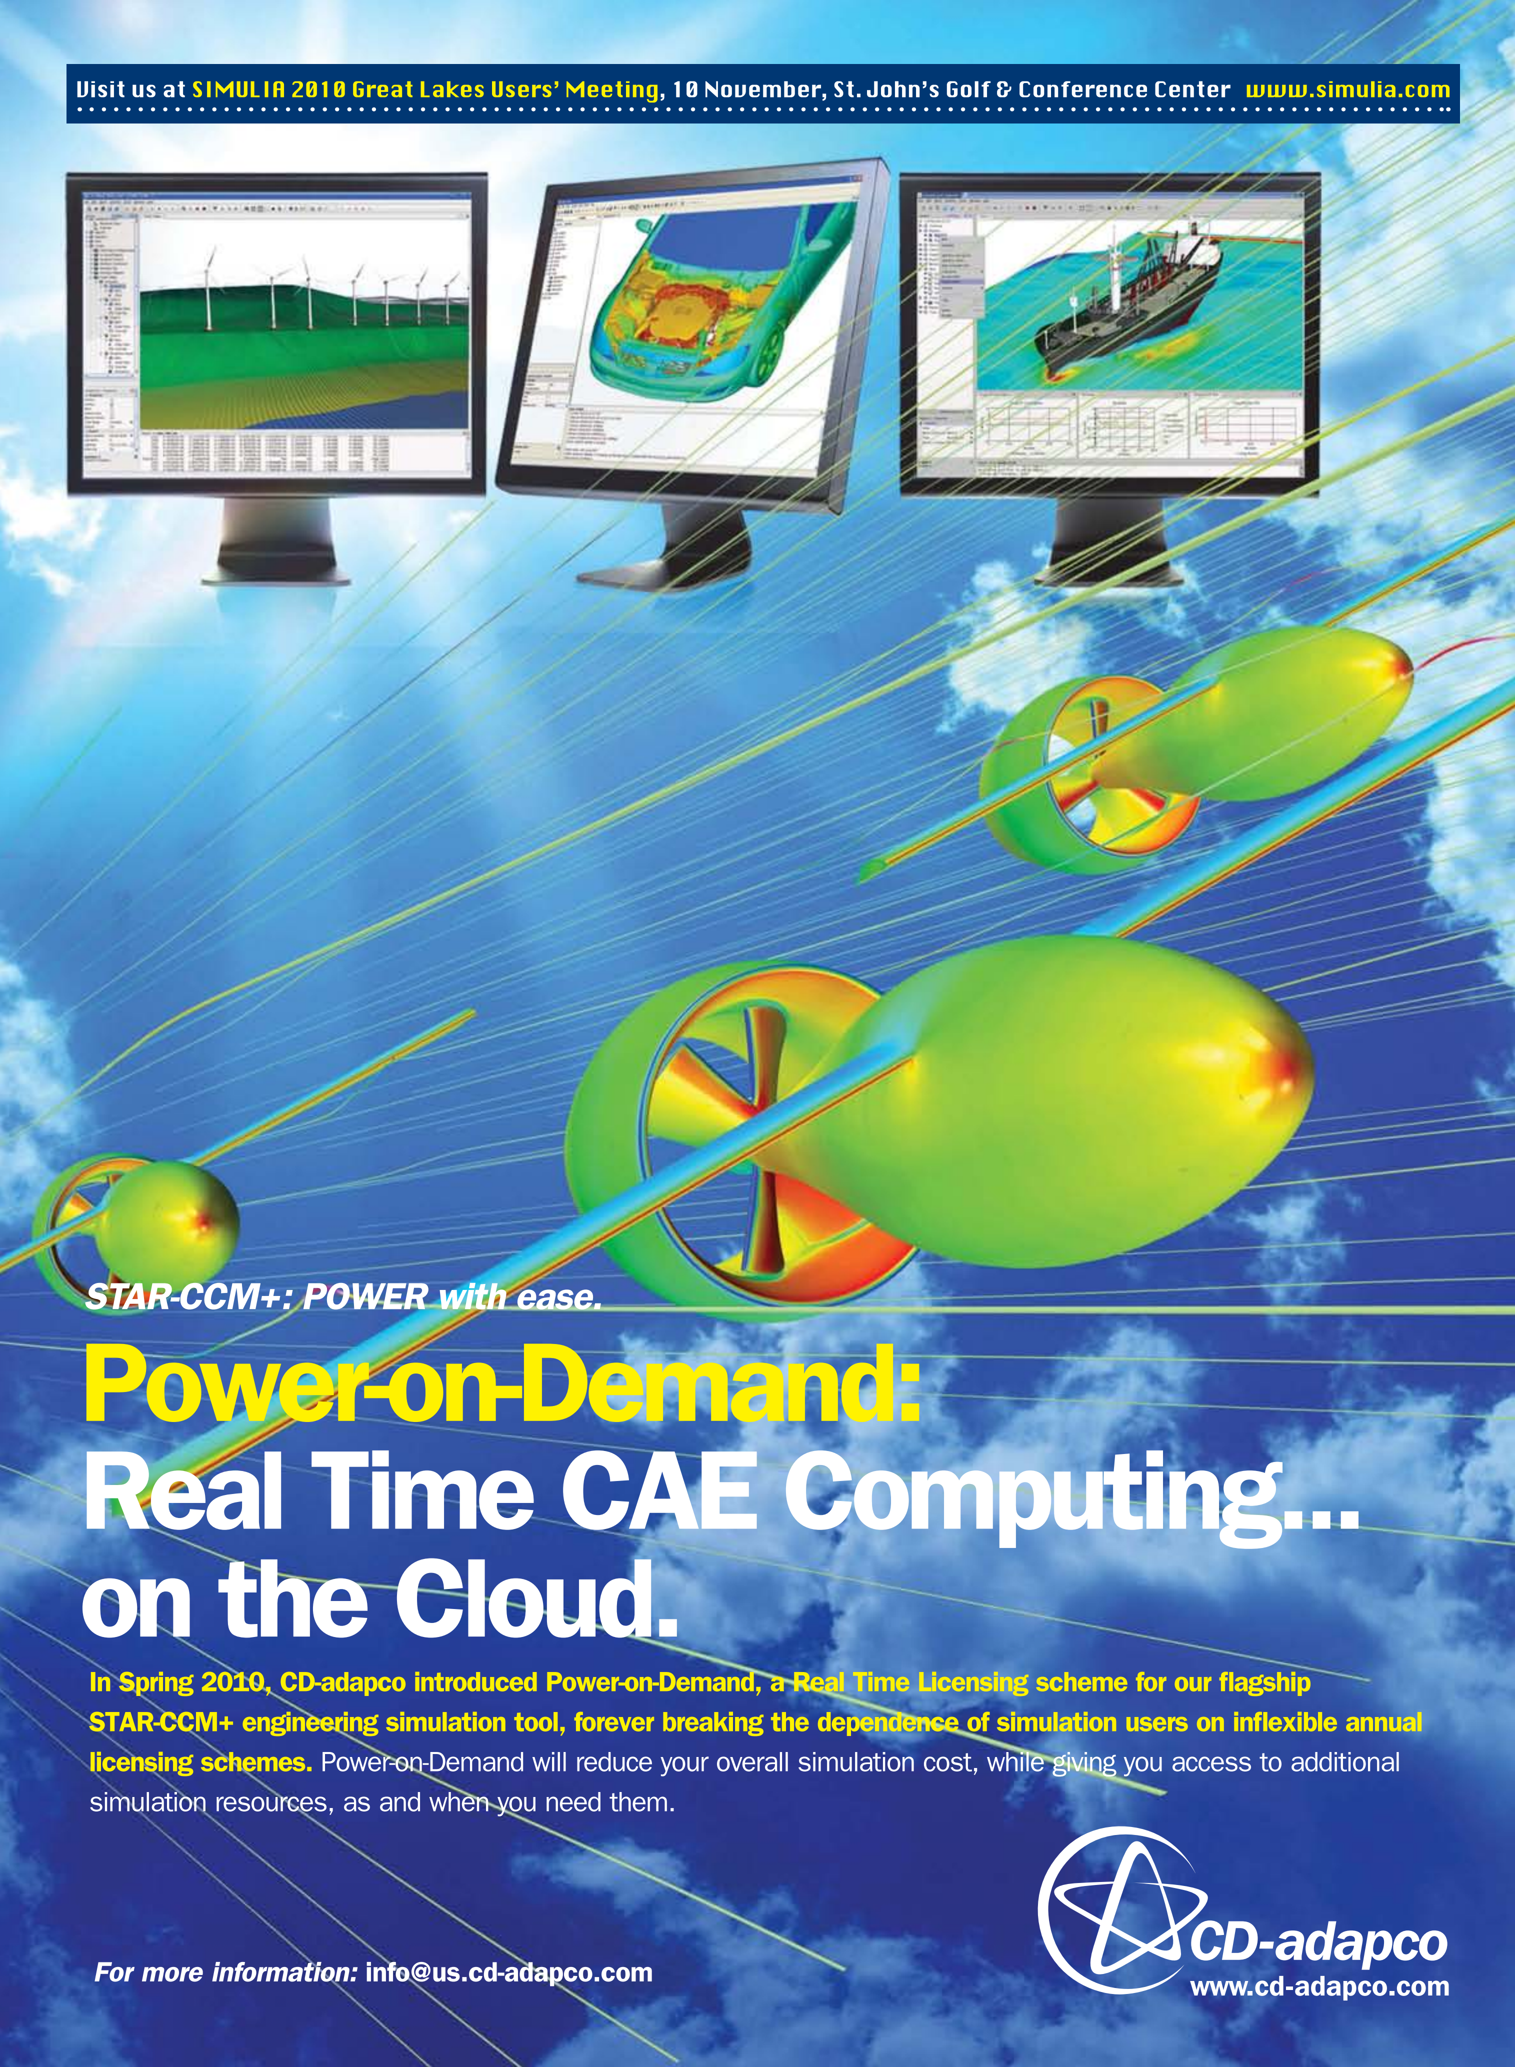
Visit us at SIMULIA 2010 Great Lakes Users' Meeting, 10 November, St. John's Golf & Conference Center www.simulia.com
STAR-CCM+: POWER with ease.
Power-on-Demand:
Real Time CAE Computing...
on the Cloud.
In Spring 2010, CD-adapco introduced Power-on-Demand, a Real Time Licensing scheme for our flagship STAR-CCM+ engineering simulation tool, forever breaking the dependence of simulation users on inflexible annual licensing schemes. Power-on-Demand will reduce your overall simulation cost, while giving you access to additional simulation resources, as and when you need them.
For more information: info@us.cd-adapco.com
CD-adapco
www.cd-adapco.com

Improving Product Development with CAE

How automotive industry simulation expertise could benefit other industries.

BY DR. BIJAN KHATIB-SHAHIDI

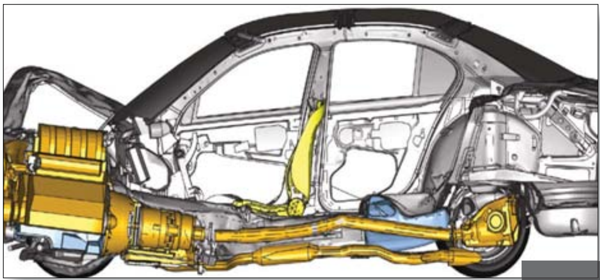

Fig. 1 and 2: The crash management system is designed and maintained during the crash scenario, including the crumple zone areas in the front and rear. *Images courtesy of BMW.*

Editor's note: *This article has been adapted from an article that originally appeared in* benchmark*, a NAFEMS publication. For more information, visit nafems.org.*

Despite the recent downturn in consumer demand and the serious economic challenges across industry worldwide, today's modern automobiles have benefited from the extensive use of computer-aided engineering (CAE) to make them quiet, durable, comfortable, stable and safe. The design and release engineers, who are responsible for signing-off on the adequacy of the component, part or system that they send to manufacturing for production, have particularly benefited from modern finite element analysis (FEA) and computational fluid dynamics (CFD) software and analysis. Their colleagues in CAE organizations conduct the analysis to better the design for the public in general. I also need to mention that in some smaller organizations, the role of the design and CAE engineer is the same, and the same engineer does both the analysis, design and release.

However, that does not take anything away from the discussion on the benefits that CAE provides. In my recent role as a vehicle CAE manager at a major automotive original equipment manufacturer (OEM), I experienced and watched the growth of FEA and CFD, in both capability and utilization, on engineering applications and design.

Using FEA and CFD, our team was able to create full-vehicle model representations of cars and trucks that included many parts. We swapped different sized tires, suspensions, multiple powertrains, and chassis systems to represent a variety of vehicle configurations. We could represent these models as if they were real prototypes. They included acoustic representations to measure noise predictions. We added sub-system representations of multiple engine/ transmissions, axles and drivelines, different trim levels, and excited them with various load conditions such as road profiles, wind loads (winter and summer modelling, and different yaw angles), engine loads, driveline imperfection scenarios, and so on.

New capabilities, such as one from Abaqus software, allowed rolling tire capabilities and sub-structuring, as well as the ability to switch from implicit to explicit analysis and vice versa to mimic road noise, vibration and harshness (NVH), and durability analysis.

I need to reiterate that these models enabled our designers and development engineers to identify, early in the design phase, any necessary "fixes" that were required to achieve consumer and regulatory targets, and also to avoid back-breaking vehicle recalls. This avoided (for the most part) the building of physical prototyping, which were then only used for validation and comparison. Such CAE simulations also enabled the teams to provide scientific cost forecasts used to report to the senior management before they agreed to sign off on any production go-ahead.

Improving Efficiency

There is huge competition among the 100+ car industries across the globe, which is due in no small part to the conver-

Fig. 3: The rail crush example is run in CZone for Abaqus, which is an add-on product to Abaqus/Explicit. *Image courtesy of Engenuity, Ltd.*

gence of design and styling, engineering content, and the never-ending demand from the costumers to have more and more included in the cars they purchase. Because of this, the companies face major economic pressures on a daily basis. In order for them to stay afloat, these manufacturing companies are more focused than ever on improving engineering efficiency, lowering development costs, and accelerating product innovation faster than the competition next door.

I have seen FEA/CFD and other simulation solutions playing an increasingly prominent role in helping every industry to design better, optimize more often, and reuse the existing design in a different platform (with the new morphing technologies) to cut costs and to become as competitive as possible. There is no other way. The convergences in the economic conditions have made this problem universal, making it more difficult to build distinctive products. This situation is no longer unique to the automotive industry.

Any industry that is in the business of manufacturing can benefit from CAE. Nowadays, every country is building cars, airplanes, washing machines and food processors, so it is necessary to be design distinctive, efficient, and fast to market. CAE tools can play a prominent role in achieving most of these goals.

Nonlinear FEA Closer to Reality

Advances in commercial FEA technology are enabling engineers to get closer than ever to simulating realistic behavior through the inclusion of nonlinear effects in processes such as rolling tires, and in materials such as rubber, plastics, and exotic metals and composites. These materials are used in design along with standard metals such as steel or aluminium. The nonlinear FEA capabilities dramatically improve the accuracy of FEA results when they are compared to linear simulations, and for instance, in noise and vibration stress and/or durability, thermal comfort and safety, and crash and blast models. Head-on and offset car impact are all represented as detailed models to fulfil the safety and regulatory missions with the simulation.

For example, consider the complex non-linear crash analysis computations done by BMW. The crash management system is well designed and maintained during the crash scenario, including the crumple zone areas in the front and rear, as demonstrated in Figures 1 and 2. Inclusion of various “non-linearities,” such as composite material behavior and the loads in the simulation models make the results come closer than ever to the physical reality. Figure 3 represents the non-linear crash simulation that is done on a composite structural member subjected to offset axial crushing.

The complex composite analysis results have been compared to the test set up in the lab with great accuracy. The rail crush example (Figure 3) is run in CZone for Abaqus, which is an add-on product to Abaqus/Explicit. Modelling conditions such as variability on what percentage of welds could be missing in the manufacturing process during the build process are important. Adding the stochastic and probability modelling on top of the aforementioned non-linearity makes the analysis more valuable than ever.

Hardware Powers Increasingly Sophisticated Software

The computer hardware on which advanced analyses—such as full crash analysis and NVH, thermal comfort, CFD, safety or blast analysis, are run —has also come a long way in recent years, and will become even more powerful in the future. An additional benefit of less expensive, parallel and faster computing is that the design of experiments (DOE), optimization, stochastic and probability analysis can now become a natural extension of the engineer’s analysis and modelling process.

In addition to the computing power, the capabilities available to the CAE engineer have increased in the last few years in tandem with the improvements in the parallelization of algorithms in FEA, CFD and multiphysics software, in combination with other enhancements, such as accurate sub-structuring that allows the models to get even larger in size and content. These refinements are allowing engineers to build ever more realistic FEA/ CFD models with increasingly finer and finer meshes.

One other noteworthy remark is that not only have software and hardware come a long way from its power of representation in multi-physics, it is also getting simpler to use. A case in point is the utilization of pre-processors such as Hypermesh (From Altair Engineering), ANSA (from Beta CAE), Abaqus CAE (From Simulia) and the Mesh Wrapping capability in STARCCM+

Fig. 4, 5 and 6: Examples of what can be done with software. You transfer results such as the pressure fluctuations to a structural software, such as Simulia's Abaqus, NASTRAN or acoustic software such as ACTRAN, for further analysis, such as noise predictions inside the cabin of a car. *Images courtesy CD-adapco.*

(CD-adapco) to create complex models quickly, and even to get them to communicate and pass results between different software packages. All in all, it is allowing the engineer to learn and adapt quickly to changing economic conditions.

I remember the first safety model I built had about 10,000 elements to represent a C or D size sedan in the late '80s. That safety model, although helpful in designing the trigger mechanisms for better crash repeatability and axial load management, missed the crush distance by quite a lot when the model's crush distance was compared to its physical prototype counterpart. That was due to the limitation of the computing power that existed within the Cray XMP, one of the most powerful computers that existed at the time. We could only get qualitative results to fine tune a model and help with the component detail design to strategically place the front rail trigger mechanisms and optimize the crush absorption.

However, we could not sign off on a design with such crash models or even totally eliminate a prototype. Today, because computers are more powerful and available, one can do wonders with non-linear analysis even on household computers.

Crash models with 2 million elements or more are common nowadays. The crash and safety engineers at Toyota have reported that they have created models reaching more than 14 million elements, with elements as small as 1 to 2 mm in size. These types of models include so much detail that they can then be used to decide, for example, where the proper location is for airbag sensors and what amount of deceleration should trigger them. CAE models provide a great deal of detailed information, which in turn improves the quality of designs, ensuring better results decade after decade. This results in a general reduction in the need for expensive physical prototypes, which cost hundreds of thousands of dollars and are destroyed in fraction of a second.

The NVH linear models of the '90s era only ever had around 100,000 linear representations with a number of super element structures to predict noise levels at the 80 to 100Hz frequency coverage. These models are now represented by more than 2 million elements in a single shot, including rolling tires and the frequency coverage of 300 to 450Hz, thanks in part to software such as sub-structuring and automated multi-level sub-structuring (AMLS), as well as the availability of inexpensive powerful computers available to the greater public. These types of models can help designers and engineers to place "Quiet Steel," multi-layer laminates in strategic locations, in the cars and trucks to quiet down the noisy panels. The models can also help in the realization of the type of trim materials, such as carpet and plastics, which can be used. This trim modelling capability was not possible or readily available in the '80s and early '90s.

On the CFD front

In previous times, some of our CFD engineers would take two to three weeks to create 2.5 million finite volume cells in an aerodynamics section, which would take a few days to run on multiple CPUs. The results would vary by about 0.020 when we compared the coefficient of drag prediction to the wind tunnel tests. A higher number of cells would provide a better prediction, but it would take even longer for the analysis to complete. For this reason, engineers often preferred to go to the wind tunnel for quicker results as well as a more accurate representation.

These days, it is extremely easy to create CFD models. I recently created an external aero model of a full vehicle representation with the inclusion of the occupants, including the windows open condition, with about 30 million cells. This took only about three hours to complete, and a few hours later (in the same day I might add) I would get the results back with accurate aero predictions, and certainly within the variability range of a shape and content. Using CCM+ for model creation is very easy; all one needs is CAD data that is fairly air- and watertight. CCM+ is also forgiving if there are few holes, which it can repair after one takes it to its pre-processing mode. Once in there, you can fill the unwanted holes and move on to wrap the surface, and then with a few clicks you can create an "automated" volume mesh. It is relatively straightforward, and as such the learning curve is very fast even for first time CCM+ users, although obviously it helps if one has some background in the theory and its application.

I could not show my analysis results and pictures for confidentiality reasons, but instead I have included some publicly released pictures from CD-adapco. Figures 4, 5 and 6 are examples of what one can do with their software and even in some more advanced applications. One can transfer results such as the pressure fluctuations to a structural software, such as Simulia's Abaqus, NASTRAN or acoustic software such as ACTRAN, for further analysis, such as noise predictions inside the cabin of a car due to structural vibrations from that of the dynamic fluid pressure induced upon it.

Those willing to invest more in realistic simulation technology and staff will certainly benefit.

Blast and Explosion

Another extension of the CAE capabilities is in the area of safety design for blast and explosion mitigation. This is a technology that combines the Lagrangian mesh that is traditionally used for typical crash modelling and analysis with that of the Eulerian mesh that is used to represent the blast of explosives that travels through the air or ground, and delivers a supersonic shock to the vehicle in question. The aim is to design a vehicle that is safe under the blast explosion through analysis and modelling. In some sense, one can think of it as the fluid-solid interaction (FSI) of the air blast and that of the vehicle that gets the shock through the fluid. Once again, this type of analysis requires a great deal of computer horsepower because the models can contain very large amounts of nodes and elements.

This type of analysis can become even more complex when one adds erosion, regularization and complex tri-axiality failure criteria on elements that erode or blast away, which could easily bring models upward of 15 million elements for blast type applications.

Automotive Advances Benefit Other Industries

The automotive industry is certainly one of the major staging grounds for large-scale FEA/CFD in terms of degrees of freedom. The use of vehicle durability, NVH and safety analyses have grown tenfold in the last decade, with enabling tools such as sub-structuring on calculating modal presentation. There are also several attempts going on to use MDO (multi-disciplinary optimization) analysis among these attributes.

We have come to the realization that the knowledge and experience gained in the automotive industry can become quite beneficial in other industries as well. The CAE engineers in the automotive industry have been eager to share their results and spread the success and the news of these capabilities at regional and world seminars and conferences. So the "lessons learned" in automotive simulation capability are feeding directly into heavy vehicle, military, defence, off-road, mining, aerospace and shipbuilding engineering.

Newer industries, such as life sciences, are learning that the design challenges they face in noise and vibration control of devices and machinery, such as the breathing apparatus design that needs to be "whisper quiet" during the patient's utilization, can also benefit from the CFD and NVH know-how that automotive CAE engineers can provide. We are also likely to see increased knowledge transfer to medical, pharmaceutical, and other industries.

A Look Ahead

I see two new frontiers. First, the explicit FEA methods that have become part of crash/safety and blast and penetration simulations in recent decades will be increasingly applied to noise and vibration applications, and obviously in blast and mitigation. This will result in an increased demand for even faster computing hardware and better integration of implicit and explicit FEA techniques to help design the lighter-weight, efficient, yet comfortable vehicles of tomorrow. Software that communicates and transfers implicit and explicit data seamlessly is well-positioned for such challenges.

I also see that it is necessary to better integrate FEA and CFD, whether this is used to couple wind loads to vehicle structure, take thermal fluid and map it to the structural surface, or to help the engineers predict the fuel sloshing noise under braking and acceleration in automobiles. Once again, the software that has this capability, or even unified software that offers both FEA and CFD, is in a better marketing position than those who do not have such capabilities built-in.

Secondly is the advent of isogeometric analysis. A technique using NURBS and T-Splines as a basis for construction of element shape functions is currently being championed by professor Tom Hughes at the University of Texas at Austin. It is a new way of thinking about how finite elements will be created or modified in the future, or at least as an element selection option. This is a tremendous way to seek better accuracy to compute stresses, velocities, pressures and

REGISTER FOR THE 2010 NORTH AMERICAN VIRTUAL CONFERENCE

NAFEMS, the only worldwide independent association dedicated to engineering simulation and the people who use it, has announced the keynote speakers and presentation lineup for its upcoming **NAFEMS North American 2010 Virtual Conference: 2020 Vision of Engineering Analysis and Simulation on Nov. 15-16, 2010**. This pioneering online event, containing most of the elements of a traditional conference, but without the costs of travel and accommodations for the attendees, will feature more than 20 live presentations in the auditorium and numerous major simulation technology vendors in the exhibition hall.

A vision of CAE's potential in 2020 has many facets associated with maximizing the value of using engineering analysis and simulation in addressing the ever-increasingly complexity of products and their life cycles. Visit nafems.org/events/nafems/2010/na2010/ for more information and to register.

buckling loads. This technique even has promise in automotive applications to overcome the limitations of modal contents in noise and vibration analysis as well as perhaps more accurate computations gained in fluid dynamics.

Those companies willing and able to continue investing more in realistic simulation technology and appropriate staff will certainly benefit as they develop the next generation of "green" products. The engineers who excel in using advanced CAE simulation technology and apply it to products will be in high demand, not only in the automotive industry but as innovation leaders who will be able to apply their knowledge about the benefits of FEA/CFD applications across many other industries as well. **DE**

Dr. Bijan Khatib-Shahidi, *principal consultant, Engineering Products, Inc., has worked with major manufacturers for 20 years. Contact him via de-editors@deskeng.com.*

INFO → **Abaqus:** simulia.com

→ **Altair Engineering:** altair.com

→ **Beta CAE:** beta-cae.gr

→ **CD-adapco:** cd-adapco.com

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

A Not-So-**Booming** Future

Engineers use computational fluid dynamics to minimize sonic booms.

BY MIKE HUDSPETH, IDSA

There was a time when flying was an adventure full of fun and excitement. You could get on a plane and expect a certain amount of prestige. It wasn't something people did every day. Not anymore.

Nowadays we are herded into coach (because the company won't spring for any better) and left to sit in uncomfortable seats while very nice flight attendants try to make the best of an unpleasant few hours. Flying across the country is one thing, but try going to Europe or Asia. Many people think planes don't need a lounge, but if you've ever spent more than eight hours aboard a plane you can appreciate the need. So how do we fix our flying dilemma? Well, we could build bigger planes. That's what Boeing and Airbus are doing. But that is only going to mean more people on the same multi-hour flight. No, we need a different solution. Go faster then. There are companies working on that too. Remember the Concorde? One problem has been that as we increase the speed of our aircraft, we create sonic booms. If we could find a way to reduce or eliminate those, we might see a new era in air travel.

The "Federales" have all but outlawed commercial supersonic air travel over the continental United States back in the 1970s because people don't like sonic booms disrupting their lives. That means we are limited to flying slower than the speed of sound (761 mph) to get where we're going. So how do we have the best of both worlds—going faster without rattling everyone on the ground?

NASA and several aerospace companies are on a grand quest to create a supersonic plane without a sonic boom. A lofty goal to be sure (pun intended). They have researched such things as wing shape and fuselage design. They found that the longer and pointier you make a plane's fuselage, the lower the turbulence around the plane. They flew several prototypes to test their theories. See figure 1. As technology advances, more and more powerful tools are brought to bear on the problem. Computational fluid dynamics (CFD) has been around for many years. It examines the flow of a liquid around an obstacle. In the '60s it was all done by hand (with slide rules) on paper. Then, they progressed to 2D software. It was faster and more accurate, but the problem was much too complex to solve that way. They needed to go 3D. One company that is involved in the quest is Tecplot, makers of Tecplot 360 software. They have studied the problem from new and more complex angles.

Fig. 1: In 2006, Gulfstream partnered with NASA to test their "Quiet Spike." The thought was to push the shockwave far out in front of the plane.

Liquid Air

Air is basically a fluid. It flows and eddies much like water. To see what that's like, just look at a wind tunnel. This is where aerodynamic studies have traditionally been done. The process is simple. Create a model of your plane, put it up on a stick, blow high-speed air at it, and for some fun, add smoke. (We've all seen those car commercials.) The problem with this approach is that you can really influence or destroy your results if you're not careful. Your model's accuracy has to be spot on. Any rough surfaces on the model will blow your drag numbers. The scale of the model has to be just right. The pylon the model sits on can introduce its own effects (something a real plane wouldn't do). Even the wind tunnel walls can make a big difference.

What CFD does is similar to a wind tunnel, but it's all done inside the computer. It brings some very definite advantages to the table. You don't have to have a pylon. Your model can truly float in the air. You can simulate any speed you want, without limit. If you see something interesting, you can run it again and again, varying the numbers to "zero in" on the results.

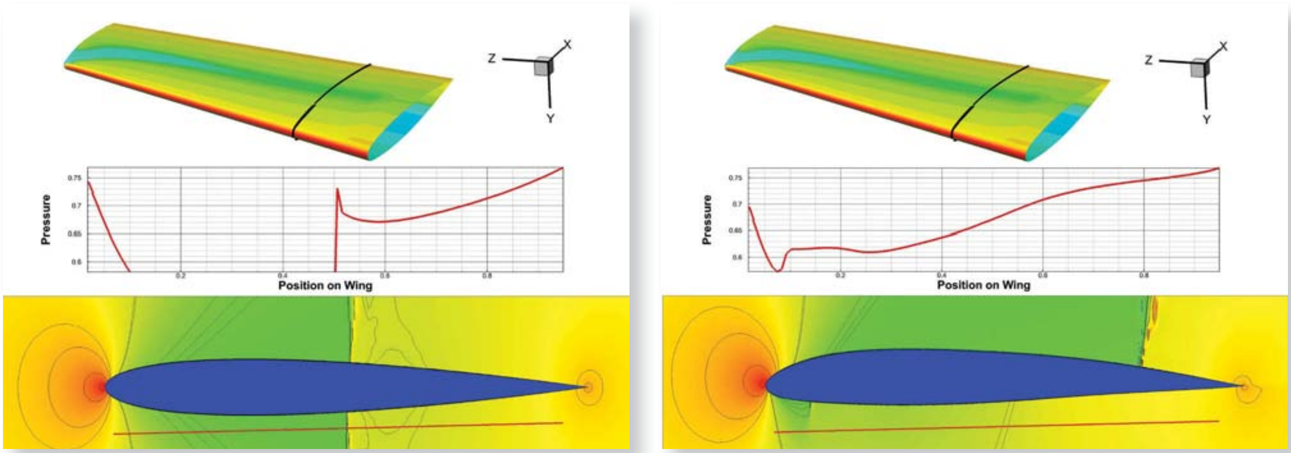

Fig. 2: These images show the changes in off-body pressures around a wing by altering its shape. At the top are pressures on the surface of the wing. In the middle is a graph of off-body pressure below the wing. On the bottom are off-body pressures surrounding the wing. The image on the left shows before the changes. The image on the right shows the wing after its shape has been changed.

"Computational Fluid Dynamics is widely used in any industry that involves fluid flow problems," says Tecplot CEO Mike Peery. "This includes aerospace and defense, shipping, turbines and engines, piping, air conditioning, medicine, environmental (rivers, dams, estuaries, glaciers), weather forecasting, climate modeling, oil and gas, etc. Tecplot has a case studies page on our website that showcases a number of different industry applications. Computational fluid dynamics' roots are in defense and aerospace, having first been used in the Manhattan Project primarily to study blast wave pressures."

If you've used mold flow analysis, you've used CFD.

Until recently, CFD software could give you questionable results. It's like FEA software was in its early days. "Garbage in/garbage out," explains Peery. You had to know what you were doing. The casual user would seldom get reliable results. You needed an expert.

"The accuracy of CFD depends upon many variables: the difficulty of the flow, the computer resources available, and the physics models that are being used. It is extremely easy to be accurate if you're looking at lift, but drag is more difficult."

Nowadays, the software is much easier to use. Tecplot uses its software in conjunction with Sculptor, from Optimal Solutions Software). Sculptor takes a 3D model and deforms it to simulate all kinds of conditions (see my reviews of Sculptor and Back2CAD in the November 2009 *DE*). It's a tool for shape optimization. CFD can save you a lot of money in testing.

"CFD replaces lab experiments," says Peery. "It can replace or minimize the use of wind tunnels, water tunnels, lab set-ups, and models. Today, CFD is on equal footing with experiments, but you typically will need some experimental data (just less of it) and you generally want to compare or 'spot check' your CFD data against your experimental data to verify that it's right.

"What CFD does do better than physical experiments," continues Peery, "is that it gives you a detailed look at what's going on in the flow. It is much, much easier to look at velocity, temperature and pressure with CFD. It is also easier to identify separations and vortices with CFD than with physical experiments." See figure 2.

BOTHERSOME BOOMS

What exactly is a sonic boom and why does it happen? Put simply, when a vehicle (even a car or bicycle) moves, it must displace the air in front of it to the space around and behind it. As a plane flies, the air spreads out around it and compresses in front of it. As the plane goes faster, it exceeds the speed that the air can easily get out of the way. The air stacks up on itself and forms waves around the vehicle. The air forms a cone behind the plane. These waves are really just rapid pressure changes. When a plane flies over your house, its shockwave will follow and you will experience it as a very sudden rise then fall off of air pressure, followed by another rise as the displaced air tries to fill in the low-pressure zone. These pressure changes can trip car alarms, break windows, annoy neighbors, etc. Why are they called sonic booms? Because, as we all know, sound is merely vibration. If the air vibrates, a sound is heard. As the pressure waves change from average to high to low and back to high, all that vibration makes a deep and powerful sound.

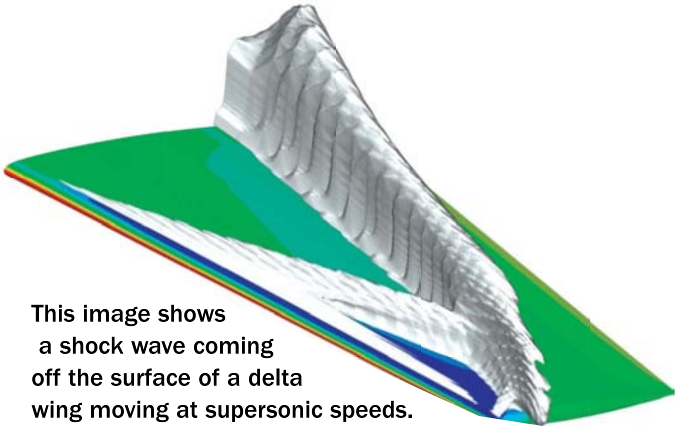

This image shows a shock wave coming off the surface of a delta wing moving at supersonic speeds.

Aerospace companies aren't throwing away their wind tunnels just yet though. There's always a danger in relying on computer simulation if you don't understand the problem or the data being used, Peery says. Computer simulation is one component of product design. Products and designs still require extensive testing before being released to the public.

The Boomless Boon

So what would be the impact of "boomless" supersonic jets?

"The ability to reduce sonic booms can pave the way for new business-class supersonic jets and help lift the FAA ban on overland supersonic travel for civil aircraft," Peery says. "It would also help minimize disruptions on the ground caused by military jets traveling at supersonic speed."

Think 2½ hour flights from New York to Los Angeles or 6 hours to Tokyo. Challenges include the need for more computer resources, speed, and bigger data sets. These are challenges that can be overcome. There are two conditions we'll need to see for supersonic cruising: aerodynamic control when transitioning to and from supersonic speeds, and regulatory cooperation. Ultimately then, it'll be up to the researchers and engineers. All they have to do is find the optimum geometry to reshape not only the shockwaves, but the flying world itself. **DE**

Mike Hudspeth, *IDSA, is an industrial designer, illustrator, and author who has been using a wide range of CAD and design products for more than 20 years. He is* DE*'s expert in ID, design, rapid prototyping, and surfacing and solid modeling. Send him an e-mail about this article to DE-Editors@deskeng.com.*

INFO ➜ **Tecplot:** tecplot.com

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

FOCUS ON Rapid Tech®

NOVEMBER 2010 /// SPECIAL SECTION

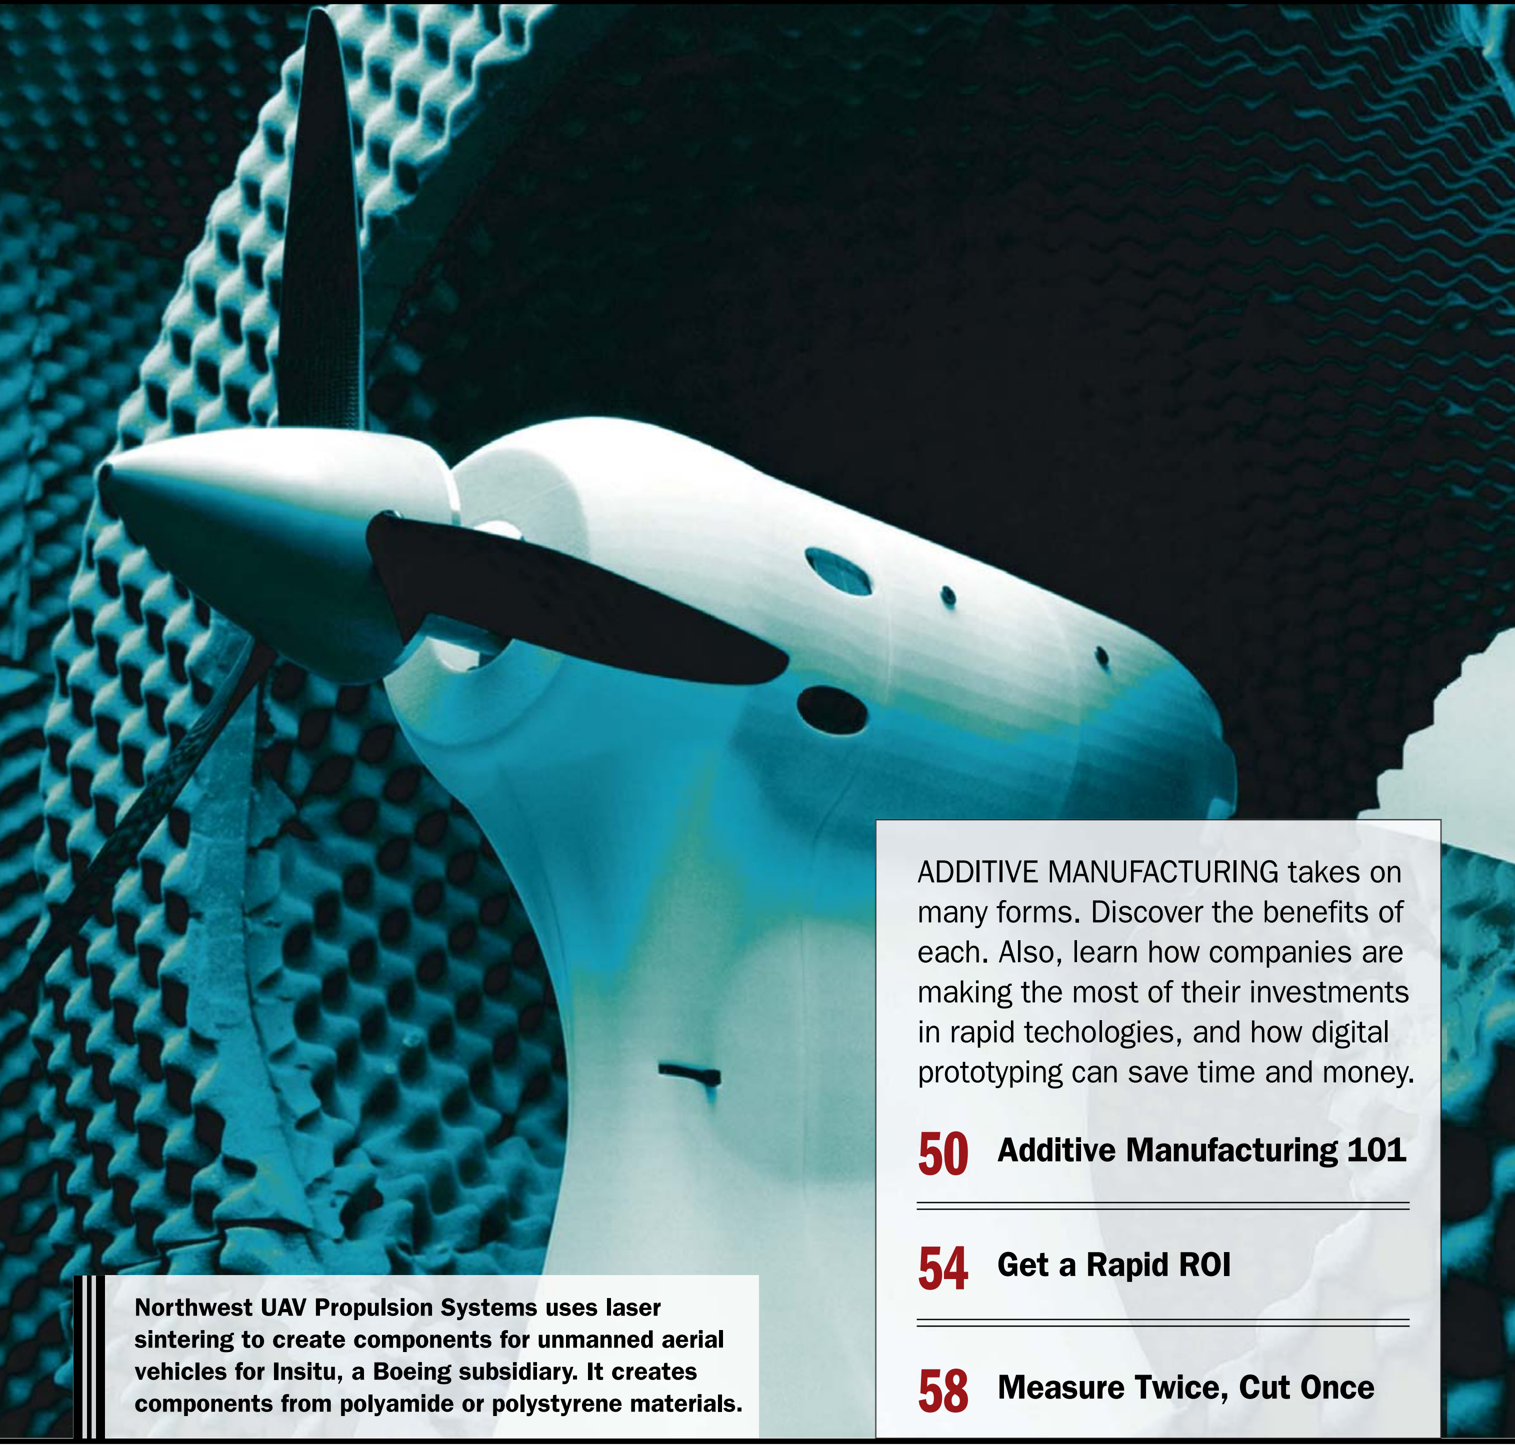

Northwest UAV Propulsion Systems uses laser sintering to create components for unmanned aerial vehicles for Insitu, a Boeing subsidiary. It creates components from polyamide or polystyrene materials.

ADDITIVE MANUFACTURING takes on many forms. Discover the benefits of each. Also, learn how companies are making the most of their investments in rapid techologies, and how digital prototyping can save time and money.

Additive Manufacturing 101

Defining your options for rapid manufacturing.

By Pamela J. Waterman

Have you heard bits and pieces about additive manufacturing (AM) but don't quite know what differentiates the technologies? This month *DE* goes back to basics and clarifies the alphabet soup of AM processes and systems.

Axel assembly of FDM parts made on a Fortus system. Individual part materials include ABS, PC, PC/ABS blend, PPSF and ABSi. Length is approximately 3 ft. *Image courtesy Stratasys.*

Variety Among AM Technologies

While machining is really "subtractive manufacturing" that cuts away an oversized solid material chunk until you get the desired shape, additive manufacturing starts with raw material, generally in sheet, liquid, filament or powder form. Guided by a 3D CAD model for a pattern, AM equipment builds parts up layer by layer to create a physical solid. The benefits of the process are many, and one or more of them apply whether you need a single prototype, a number of very different final parts, a thousand almost-identical parts, or a custom one-off piece.

You could say that casting and injection-molding represent a third type of manufacturing process, but the traditional approach to creating these intermediate-step forms still depends on machining.

Additive manufacturing is an umbrella term for the growing market of international systems that produce a mind-boggling variety of parts and assemblies. In just over 20 years, the variety of attractive options has expanded from the original technology—a stereolithography apparatus (SLA) approach—to dozens of astoundingly creative and different approaches. For an in-depth look at systems and the business market, check out Wohlers Report 2010 by Terry Wohlers (see sidebar).

The ABCs of AM Systems

When looking into various AM systems (hardware, control software and build materials), you'll find that the basic differences lie in speed, accuracy and material types, which affect both initial and long-term usage costs. Some systems are more appropriate for high-volume manufacturing, others for office-environment production, and still others for desktop prototyping, yet even there the lines blur. Be sure to ask about material reuse, as-built surface quality, any optional (or required) finishing steps, and properties such as thermal stability or chemical resistance.

Add to this mix the possibility of building or using very low-cost equipment (see Sidebar "Low-Cost 3D Printers Guide" on page 54) and you'll see it pays to become an informed consumer.

THE ANNUAL AM ROUNDUP

If you're serious about weighing the attributes of any AM system, particularly toward a purchase, *Wohlers Report 2010: Additive Manufacturing State of the Industry Annual Worldwide Progress Report*, is a resource guide. Now in its fifteenth year, this 250-page review by analyst Terry Wohlers covers industry growth estimates, new developments and business opportunities, emerging applications, research and development and a view to the future. Other documents and free information are available at wohlersassociates.com.

A full-size, full-color joystick model created with the 3D printing process on a ZPrinter 650 from ZCorp. *Image courtesy Z Corporation.*

3D printing is an umbrella term for AM because it presents an easily understood vision of transitioning from paper copies to solid models. A decade ago this often referred to a process licensed by MIT to five different end users, two of which still exist with a strong market presence: Z Corporation and ExOne. The 3D printing process uses a standard 2D inkjet print head to spray binder liquids onto powdered base materials, tracing and hardening the layers of a 3D CAD model.

Z Corp. offers two composite materials: plaster-based and cellulose-based. The cellulose-based material is suited to infiltration with an elastomer to create rubber-like parts. ExOne creates both direct-metal parts (in stainless steel infiltrated with bronze or gold) and sand cores and molds for metal casting.

Aerosol Jet technology, from Optomec, differs from inkjet processes. It precisely delivers fluid and nanomaterial formulations, producing electronic, structural and biological patterns on almost any substrate. The approach uses aerodynamic focusing to deposit atomized materials, which can be optionally post-treated with a focused laser or other sintering methods.

Digital Light Processing (DLP), a Texas Instruments process, is the technology behind Envisiontec's line of AM equipment (Z Corp. also markets a DLP system based on an OEM agreement with Envisiontec). This approach uses a series of micro-mirrors to reflect light (UV) through projected, 3D-CAD-created masks to solidify layers of various photopolymers.

Direct Metal Deposition (DMD) describes the approach that the POM Group takes to create and also repair metal parts made in tool steel, titanium, Inconel and other alloys. One system offers an evacuated chamber for special alloy processing, fabricating fully dense metal "from the ground up" using powdered metal and a focused laser. Another system is configured as a portable robotic arm that brings the machine to the part for shopfloor repair of massive tooling. Dissimilar metals can be combined as needed.

Direct Metal Laser Sintering (DMLS) systems from EOS produce parts made from metals, plastic polymers or sand, depending on the model. The technology fuses powdered material into fully dense parts formed by melting it with a laser beam. Special process-control software and other quality assurance functions allow these systems to even process so-called "uncastable" metals, particularly those with high-temperature applications.

Drop-on-Demand (DoD) describes the Solidscape AM process typically used to create highly detailed casting models for the jewelry and dental industries. It works by spraying heated wax material (which starts in granular form) through print heads, as with 2D inkjet printers. Two types of wax serve the build and support functions.

Electron Beam Melting (EBM) represents another approach to creating fully dense metal parts by melting layers of powder. This process, developed by Arcam, employs a MultiBeam technology that uses fast deflection electronics (rather than optics and moving mechanical parts) to melt the powder simultaneously at multiple locations. A sequence of scans first heats the region to a certain temperature specific for different alloys, then melts the part contours, and lastly melts the bulk of the material.

Fused Deposition Modeling (FDM) is the term for Stratasys' technology that heats and extrudes a fine-diameter filament of thermoplastic material and deposits it bit by bit along with any necessary support structures. The latter material is soluble and can be dissolved in a water-based solution. Available build materials are a number of ABS plastics, polycarbonates and polyphenylsulfone.

Film transfer imaging (FTI) produces hard plastic parts on 3D Systems' VFlash Personal 3D Printer, curing one full layer at a time. This desktop modeler operates with a single-use recyclable cartridge that contains both the raw build material and such parts as pumps, measured dispensers, re-coaters and transporters.

Laser consolidation (LC) refers to the net-shape direct-metal process developed by Canada's National Research Council's (NRC) Integrated Manufacturing Technologies Institute (now marketed through Accufusion). LC bonds high-strength alloy onto edges and other features of inexpensive metal parts and tools, reducing the amount and cost of alloy use for part production and repair.

Laser engineered net shaping (LENS) systems from Optomec use a high-power laser to fuse powdered metals into fully dense 3D structures. LENS systems can process a wide variety of metals including titanium, nickel-base superalloys, stainless steels and tool steels.

Laminated object manufacturing (LOM) is a general term for building a part from flat sheets of material stacked and "glued" together. The sheets are cut to shape by a laser or a blade after each layer is bonded. Previously used by now-defunct Helisys, LOM (or just "layered manufacturing") describes the systems from two current companies.

Solido's lamination system bonds colored or clear PVC plastic sheets (cut from a roll) with glue where desired and "anti-glue" where not wanted, an approach that simplifies manually peeling away the unbound sections. An extremely low-cost LOM system is the paper-based Matrix from Mcor Technolo-

LOW-COST 3D PRINTERS GUIDE

So you've researched the dozens of options for AM systems, have fairly simple needs and would like to spend less than $10,000. Where do you go for help? Check out the new report from Castle Island, "Everything you should know about open-source 3D printers — but didn't know you should ask." Researched and written by long-time industry pundit Ed Grenda, this document (29-page PDF, $9.95) covers both commercial 3D printers, such as from Stratasys and Solido, as well as the do-it-yourself open-source kits and assembled printers now on the market, from RapRep to Fab@Home to lesser-known systems. A fascinating read for anyone in the AM field, this report also offers web resources, supplies information and a look at the legal/patent angle of the business.

Fully detailed turbine blade created from sheets of PVC plastic, laminated on an SD300 Pro LOM system from Solido3D.com. *Image courtesy Solido3D.com.*

gies. This unit builds concept parts using paper sheets from a standard copy machine. Even scrap paper will work.

MultiJet Modeling (MJM) is the basis for the ProJet line of equipment from 3D Systems. Similar to inkjet 2D printers in concept, the MJM compact manufacturing systems create high-definition parts or wax castings by "jetting" various plastics, waxes and support materials.

Micro Light Build (MLB) systems have been introduced by Huntsman Advanced Materials. Using MicroLightSwitch technology based on a micro-electronic mechanical systems (MEMS) shuttering system, this AM approach enables a larger surface area of radiation-curable resin to be selectively exposed at one time, with conventional UV light instead of a laser beam.

Metal Printing Process (MPP) is the name for a manufacturing method in development at the UK/Norwegian research institute SINTEF. It involves an electrostatic (xerographic) photoreceptor that attracts powdered build material and support material to create one layer of a CAD-driven pattern. The material is transferred from a roller to a punch/consolidation system where each layer is sintered with pressure and heat.

PolyJet manufacturing from Objet Geometries involves jetting or squirting a liquid, acrylic-based photopolymer thorough fine nozzles in the CAD-driven shape, then curing each layer with UV light. No post-curing is necessary, and the support structure (a gel-like material) is removed easily by hand or with water. This is the only system that builds parts incorporating multiple dissimilar (transparent, hard, rubbery, etc.) materials in a single build, through a dual-jet multi-print-head configuration.

Stereolithography Apparatus (SLA) refers to the original, high resolution AM process developed by 3D Systems. It produces fine-finish parts from liquid photopolymer by solidifying it with a high-definition UV laser. More than a dozen clear or opaque material choices create parts with properties of ABS and polycarbonate plastics, plastic-ceramic composites and/or high thermal resistance.

Selective Laser Melting (SLM) describes the process taken by MTT Technologies that uses a laser to melt and fuse metal parts. Materials include tool steel, aluminum, Inconel, gold alloy, titanium and cobalt chrome.

Selective Laser Sintering (SLS) uses a laser to heat and fuse powdered material into solids. The process, from 3D Systems, can be applied to various plastics, glass-filled polymers, metals, alloys, ceramics and green sand, to produce hard, flexible, fully dense or even flame-retardant parts.

Selective Mask Sintering (SMS) is another process that sinters thermoplastic powdered material. However, this approach uses an IR laser to solidify powder through a series of physical masks. The masks' cut-out shape is generated on the fly from a roll of stock and allows exposure of a complete layer in a single flash. Sintermask Technologies is the vendor.

Ultrasonic Consolidation (UC) was developed by Solidica as a way to embed sensors, diagnostics and other electronics in rugged metal packaging. It uses sound waves to merge layers of metal rolled off narrow sheet stock. The combination of bonding plus machining for each layer creates full-density parts with metallurgical bonds and works with a variety of metals. DE

Contributing Editor **Pamela Waterman**, DE's *simulation expert, is an electrical engineer and freelance technical writer based in Arizona. You can send her e-mail to DE-Editors@deskeng.com.*

INFO → **3D Systems:** 3dsystems.com

→ **Arcam:** arcam.com

→ **EOS:** eos.info

→ **Envisiontec:** envisiontec.com

→ **ExOne:** prometal.com

→ **Huntsman Advanced Materials:** huntsmanadvancedmaterials.com

→ **Mcor Technologies:** mcortechnologies.com

→ **MTT Technologies:** mtt-group.com

→ **Objet Geometries:** objet.com

→ **Optomec:** optomec.com

→ **POM Group:** pomgroup.com

→ **Sintermask Technologies:** sintermask.com

→ **Solidica:** solidica.com

→ **Solido3D.com:** solido3d.com

→ **Solidscape:** solid-scape.com

→ **Stratasys:** stratasys.com

→ **Z Corporation:** zcorp.com

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

Get a Rapid ROI

What gets the job done most efficiently and cost effectively in rapid manufacturing may be more important than sticker price.

By Susan Smith

Rapid prototyping and manufacturing costs are generally calculated in terms of how much the equipment costs, minus how much time and money is saved with the equipment, and what option is chosen—whether it be to purchase 3D printers, systems platforms or to use service providers. Other benefits include being able to make several iterations of a prototype until you arrive at the right design in-house, rather than outsourcing and possibly not being part of the whole prototyping process.

The costs of rapid prototyping and manufacturing are commensurate with what value your organization can derive from the investment. Is a low cost 3D printer what you need, or do you need to go with a higher end selective laser sintering (SLS) or 3D system to meet your manufacturing needs? Would your organization prefer to try a service provider first to see what processes work best, then invest in your own equipment?

Low-Cost Options

Organizations need to determine if they want to do prototyping or whether they want to create end-use parts or products. Many low-cost 3D printers are capable of producing good, accurate prototypes. Multi-material machines from Objet Geometries can enable customers to create prototypes using numerous materials that have the look and feel of the finished product. Some customers find they want to do manufacturing of end use parts, and can justify the cost of higher end machines. Prototyping requirements for some organizations are very specific, such as those in the medical or aerospace industries.

The number one way to keep costs down in manufacturing is to perfect your design with prototyping before going to tooling, to reduce tooling rework, according to Joe Hiemenz, technical communications and public relations manager at Stratasys Inc. In terms of manufacturing, you want to be able to eliminate tooling altogether.

The low cost of some 3D printers have made them attractive to other industries such as architecture and art. Some models come in the $14,000 to $24,000 range, such as two new machines recently introduced by Z Corp. at an entry-level price point: the ZPrinter 150 (monochrome, $14,900) and ZPrinter 250 (multicolor, $24,900). These machines have a small physical footprint that is suited for the classroom, department, or small business, but are not aimed at manufacturing.

The Objet Eden reduced approximately 20% of ADMTronics' design cost in getting its part readied for a clinical trial.

In the low-medium price range is the new ZBuilder Ultra ($34,000) ranging upward through the Objet ($100,000 to $240,000) to the high end Stratasys fused deposition modeling (FDM) ($250,000 to $350,000) and EOS' $250,000 to $850,000 laser sintering (LS) offering.

Measuring ROI

Measuring return on investment can be a no-mans-land for some customers, but can be quantified by greater accuracy, time savings and production costs. One way of looking at

A manifold created with a Dimension 3D printer.A contractor for the Missouri National Guard reported they saved enough money to pay for their Dimension 3D printer, which is priced around $30,000.

ROI is that a penny over the purchase price is technically how ROI is calculated. Once you make your money back after purchase of the machine, you make one cent of return. But this is not counting the cost of materials. Another perspective is that once you have made a purchase, from that time on you stand to save money as the use of the machine is helping to save from that day forward. Costs are probably easier to calculate if the rapid prototyping is outsourced.

Hiemenz says the quickest way to measure ROI after purchase of a 3D production machine is to manufacture assembly jigs and fixtures on it, because it cuts out machining and tooling which saves weeks or months of work.

"You're going to save 50 to 99% in production costs, typically more like 70 to 90%," he says. "In the manufacturing world, if you can save 5 or 10% you're pretty happy."

Customers can save the same amount in production costs by making jigs and fixtures with their Dimension 3D printers. Di Mino says they can use Objet Eden to create tools such as jigs and fixtures for their own use in the manufacture of the devices.

"I've seen it (ROI) as little as one project," says Hiemenz. "On a Dimension 3D machine, a contractor for the Missouri National Guard made 40 camera mounts for armored tanks, a manufactured end-use product." They reported they saved enough money to pay for their Dimension, which is priced around $30,000.

At Sheppard Air Force Base in Wichita Falls, TX, the Trainer Development Flight (TDF) is a facility that designs, develops and manufactures trainers and training aids for the Air Force and all branches of the Department of the Defense (DoD). These products are either original products or replicas of existing ones. One of their designs is an exact replica of a UAV for training repair technicians. The TDF uses four Stratasys Fortus machines. For just the production of the UAV's large antenna, for example, the company reported that it took one-tenth the time that it would have taken to accomplish with conventional methods, with an ROI of more than $12,000.

Mitchell Weatherly, chief of the TDF said "for our first FDM machine purchase, we projected ROI in four years, but it took only 18 months. For our second FDM machine purchase, we saw ROI in only nine months."

The cost of the four machines that TDF purchased was approximately $350,000 each.

Stratasys Fortus 3D Production systems range from $90,000 to $450,000. Their Dimension 3D Printers range from $15,000 to $35,000.

Perfecting your design with prototyping before going to tooling reduces tooling rework.

Intangible Benefits

Andre DiMino, executive vice president and CTO of AD-MTronics (formerly Ivivi Technologies) says not all value in rapid prototyping can be quantified. For example, on one project they calculated that the Objet Eden reduced approximately 20% of their design cost in getting the part readied for a clinical trial.

"As far as costing, that was 20%, but there was more like a 40% reduction in time, which was even more valuable to us," says Di Mino.

DiMino says AD-MTronics develops, designs and manufactures medical electronic equipment with some ergonomic functions using Objet's technologies. They work mostly with non-invasive electrotherapies such as nerve stimulators and ultrasound devices.

With previous processes, DiMino says the company used stereolithography combined with getting molds made and sculpting in clay. The molds had to be sent out to a service bureau.

"Being able to draw the units, producing them in our Objet Eden 350 machine and having that piece be right to install the electronics has cut months off the process," says DiMino. "It also allows us to try many different things, print them on the printer, install the electronics, and physically see it operating within days, which we wouldn't be able to do with previous processes because it would take so long to do different variations."

Medical devices must go through clinical trials, and DiMino pointed out that you want to put the product into the clinical trial that is as close to the finished production piece as possible.

"We will print out 40 to 100 units for clinical trials, which keeps us from having to make an actual finished mold for the clinical trial," DiMino says. "Once the clinical trial is done, it gives us some feedback as to whether changes are needed before we go into the actual expense of molds. We're actually able to change the ergonomics before we go into making a final mold."

Object's high-end Eden 3-Dimensional Printing Systems offer printing with the accuracy of 16µ layers, and all Eden systems produce models with very fine details and smooth surfaces. The Eden 350 and 260V are similar in nature with only a difference in bed size. The 260V is the same as their 500V with only a different bed size, and retains all the features and functions but at a lesser price.

The Eden260V costs approximately $100,000. A Connex costs approximately $240,000.

The ability to make parts with different facets that could not be machine cut or injection molded is a real plus that is hard to calculate in terms of time savings. EOS points out that Northwest UAV Propulsion Systems (NWUAV), an aerospace provider, is now able to make parts with multiple angles that couldn't be machine cut.

Unmanned vehicle craft are very weight sensitive and require a lot of design iterations, and new functionality for which EOS' laser sintering (LS) process works well. Lightweight nylon parts are required for ducting and shrouding as well as fuel tanks and a part can be designed with a unique shape to fit a complex space with many interfaces.

According to Alex Dick, of Northwest Rapid Manufacturing, which is part of the Northwest UAV family of companies, from a production standpoint, "we're most often using the latest sintering process to make shrouds that direct cool air over the engine ribs, and the top of the engine to make sure it's staying cool enough. That's a three- or four- part assembly with different LS parts.

"The laser sintering process also gives us a large flexibility in the volumes that are produced. We can produce a single part, or we can produce several hundred parts quickly and cost effectively dependent on part size," says Dick. "You can fit a lot of features and functionality, holding reasonable tolerances, etc. with parts with x/y/z dimensions of 4 to 6 in. Too small or too large and the tolerances fall off but there are good applications on both ends of the part-size spectrum (0.030 in. minimum to 28 in. maximum in a single piece)."

NWRapid Manufacturing does all the component production for laser-sintered parts. NWUAV is doing all the assembly, and handling the whole engine subsystem, which includes the controller cards and propeller and everything in between.

"As compared to other additive technologies, with laser sintering there are time and labor savings because there is no support creation or removal process," says Dick. "Parts can be nested for increased productivity. Compared to conventional technology, high-complexity, low-volume parts can be created without costly molds or assembly."

Dick said that for the quantity and complexity and frequency of design changes, "I don't see another solution."

NWUAV initially started additive manufacturing with the EOSINT P390 system (currently priced at around $350,000) and it quickly discovered a need for increased capacity leading them to the purchase of a second system, EOS's larger frame plastic sintering system the EOSINT P730 (currently priced at around $850,000).

With the larger build envelope of the Stratasys Fortus machine, for example, it is possible to make bigger parts in one go, and the throughput is much faster: 200 to 1,000% faster than that of Dimension 3D printers, according to the company. The 3D production machine also has higher accuracy and more options in terms of software and user control over the parameters, plus higher resolution. DE

RAPID MANUFACTURING BUREAUS AT YOUR SERVICE

One way customers can get rapid manufactured parts built without owning their equipment is to first use a service bureau, which maintains the latest in machines and materials, and can advise which ones will work best for their needs.

Service providers can address customers who wouldn't otherwise be able to afford to buy their higher end systems. Joe Hiemenz, public relations manager at Stratasys, says some of these customers later purchase 3D printers or high end machines for their own organizations.

Cathy Lewis, vice president of global marketing for 3D Systems, said she sees "significant potential" with the company's 3Dproparts service bureau in the defense and aerospace market, where these companies and their subcontractors don't necessarily want to own and operate every type of additive manufacturing system available, particularly with their myriad application requirements. The preference seems to be growing for access to a qualified service provider with a broad array of rapid manufacturing solutions and expertise in AS9100 and ISO:9001 processes and certifications.

But many come to rapid technologies to have more control over their design process, and to accomplish in prototyping or manufacturing, what cannot be done with machining, tooling or injection molding. Maintaining machines and materials in-house allows them to produce volumes of products in their own timeframe, without having to wait for them to be completed by a service provider.

Potential customers can research vendors who offer different machines, or visit customers who are using them successfully to see what might work for them. A lower cost machine might be a good starting point for a new customer.

Susan Smith *has been working as an editor and writer in the technology industry for more than 16 years. She is a contributing editor for* Desktop Engineering *magazine.She can be reached via de-editors@deskeng.com.*

INFO → **Company Name: Dimension Printing** uprint.dimensionprinting.com

→ **Stratasys** stratasys.com

→ **3D Systems** 3dsystems.com

→ **Objet Geometries, Ltd.** objet.com

→ **EOS** eos.info

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

Commentary By Scott Crump

Measure Twice, Cut Once

In these uncertain economic times, I remind myself to "measure twice; cut once," a lesson my father taught me as a boy.

In the R&D arena, that translates to "prototype often" before cutting a production tool. Prototyping offers the confidence that there will be only one tool cut and no rework. For very little financial gain, companies take a big risk by ignoring the wisdom of prototyping early and often.

In a tough economy, an early casualty for many companies is the number of prototypes built. Faced with budget cuts, prototyping is incorrectly viewed as an optional expense. I have seen companies take this route, and I have seen them suffer the consequences. Over the years, we have suffered the same fate when we assumed the design was perfected, and another prototype wasn't needed.

No one is immune to the struggle to cut costs and increase productivity. And even though our company manufactures prototyping systems, we too feel the temptation to cut back on prototyping to save a few bucks and a little time.

> **"Prototyping is incorrectly viewed as an optional expense."**

Gambling on Prototypes

In part, this happens because there is no direct correlation, no cause-and-effect relationship between the prototype and the success of the product. Building prototypes does not guarantee problem-free designs, market demand or profitable sales. Eliminating a prototype does not mean that there will be rework and manufacturing problems. The only sure thing is that the prototype will cost you some time and money.

So, the prototyping decision is a gamble. As with an insurance policy, you are betting on the outcome. If a prototype is eliminated and you don't run into any problems, the gamble pays off. But if a design flaw is discovered when tooling is being cut, a $1,000 prototype savings can turn into a $10,000 tooling mistake that derails production schedules. I guess the question comes down to "Do you feel lucky?"

Before answering that question, reflect on your past design projects. How many prototypes have you made that didn't reveal the need for any revisions: 10%, 5%, 1%? For most, it is rare that a prototype doesn't lead to at least one discovery. The truth is that the odds of having a flawless design are against you. In this light, prototyping isn't an option, it's a necessity.

Work Smarter

While a recession demands change in the way we do business, we have to be smart about the alterations we make. If budgets must be cut, it must be done wisely. This means being frugal, not cheap. It means cutting back, not cutting corners. It means being conservative, not timid. And it means being progressive, not stagnant.

To be frugal, conservative and progressive—all at the same time—we have to work smarter. We have to use the tools that are available to make the most of our time and money. From my experience, prototyping is one of those tools.

In my mind, prototypes are absolutely essential. Yet, being a realist, I know that the economic climate will have an impact on when and how often they are used. I know that some sacrifices have to be made. But I am convinced that the only reason not to prototype is when there is no R&D activity, which is another gamble with very poor odds. **DE**

Scott Crump *is CEO of Stratasys, a maker of additive manufacturing machines for prototyping and producing plastic parts. Send e-mail about this commentary to de-editors@deskeng.com.*

CONTRIBUTE YOUR COMMENTARY

This commentary page is your chance to share your opinions and insights with *DE*'s readers. Design engineers, vendors, engineering students, consultants and teachers are all welcome to submit commentaries to *Desktop Engineering* for consideration.

Just send an 800-word commentary to de-editors@deskeng.com. It should focus on ideas that can help design engineers work more efficiently, innovate, or take advantage of a specific technology. If it is accepted, we'll notify you of its publication date.

Converting Airbus A350 Models

GLOBAL VISION SYSTEMS is a vendor of solutions for industrial supervision. The start-up was founded in 2007 by two engineers, David Croenne and Baptiste Gendron Canion, who worked together as consultants at the German Space Operations Center (GSOC). The company's flagship product is called EasyMonitoring. It's an ergonomic platform that facilitates the supervision and analysis of complex systems in their working environment.

EasySensorManager is the second part of the Global Vision Systems offering. Presented as "PLM for instrumentation," the application offers a simple graphic human-machine interface in which to implement test equipment in a digital mockup of a supervised environment.

INFO ➜ deskeng.com/articles/aaaysp.htm

Xcelaero Improves Radial Fans with Concepts NREC

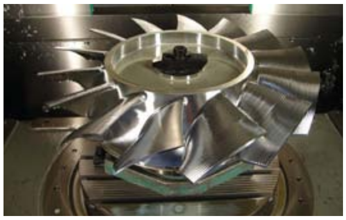

XCELAERO CORPORATION of San Luis Obispo, CA, specializes in custom-vane axial and radial fans for the commercial and defense industries. It uses techniques and technologies adapted from the aviation and gas turbine industries to achieve fan efficiencies.

Xcelaero's axial fan designs use complex shrouded vane sets, which present manufacturing challenges due to a lack of tool access and minimal spacing between vanes. Xcelaero is successful, in part, because it purchased 5-axis computer-aided machining software from Concepts NREC (CN).

"To prototype our fans we purchased a 5-axis vertical machining center and coupled it with CN's 5-axis CAM software family including the MAX-PAC and MAX-SI packages specifically for shrouded vane sets," says Bill Denzel, engineering manager at Xcelaero. "Additionally, we chose to use the CN CAM family because it merged with our impeller design software."

INFO ➜ deskeng.com/articles/aaayss.htm

FDM Cuts Jet Engine Prototype Time

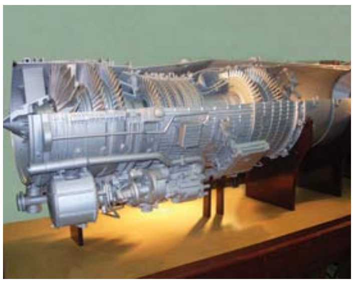

THE GAS TURBINE Research Establishment (GTRE) of Bangalore, India, is a government laboratory whose primary function is research and development of marine and aeronautic versions of gas turbines. Development of GTRE's flagship product, the Kaveri jet engine, was commissioned for the HAL Tejas aircraft, which has an all-terrain capability that spans from hot deserts to the world's highest mountain range.

One of the greatest challenges in designing the Kaveri was positioning its many piping runs and line replaceable units (LRUs) on the outside of the aircraft. Many of the LRUs are connected to the interior of the engine with pipes that carry hydraulic fluid, fuel, and lubricants. It was a major challenge to design each piping run to minimize length to reduce weight and cost while avoiding interference.

INFO ➜ deskeng.com/articles/aaayst.htm

Valiant Takes Flight with Autodesk Software

IN 50 YEARS of business, Valiant Machine & Tool Inc. has forged a reputation for out-innovating its competition. While the company traditionally has developed manufacturing products for automotive, construction, and forestry industries, Valiant began applying its creative engineering expertise to aerospace industry projects with impressive success.

Due to this success, Boeing recognized Valiant as a "Supplier of the Year" for 2007 in the "Non-Production/Shared Services" category.

At the core of Valiant's innovative designs are Autodesk Inventor Professional software and Autodesk Algor Simulation software, which enable Valiant to more quickly assess whether a new idea is a dead end or a promising beginning.

INFO ➜ deskeng.com/articles/aaaysw.htm

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

SPOTLIGHT

Directing your search to the companies that have what you need.

Advertising Index

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

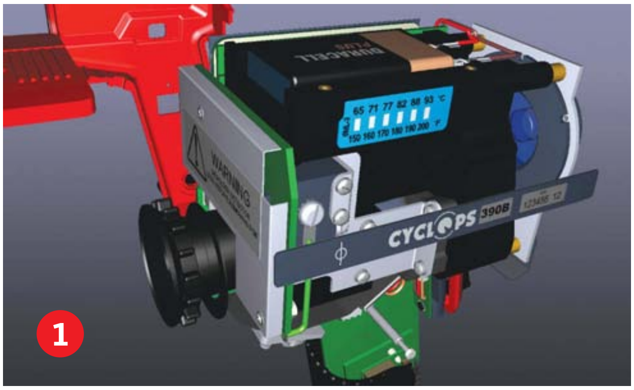

1 Autodesk, Inc. has introduced the latest release of Autodesk Inventor Publisher technical communications software. Autodesk Inventor Publisher 2011 software helps manufacturers explain and differentiate their products with 2D and 3D product documentation. The software incorporates many new enhancements, including the ability to publish 3D interactive instructions directly to iPhone, iPad and iPod touch mobile devices. autodesk.com

2 Cray Inc. has announced that the new Microsoft Windows HPC Server 2008 R2 operating system is now available on the Cray CX1 line of deskside supercomputers, including the Cray CX1 and the Cray CX1-iWS systems. The Cray CX1 line of deskside supercomputers is designed for individuals and departmental workgroups who want to harness the compute power of an HPC cluster. cray.com

3 Mersive Technologies, Inc. has announced the M-Series, an 11-megapixel display that delivers a 157 in. diagonal image while using a footprint that is 24 in. deep. The company says the image is almost six times the resolution of HD and enables collaborative visualization for applications such as engineering design and scientific visualization. At less than 2 ft. deep, the M-Series eliminates the need for room reconfiguration and specialized HVAC and building infrastructure. It can be maintained by IT or business staff. It uses Mersive's SOL Harmony display management software engine that automatically aligns, color balances, and maintains image quality across multiple projectors throughout the life of the display. The company says the M-Series display system sets up in less than half a day and requires minimal ongoing maintenance. It works many video sources, including PCs, or video wall controllers. Image quality is obtained by combining multiple HD resolution projection units into a single, seamless image of over 11 megapixels. mersive.com

4 FARO Technologies, Inc. has introduced the FARO Laser Scanner Focus3D. Focus3D is designed for detailed measurement and documentation with touch screen control. It is four times lighter and five times smaller than its predecessor. It uses laser technology to produce three-dimensional images. The resulting image is an assembly of millions of 3D measurement points in color, which provides a digital reproduction of existing conditions. The company says the Focus3D has millimeter accuracy and 976,000 measurement points/second. It is entirely self-contained, meaning no additional devices, cables or laptop are needed. Its dimensions are 9.5x8x4 in. and it weighs 11 lbs. faro.com

5 Mentor Graphics Corporation has announced new capabilities in the next generation of its FloTHERM 3D computational fluid dynamics (CFD) software for electronics cooling applications. The technology in the FloTHERM software provides Bottleneck (Bn) and Shortcut (Sc) fields so that engineers can now identify

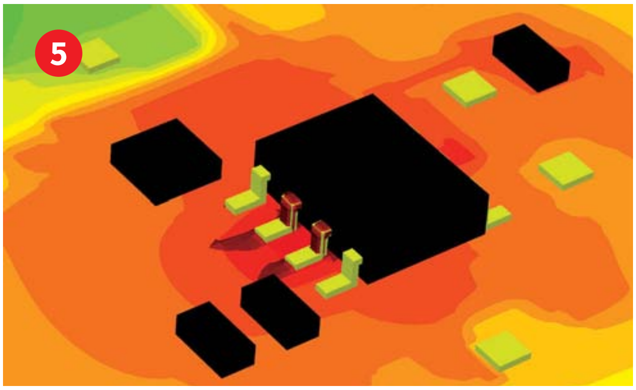

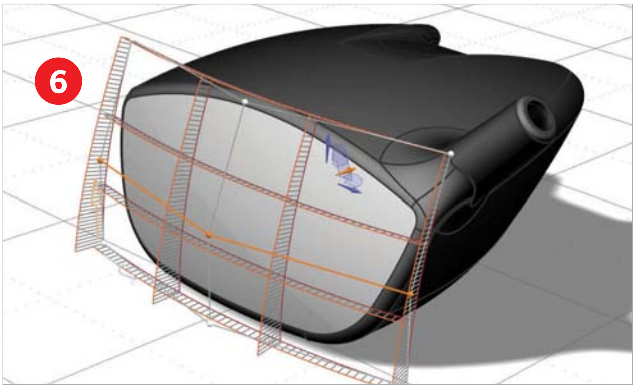

where heat flow congestion occurs in the electronic design and why. It also identifies thermal shortcuts to help quickly and efficiently resolve the design problem. Together, the Bn and the Sc fields are designed to elevate the use of simulation from an observation tool, which identifies heat management problems, to an effective thermal design problem-solving tool that suggests potential solutions to the designer. The end result is faster and more efficient resolution of heat management problems, according to the company. mentor.com

6 NX product design software, from **Siemens PLM Software**, now supports synchronous-enabled freeform design. With these new tools, users can create and modify complex freeform geometry. The shaping toolkit works with any geometry, including imported models. Push and pull shaping techniques mean users don't have to be a surface design expert to create complex freeform models. The freeform modeling tools allow users to begin with solid or surface, analytic or B-rep geometry. Users can insert or model organic forms by moving constraint points, surface poles and handles. Surface selection and analysis tools help users control continuity, blending and inflections. These freeform models are re-usable, from concept through production. plm.automation.siemens.com

MathWorks Delivers GPU Support for MATLAB

MathWorks (mathworks.com) has announced support for NVIDIA graphics processing units (GPUs) in MATLAB applications using Parallel Computing Toolbox or MATLAB Distributed Computing Server. This support enables engineers and scientists to increase the speed of many of their MATLAB computations without performing low-level programming. Now more engineers can take advantage of NVIDIA's CUDA-enabled GPUs, including the latest Tesla 20-series GPUs, all from within MATLAB. Parallel Computing Toolbox users can access the NVIDIA CUDA library without having to learn CUDA programming or significantly modifying their applications.

Omega Releases RDXL8

Omega's (omega.com) new line of thermometer/data loggers can be used for remote logging and as a mobile measuring device. The RDXL8 can perform integrated recording, profiling and archiving functions. This CE-marked product features a 5.7-in. display, TFT LCD or monochrome 5-in. display, eight configurable isolated inputs for thermocouples and DC millivolts, graphic display of each input, and programmable "hi or lo" alarm for eight inputs. Prices start at $1,900.

New Parallel Nsight and CUDA Developer Tools

NVIDIA's (nvidia.com) Parallel Nsight enables creation of GPU-accelerated applications. Version 1.5 includes support for Microsoft Visual Studio 2010, Tesla Compute Cluster (TCC) debugging, CUDA Toolkit version 3.2, and support for NVIDIA's Fermi GPU architecture. The new CUDA Toolkit 3.2 release includes two new math libraries, performance improvements and support for the new 6GB Tesla and Quadro products. DE

Statement of Ownership, Management and Circulation

1. Publication Title: DESKTOP ENGINEERING 2. Publication No. 0014-476. 3. Filing Date 10/11/10. 4. Issue Frequency: MONTHLY. 5. No. of Issues Published Annually: 12. 6. Annual Subscription Price: NONE. 7. Complete Mailing Address of Publication: LEVEL 5 COMMUNICATIONS, Inc., 1283D MAIN ST., PO BOX 1039, DUBLIN, NH 03444. Contact Person: DARLENE SWEENEY, Telephone: 603-563-1631. 8. Complete Mailing Address of of Publisher SAME 9. Full Name/Mailing Address of Publisher: TOM CONLON, address: SAME. Executive Editor: STEPHEN ROBBINS, address: SAME. Managing Editor JAMIE GOOCH, address: SAME. 10. Owners STEPHEN ROBBINS, THOMAS CONLON, address: SAME. 11. Known Bondholders, Mortgages, and Other Security Holders Owning or Holding 1% or More of Total Amount of Bonds, Mortgages, and Other Securities: NONE. 13. Publication Name: DESKTOP ENGINEERING.
14. Issue Date for Circulation Data Below: SEPTEMBER 2010.
15. Extent and Nature of Circulation:
a. Total Number of Copies: Average no. each issue: 54,491. No. copies of single issue: 51,938.
b. Legitimate Paid and/or Requested Circulation stated on PS Form 3541:
(1) Outside-County: Average each issue: 53,318. Sept 2010 issue: 51,078.
(2) In-County: Average each issue: 0. Sept 2010 issue: 0.
(3) Sales Through Dealers, Carriers, Street Vendors, Counter Sales, & Other Dist. Outside USPS: 0.
(4) Requested Copies Distributed by Other Mail Classes Thru USPS: 0.
c. Total Paid and/or Requested Circulation: Average each issue: 53,318. Sept 2010 issue: 51,078.
d. Nonrequested Distribution stated on PS Form 3541:
(1) Outside-County: Average eash issue: 775. Sept 2010 issue: 774.
(2) In-County: Average each issue: 0. Sept 2010 issue: 0.
(3) Other Classes DistributedThrough the USPS by Other Classes of Mail: 0.
(4) Outside the Mail: Average each issue: 398. Sept 2010 issue: 131.
e. Total Nonrequested Distribution: Average: 1,173. Sept 2010 issue: 905.
f. Total Distribution: Average: 54,464. Sept 2010 issue: 51,957.
g. Copies not Distributed: Average: 505. Sept 2010 issue: 520.
h. Total: Average: 54,491. Sept 2010 issue: 51,983.
i. Percent Paid and/or Requested Circulation: Average: 98%. Sept 2010 issue 98%.
16. This Statement of Ownership will be printed in the NOVEMBER 2010 issue of this publication. 17. Signature and Title of Owner STEPHEN ROBBINS, Publisher. Date 10/11/10. I certify that all information on this form is true and complete. I understand that anyone who furnishes false or misleading information on this form or who omits material or information requesed on the form may be subject to criminal sactions (including fines and impresonment) and/or civil sanctions (including civil penalties).

LCD Workflow Assistant

1 A full-color LCD shows the status of the function keys, navigation setting and workflow applets. It is based on an open software architecture, so customized applets can be developed. The following applets are supplied by default: Function Keys, Outlook Mail, Outlook Calendar, Outlook Tasks and RSS Feeds.

QuickView Navigation Keys

2 The QuickView Navigation Keys provide one-touch access to 32 standard views, enabling users to view models from any angle. Three buttons enable you to select each side of your model (six views) in conjunction with the center ISO1 and ISO2 button. Pressing the rotate button rotates the view by 90 degrees clockwise or counter-clockwise, depending on the length of the button press.

Function Keys

3 The SpacePilot PRO recognizes the active application and environment, and assigns 10 application commands to its Intelligent Function Keys. Users can also customize the function keys.

3Dconnexion Cap

4 You can push, pull, twist or tilt the 3Dconnexion cap to pan, zoom and rotate the model on the screen.

Navigation Setting Keys

5 The navigation setting keys allow users to customize how the controller cap works by restricting navigation to certain axes. The navigation keys also control how much force is needed to move a model on the screen.

TECH SPECS

- Six-degrees-of-freedom sensor
- 1.96x1.48-in. color LCD with a 320x240 resolution
- Keyboard modifiers (Ctrl, Alt, Shift, Esc)
- 21 programmable keys in total (access to 31 commands)
- Dimensions: 9.1x5.9x2.3 in. (LxWxH)
- Weight: 1.94 lb.
- 3-year warranty
- Minimum system requirements: Windows XP Professional SP2 (x86 and x64), Vista, and 7; Linux Red-hat Enterprise WS 4; SuSE Linux 9.3; Sun Solaris 8 (SPARC); Sun Solaris 10 (x86)
- Microsoft Outlook 2003 or later on host PC for accessing email, tasks and calendar information on the LCD
- USB 1.1

For more information, visit www.3dconnexion.com

→ **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.


MAKE YOUR IDEAS REAL
dimension.
3D Printers
3D PRINTERS
Create a more productive design process with Dimension 3D printers.
Dimension 1200es
DEAR BOSS,
PROJECT DELTA WOULD GET TO MARKET SOONER USING ONE OF THESE DIMENSION 3D PRINTERS!
- WE'D SEE & TEST MORE ITERATIONS AND DECIDE ON THE FINAL DESIGN FASTER.
- FUNCTIONAL MODELS ARE IN TOUGH ABS.
- LEASING STARTS AT $275/MO*
- CHECK OUT: WWW.DIMENSIONPRINTING.COM/DE
THANKS! — LARRY
HERE'S THE TEST PRINT OF OUR NEW STRUT. COOL!
*Price for US leasing. Additional options, shipping, applicable taxes/duties not included. ©2010 Stratasys, Inc.
STRATASYS