January 2013 / deskeng.com

DE Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

AMD and NVIDIA Graphic Card Reviews P.40

Prototype Embedded Systems P.36

Innovations in 3D Scanning P.28

FOCUS ON: Test & Measurement

Mesh with the Best

SCOPING OUT OSCILLOSCOPES P..34

SC12 RECAP P.38

Degrees of Freedom

by Steve Robbins

Welcome to 2013

Product design has evolved over the years I have been involved in technology publishing, and the growth of technologies that enable collaboration has had a lot to do with it.

When we first published *Desktop Engineering* in 1995, in small- to mid-size companies there wasn't much happening in terms of collaboration and team building. In fact, mechanical engineers, who were just starting to embrace 3D MCAD on their workstations, were often totally separated from the electrical and manufacturing engineers. They were charged with creating a mechanical model under the constraints of management and industrial engineers. While enterprise companies were the early adopters of 3D MCAD and CAE, especially automotive and aerospace companies, major engineering efforts by the mid-sized companies were often a compromise of disparate individuals with a lack of communication with each other.

> **Processes you have thought were beyond your expertise will be available to you soon.**

Tear Down That Wall

I remember Tony Lockwood, the founding Editor-in-Chief of *Desktop Engineering*, often referring to design engineers "throwing the model over the wall" to manufacturing. More likely than not, the manufacturing engineering team would include the electrical engineer to make sure all of the electronic components worked within the design.

About 10 years ago we started to see the walls crumbling. Part of this was because of analysis and simulation software being validated by real-world testing. Analysts opened their doors to the design engineer, and some simulation software was democratized. It also moved forward in the design process. Still, in the mid-market, simulations were often accomplished just before the actual prototype was manufactured and tested. Multiple iterations of simulations were only available to those that possessed high performance computing systems, and even then the process of completing multiple simulations was a manual process. Of course there was optimization software, but it wasn't widely used.

During the last 10 years, with compute power increasing and following Moore's Law, the use of simulation software has increasingly been used in the very beginning of the design process. Complex assemblies and systems have been modeled and meshed, then simulated. Collaboration has become the prevailing method of working on and completing a design.

Industrial engineers who used pens, crayons, clay and other materials to describe to the designer what the CAD model should look like are using readily available design software to express ideas that can be simulated before the parametric model is created. Prototypes can be printed and tested. Complete, large systems that just a few years ago could only be simulated as parts or small assemblies, are now being virtually tested, giving the engineering team a high expectation of what the outcome will be when the prototype is built. Electronics are designed with the mechanical model. Electronic components are often specified in the beginning of the design process and simulated as part of the complete system.

As engineering technology and processes have improved, so has the outcome of the designs. The Mars Curiosity Rover demonstrates how all of the design processes are coming together to create an optimized and successful design that cannot be physically tested as a system until it is deployed in, do I dare to say, a real-world environment.

Looking Ahead

It's easy to look back and see the changes that have influenced design engineering, but I would like to stick my neck out and look into the future of design engineering technology.

First, whether you like it or not, cloud computing is here to stay. We will be using it in the next few years in ways that even the experts cannot predict.

Second, your desktop engineering workstation will not be going away. It will not be replaced by an iPad or other tablet. It will continue to become more powerful. Local computing will be a better option for a long time to come. You will be accomplishing design challenges on it that you never thought possible today.

Third, while your jobs are not going to be getting any easier, your tools will be easier to use. This will be driven by design tools that are being made for non-professionals, like the maker market and gamers.

And last, processes that you have thought were beyond your expertise or means will be available to you in a very short time. Things like collaboration systems, sharing of designs, and optimization software to handle all those simulations you will be doing on your tricked out workstation.

And you will be designing things you never even imagined. I guarantee it. Happy New Year! **DE**

DE January 2013 VOLUME 18/ISSUE 5

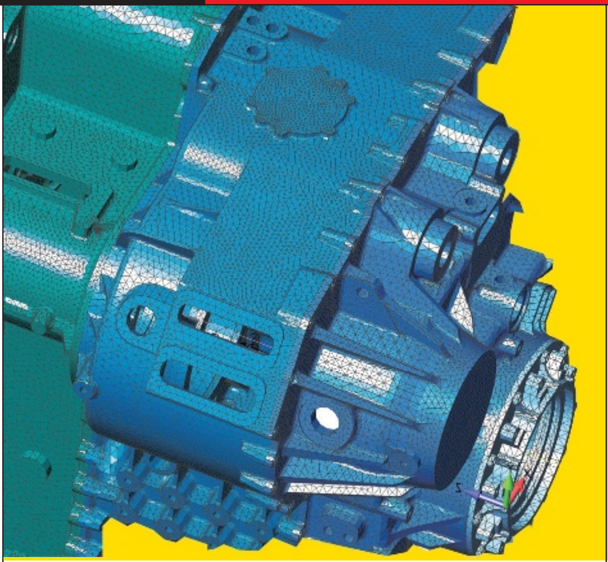

SIMULATE

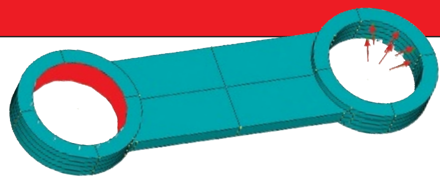

ENGINEERING COMPUTING

COVER STORY

Meshing Your Design for Analysis

FOCUS ON TEST & MEASUREMENT

ON THE COVER: Evaluating the best approach to meshing your design for analysis. *Images courtesy of Altair Engineering, ANSYS, COMSOL, PTC, and Siemens PLM Software.*

DEPARTMENTS

deskeng.com

RAPID READY TECH BLOG
Read all about making digital designs physical @ rapidreadytech.com.

ENGINEERING ON THE EDGE BLOG
See the future of engineering technology @ engineeringontheedge.com.

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

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

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

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

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

DE Desktop Engineering®

EDITORIAL
Steve Robbins | Editorial Director
Jamie J. Gooch | Managing Editor
Kenneth Wong | Senior Editor
Anthony J. Lockwood | Editor at Large
Heather Pittinger | Copy Editor

CONTRIBUTING EDITORS
Tony Abbey, Brian Albright, Mark Clarkson, David S. Cohn, Barbara Goode, John Newman, Susan Smith, Beth Stackpole, Peter Varhol, Pamela J. Waterman

PUBLISHER
Thomas Conlon

ADVERTISING SALES
603-563-1631 • Fax 603-563-8192

Erich Herbert | Sales Executive (x263)
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.

by Kenneth Wong & Beth Stackpole

Faster Rendering at Lower Watts with Dedicated Ray-tracing Unit

Rendering on the ray-tracing unit (RTU), a dedicated hardware for ray tracing, is expected to deliver faster ray tracing at lower power consumption. Caustic Professional, a division of Imagination Technologies, will deliver RTU-based rendering with plug-ins to popular rendering packages.

Alex Kelley, director of business development at Caustic Professional, a division of Imagination Technologies, has been waiting patiently to unleash the company's creation on the hardware market for quite some time. It's a processor, but neither a CPU nor a GPU. It's developed specifically for rendering. Kelley and his team call it a ray-tracing unit, or RTU.

Last month, at Autodesk University, Kelley and his colleagues decided it was time for the RTU to make its debut. So they set up a booth and demonstrated the latest generation of their hardware, branded as Caustic Visualizer. The RTU works with the CPU to deliver optimized rendering power.

The story of the RTU begins with a San Francisco start-up, once called Caustic Graphics. In December 2010, Imagination Technologies, an embedded processor developer, acquired the firm for $27 million. The RTU's evolution was shaped by, among others, Michael Kaplan, who previously worked on NVIDIA's professional rendering solutions. Kaplan currently oversees the RTU project as director of product management.

"We're not trying to be a high-performance computing (HPC) system that happens to do ray tracing," Kaplan clarified, noting that what other CPU and GPU makers are doing is "to cram a lot of cores into a unit to do parallel computing. But that's not optimized for ray tracing."

It is, in Kaplan's view, inefficient because ray tracing doesn't need the volume of CPU-based general-purpose computing power provided by HPC systems. The RTU is more efficient than ray tracing on GPU because, Kaplan said, "We're doing as much ray tracing as multiple graphics cards, but our card takes 55 watts. It doesn't need a big power supply."

At press time, rendering by the RTU is made possible by a beta plug-in developed for Autodesk Maya. The plug-in is slated to officially become available later this month, and another plug-in for Autodesk 3ds Max is in now in development, targeted for March 2013.

Whereas Autodesk Maya is the standard for entertainment and media content, Autodesk's professional design and engineering software packages (Autodesk Inventor and AutoCAD, to name but two) are usually bundled with Autodesk 3ds Max in most Autodesk Suites.

Dedicated to Ray Tracing

The RTU is designed specifically to address ray tracing, often the most compute-intense portion of the work involved in photorealistic rendering. The RTU won't provide the same general purpose computing horsepower found in professional GPUs or HPC systems. Therefore, the RTU won't be suitable for those seeking hardware to run simulation and analysis programs.

According to Imagination Technologies' announcement, "Caustic Visualizer is built entirely on the Imagination PowerVR OpenRL platform. Through OpenRL's highly optimized CPU implementation, Caustic Visualizer can deliver an interactive photorealistic viewport experience on a CPU-only desktop that was previously only possible in competitive solutions based on expensive quad-GPU hardware configurations … When coupled with OpenRL running on the new PowerVR ray-tracing hardware reference platform, Caustic's Visualizer software can deliver true full-frame (1,024x1,024) resolution imagery up to 15 photorealistic frames per second, without the need to de-res the image or add significant noise during interaction."

Currently Caustic Professional's RTUs are sold as separate accelerator cards, which fit into standard workstations and desktop PCs. Kelley envisions that, in the long run, the RTU's technology will become part of the GPU.

—K. Wong

aPriori Kicks Costing up a Notch

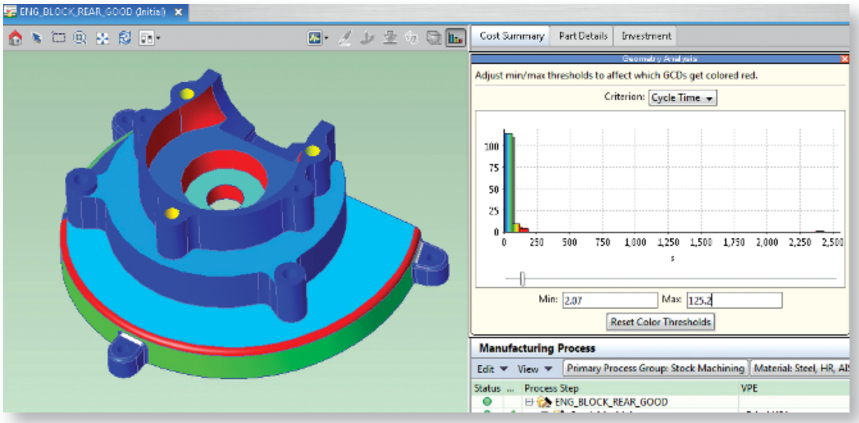

Through the use of an expanded geometry analysis tool, aPriori 2012r2 employs heat maps to deliver a quick visual overview of a product design, highlighting target areas for cost reduction. *Image courtesy of aPriori.*

As part of its vision to drive product cost awareness upstream into the earliest stages of product development, aPriori is stepping up its interpretation of the discipline with a new version of the software, which considers costing not just on a component or small assembly level, but takes an aggregate view of costing across the entire product.

The company's latest release, aPriori 2012r2, is stocked with an array of new capabilities to help engineers, as well as sourcing and manufacturing professionals, evaluate costs associated with building a product—be it industrial equipment, cars or something as complex and grand-scale as an airliner. While challenges vary across industries, global competition, ongoing regulatory requirements and a still-lumbering economy are driving a near-universal need for keeping product costs down in an effort to boost profitability.

"With more and more competition, there's a focus on trying to figure out what features drive the most value for customers," explains Julie Driscoll, aPriori's vice president, strategic marketing and product management. "Being able to quantify those features, make tradeoffs, and look at the product at higher levels to make strategic decisions is what is driving this need."

Import and Compare Costs

To address that need, aPriori is taking costing to a new level in this release, allowing users to import complex engineering bills of materials (BOMs) from enterprise resource planning (ERP), product lifecycle management (PLM) or other enterprise applications. Users can utilize that data against the aPriori cost engine to calculate costs for new components and subassembly designs while comparing those costs to initial targets.

While the original product allowed a design engineer in an automotive firm to evaluate costs around his or her specific responsibility—the design of shock absorbers, for example—the newer release builds on that component-level functionality. For example, it may provide additional views for engineering managers to consider costs around the greater chassis assembly—or at an even higher level, for a vice president of engineering to drill down into costs associated with the entire vehicle program.

Having the ability to view costs at an aggregate level is critical for providing visibility into what products are going to cost and what type of margins companies can hope to achieve, according to Rick Burke, aPriori's vice president of marketing. Viewing product costs at a more holistic level also expands aPriori's use from a departmental tool to an enterprise-class solution, he says.

"Having individual component cost information is important to guide specific decisions relative to those components," he explains. "Understanding where that sits at an aggregate level can help determine what a product costs—and if it is over cost targets, what areas to focus on for cost reduction."

Visualize Data

In addition to the BOM loader and other product cost roll up features, aPriori 2012r2 has a number of other notable enhancements, including a new geometry analysis tool that, through a color-coded heat map, provides an easy-to-understand visual roadmap as to where to most effectively focus on reducing costs. There are also enhancements to aPriori's Virtual Product Environments (VPEs), a digital representation of a real-world factory floor, for improved exploration of cost scenarios as well as streamlined system maintenance; an enhanced Progressive Die model for tooling estimates; and additional cost calculation functions for precision machining, including increased sensitivity to geometric tolerances.

—B. Stackpole

A Future with Driverless Vehicles Requires Sensory Adjustments

I'm lucky enough to live in San Francisco, a city with a fairly reliable public transit system. I can get around by jumping on a bus, a streetcar (not named Desire), or a train. So usually I read, fidget with my iPhone, or daydream while I let someone else do the driving.

A few years ago, I discovered that, once the trains go into the tunnel to complete the downtown route, the conductors relinquish control to a computerized system. Essentially, I have been riding in driverless trains for the past several years now—clearly a prelude to the not-so-distant future where private vehicles will also become driverless, or self-driven.

Running a train on a predefined track inside a tunnel, of course, is much simpler than driving a car on a highway shared by others. An autonomous car must have the capacity to make nearly all decisions currently made by a human driver—when to merge, when to get out of the way of another car, when to exit, when to brake, and so on.

Autonomous Autos

The underlying technology in self-driving cars is the specialty of Dr. Sandeep Sovani, manager of global automotive strategy at simulation software maker ANSYS. "A driverless car needs to have a good understanding of, and must keep track of, the positions of nearby objects, such as other cars, highway structures, and exits. One of the technologies that enable a car to do that is radar," says Sovani.

Radar simulation—specifically, simulating the reach and transmission of radar—falls under the electromagnetic segment. ANSYS added this capability to its repertoire when it acquired Ansoft, an electric design automation software maker, for $832 million in 2008. "You're actually simulating the electromagnetic waves," says Sovani.

This includes, for example, the shape of the transmission device and how that influences the way the waves travel in space. To be good enough to replace a human driver who can make instantaneous decisions about navigation, the complex sequence of radar signal processing has to be done in "milliseconds, or less," says Sovani. The other important aspect of the self-driving car is its ability to communicate with nearby cars and nearby structures transmitting navigation data. "It's all about antenna design," says Sovani.

Eyes on the Road

Because a self-driving car must be able to "see" its surroundings, the car's vision will comprise, Sovani predicts, radars that detect nearby objects, cameras that capture road conditions, and GPS devices that tell the car its own location at all times.

Simulating a self-driving car is more than simulating radars and onboard devices. It's about simulating the entire car as an interconnected system—something often referred to as model-based simulation. "To make sure that the system performs well, you have to simulate the system, not just its components," Sovani cautions.

The complexity of system simulation often requires far more computing power than what's generally available in workstations. It suggests high-performance computing (HPC) clusters will be an inevitable part of such exercises. Confidentiality clauses in contracts prevent Sovani from naming names, but he says nearly all the leading automotive makers are partnering with ANSYS to explore self-driving vehicles.

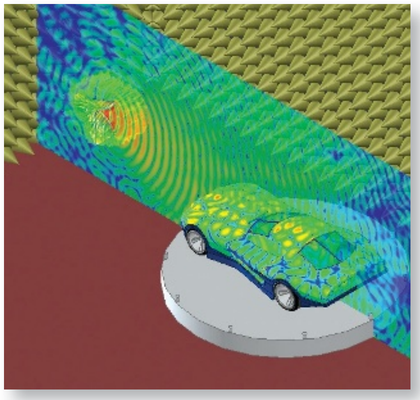

Simulating the electric field distribution and antenna radiation far field pattern at 1 GHz for a complete vehicle simulation, according to ISO 11451-2, using the conventional finite element method approach. The simulation was done using ANSYS software.

A Long Way to Go

"The currently available driver-assistant technologies [such as GPS-based navigation] ultimately aim to remove the driver altogether. They will continue to appear in cars at an accelerated pace from now on," he says. But don't expect a commercial autonomous vehicle to show up at a dealer near you in the next few years.

"My personal take: We're still a very long way from commercial application, probably 10 years away," Sovani says. "As a general rule, the last 20% of the development is solving 80% of the problems. And the remaining problems to solve are enormous. Just think of something like a car suddenly getting a flat tire, or a little rock lying on the road that needs to be avoided. Thousands of such scenarios need to be resolved properly."

Though still far away, a self-driving car will be a welcome change for a social media addict like me. When the technology is reliable enough to detect and avoid obstructions without my help, I can surf Facebook or write a blog post while riding in a car. DE

—K. Wong

Autodesk Rounds out 360 with Cloud-Based CAD

If there were lingering doubts about whether Autodesk was fully committed or merely experimenting with the cloud, that ambiguity should be put to rest with its latest announcement: Autodesk Fusion 360, what the company claims is the first comprehensive 3D CAD program to support the emerging delivery paradigm.

Obviously, Autodesk is no stranger to the cloud. In the past several months, it has debuted a full lineup of services, from consumer-oriented 3D modeling and 3D print solutions to full-blown professional tools, including its Autodesk 360 PLM and Autodesk 360 Simulation offerings. Still, the big holdout in Autodesk's cloud-based lineup was traditional 3D CAD. While users and experts have slowly been warming up to the new software delivery model, they've remained skeptical about CAD delivered in this fashion—stemming from concerns about the security of design IP, as well as issues related to performance when modeling complex 3D geometry over the web.

Dispersing Cloud Concerns

Keith Perrin, Autodesk's senior industry manager, says those concerns (myths, he calls them) are fast abating, and companies are coming around to the benefits associated with using CAD and other design tools in the cloud, including lower costs, ease of administration and the ability for easy collaboration and scalability, particularly in the face of globally extended design teams.

"We are seeing many more people become open to the idea of the cloud, and exploring how they might use it in a more intelligent way then when we first started these efforts," says Perrin, adding that Autodesk's cloud-based product lifecycle management (PLM) and simulation offerings have paved the way for this latest journey around CAD.

Autodesk Fusion 360 brings less-skilled CAD users along with built-in guidance to speed the learning curve, while there is deeper functionality for design experts. With the combination of the revamped user interface and cloud paradigm, Autodesk sees the tool appealing both to its traditional user base potentially looking for a complementary use case mode, in addition to smaller design shops that may have been shut out from the traditional CAD experience because of cost or lack of expertise.

Autodesk Fusion 360 is slated to be commercially available early next year.

—B. Stackpole

Hoosier Sets a New Pattern

With the help of Cimatron NC software, one Indiana toolmaker focuses on quality products and services—while growing the next generation of American engineers.

Hoosier Pattern is a leading pattern shop, with 25 employees and a 42,000-sq.-ft. facility in Decatur, IN. Founded in 1997, Hoosier has developed a reputation of being the preferred pattern provider for local foundries serving the automotive, agriculture, marine and consumer products industries.

Staying a Step ahead

Hoosier has developed close relationships with the foundries it serves, but that doesn't mean it can rest on its laurels, says Keith Gerber, one of the original company founders and its current owner and president.

Responding to numerous quotes a day is a matter of routine for Gerber.

"Cimatron is very helpful in the quoting process," he notes. "I use it to do a quick split, direction analysis, figure out a parting line."

The streamlined process and collaboration combine to keep customers happy, Gerber says.

Dave Rittmeyer, CAD specialist/supervisor for Hoosier, explains that process: "Once we get the job, we rely on Cimatron to figure out the best way to manufacture the part, make sure the parts are dimensionally accurate, split correctly, and with the proper draft. From that point, we'll use Cimatron NC to machine all the tooling."

MORE → deskeng.com/articles/aabhnp.htm

Taming the Wind

Cobham Technical Services and Boulder Wind Power collaborate to accelerate development of new wind turbine technology.

Aurora, Ill.-based Cobham Technical Services is helping Colorado's Boulder Wind Power (BWP) accelerate development of a new permanent magnet generator through use of advanced electromagnetic design software. Based on a permanent magnet, direct-drive design, BWP's generator uses an exclusive axial flux, air-core architecture that increases efficiency and reliability, and is believed to reduce the cost of wind-generated electricity to compete with fossil fuels.

In support of initial engineering studies taken on in early 2010, BWP chose to use the 3D version of the Opera electromagnetic simulator from the Vector Fields Software range of Cobham Technical Services. BWP cites the software's accuracy, the ease with which its analytical capabilities can be adapted to suit specific requirements, and Cobham's willingness to collaborate in further developing key areas such as dynamic modeling as the primary factors behind its choice.

How It Works

Like all permanent magnet direct-drive wind turbines, the generator rotor of BWP's 3.0MW design turns at about 13 rpm, necessitating a high pole count. Opera's advanced solvers allow high periodicity to be leveraged, so that the analytical model can be a fraction of the size of the complete generator—significantly reducing simulation times.

MORE → deskeng.com/articles/aabhnr.htm

A Piece of Quiet

Donaldson hits noise reduction targets with LMS acoustic simulation, an ideal solution for the design changes caused by looming Euro 6 emission standards and tougher OEM requirements.

Business is changing at Donaldson, one of the world's largest manufacturers of filters for trucks, buses, construction equipment and industrial machinery. In addition to orders for standard and custom air filters, a growing number of requests are for designing and building complete air intake systems for long-haul trucks. There's a sense of urgency, because truck engines are being redesigned to meet deadlines for strict Euro 6 emission standards taking effect shortly.

This shift in research and development work from the original equipment manufacturer (OEM) to the supplier follows the trend in other segments where responsibility for system development is increasingly delegated to subcontractors. The result is a win-win situation, with OEMs better able to focus on total vehicle design—and capable suppliers gaining new business from their system-level know-how.

Mind-boggling Constraints

As Donaldson engineers can attest, landing these major contracts is no simple matter.

MORE → .deskeng.com/articles/aabhns.htm

→ For the complete application stories, visit deskeng.com/fastapps

Turbocharging Engineering with the Intel Xeon Phi

The Intel Xeon Phi Coprocessor adds an exciting new dimension to HPC servers, giving engineers more options for analysis and simulation.

BY PETER VARHOL

Engineers are using increased Intel® Xeon® processor and system performance to simulate and analyze designs more frequently. Rather than build physical prototypes for testing, it has become possible to simulate designs iteratively and make incremental improvements quickly and with certainty. Because prototypes aren't being constructed, engineers can simulate and analyze, refine, and repeat any number of times during the design cycle.

The result is better designs, optimized for efficiency, cost, or other product characteristic. And the designs are completed more quickly, because there is no pause in the design process in order to build physical prototypes. If such prototypes are built at all, they are used at the very end of the process as a confirmation of manufacturability, or because of regulatory requirements.

Now it's possible to speed up simulations even more, using the just-introduced Intel Xeon Phi coprocessor. The Xeon Phi is based on Intel's Many Integrated Core (MIC) architecture, which uses Intel's industry-standard architecture and instruction set across multiple high-performance processor cores.

The fact that the Xeon Phi uses the same instruction set as the Intel industry standard architecture is perhaps its most significant advantage, because that enables engineers to run existing applications and custom code without modification. This provides an immediate path for increased performance and productivity for engineers, as simulations complete in a much faster amount of time, and with a much higher fidelity, than in the past. Using Intel Xeon processors with the Xeon Phi coprocessor, engineers can immediately boost their simulation performance.

As vendors and engineering software developers recompile and add new instructions to their applications, performance improvements are likely to be much higher. Software optimized for the Phi will be able to take advantage of specialized instructions to improve parallelism and increase execution speed.

Systems and Boards Already Available

At the Supercomputing 2012 conference and show, a number of the system and board vendors demonstrated systems that included a Xeon Phi slot, making them ready to go as soon as the coprocessor was available. The proliferation of Xeon Phi solutions indicates that this product will be adopted quickly and decisively by the engineering community.

Computational performance in parallel computations is where the Xeon Phi really shines. The Xeon Phi has up to 60 cores, with four threads per core, for a total of 240 threads. This design delivers up to 1 TFLOPS of double precision peak performance. Using the Xeon Phi in conjunction with one or two Intel Xeon E5 multicore processors will provide supercomputer-like capabilities for engineering design, analysis, and simulation work entirely on the desktop. A system configured with these processors, high-performance ECC memory, and flash storage will make engineers much more productive at design activities.

Putting the Power to Use

Engineering design is taking new computing power provided by the Xeon Phi and other Intel Xeon processors and radically changing the design process. Most and in some cases all design activities can occur on the Intel-based HPC server, with less time out for prototyping or simulating on large clusters and supercomputers.

Engineers are making significant strides with simulation on their individual workstations, from high-fidelity models of individual components to working virtual prototypes of entire systems and designs. Emerging HPC servers using the Intel Xeon E5 processors and Intel Xeon Phi coprocessor are driving the ability to accelerate and optimize the design process through an unparalleled degree of computation performance, supplementing the computation done on the workstation. This level of power lets more engineering design work occur on systems closer to the engineering group and individual engineer. This means more interactivity in the design process, optimizing the design in less time than ever. DE

INFO → Intel Corp: intel.com/go/workstation

Cooking up Self-Healing, Conductive Polymers

Researchers at the Stanford School of Engineering have come up with a mix of plastic polymers and nickel particles that is both conductive and touch-sensitive—and can even self-heal when damaged.

The material is made up of long polymers connected by weak hydrogen

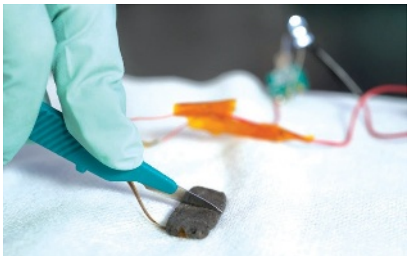

bonds, which enables its healing properties. The nickel particles increase its mechanical strength, and allow for conductivity. Not only can the material repair itself at room temperature, it can do so multiple times with no effect on its conductivity. According to a report in the *Stanford Daily*:

"To test the material's conductive properties, the researchers cut a strip in half with a scalpel. Seventy-five seconds later, they rejoined the material at the incision site and observed that the material retained 75% of both its original strength and electrical conductivity within seconds. In half an hour, the material regained nearly 100% functionality. The team also found that the same material could be cut and repaired as many as 50 times while maintaining its original bending and stretching capabilities."

MORE ➔ engineeringontheedge.com/?p=3201

Cray Laps Itself in Supercomputing Race

Cray continues to push the boundaries of supercomputing. Not only has the company's Titan system at Oak Ridge just nabbed the No. 1 spot on the Top500 list, but Cray is already shipping systems that promise to leave Titan in the dust.

The company has launched its next-generation XC30 supercomputer (previously code-named Cascade). It's designed to achieve HPC workloads of more than 100 petaflops and scale up to 1 million cores.

A number of major customers have already signed contracts to purchase XC30s, including the Swiss National Supercomputing Center,

the Finnish IT Center for Science, and the Department of Energy's National Energy Research Scientific Computing Center.

MORE ➔ engineeringontheedge.com/?p=3189

Compost that Server

The folks at Facebook would like to take green engineering to the next level and make biodegradable computers. The company has launched a contest as part of its Open Compute Project, asking students at Purdue University to come up with a biodegradable, compostable computer chassis. Teams will receive a server to use for the design process.

Students will participate in a Computer and Information Technology (CNIT) course at Purdue during the Spring 2013 semester, then break into teams to develop proposals. At the end of the semester, winning teams will receive funding and support to build a prototype.

MORE ➔ engineeringontheedge.com/?p=3182

Making Gasoline from Air

UK-based Air Fuel Synthesis claims to have created around 5 liters of petrol literally out of thin air—using a process in which carbon dioxide is combined with hydrogen molecules separated from water vapor. That process creates methanol, which is then passed through a gasoline fuel reactor to create petrol.

The process not only removes carbon dioxide from the atmosphere, but produces a clean-burning fuel that can be used in standard engines. The downside: The process requires massive amounts of electricity and water.

MORE ➔ engineeringontheedge.com/?p=3123

Autodesk Points Magic Finger at New Interface Technology

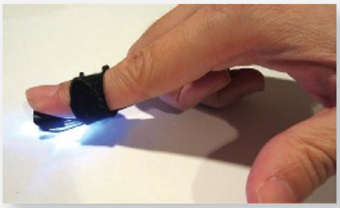

Autodesk Research has unveiled a solution for converting standard surfaces into touch interfaces. The Magic Finger is a small device that combines surface and motion detection in a thimble-sized form factor that can interpret finger gestures.

According to the company, the device senses touch through an optical mouse sensor. A micro RGB camera allows "contextual actions to be carried out, based on the particular surface being touched."

Magic Finger can recognize 32 individual textures, meaning it can tell whether it's being run across a table vs. a pair of pants, and commands could be programmed that are specific to a surface. It can also read 2D bar codes.

Autodesk Research developed the device in collaboration with the University of Toronto and the University of Alberta.

MORE ➔ engineeringontheedge.com/?p=3006

➔ For daily coverage on the future of engineering technology, visit EngineeringOnTheEdge.com

Objet Announces New Objet1000

Just before its merger with Stratasys was complete, Objet took a big step in additive manufacturing (AM) system with the release of the Objet1000. The new printer operates on the same principles as the company's Connex line, offering more than 120 different digital materials from which to choose. According to Objet, the system is also capable of combining up to 14 different materials in the same build. The Objet1000 comes bundled with Objet Studio.

The Objet1000 has a 39.3x31.4x19.6 in. build envelope, the largest yet offered by the company. To give you an idea of the size, a handcart is included in the package to facilitate unloading of the model.

MORE ➜ rapidreadytech.com/?p=2941

Creaform Releases Portable Go!Scan 3D Scanner

Creaform's new Go!Scan is a handheld 3D scanner that is being marketed to industrial and professional users. Go!Scan uses a white LED light source, with a resolution of .500 mm, and an accuracy of 0.1 mm. It can cover a 15-sq.-in. scanning area, capturing up to 550,000 measurements per second. Creaform claims the speed of the Go!Scan is 10 times faster than other 3D scanners.

The new 3D scanner weighs 2.4 lbs.—making it heavier than an iPad, but lighter than most cordless drills. It is point-and-shoot, or the scanner can be fixed in place to capture data while an object is rotated in front of it.

MORE ➜ rapidreadytech.com/?p=2691

Stratasys and Objet Complete Merger

As *Desktop Engineering* reported previously, two of the largest companies in 3D printers, Stratasys and Objet, announced their intent to merge in April last year. The boards of both companies have now approved the deal, which will see the combined stock hit an equity value of around $1.4 billion. The new company retains the Stratasys name with dual headquarters in Eden Prairie, MN, and Rehovot, Israel.

Scott Crump, CEO, chairman and co-founder of Stratasys, will act as the chairman for the combined company. A four-man executive committee, headed by Elchanan Jaglom, current chairman of Objet, will oversee integration and implement business strategy. David Reis, CEO of Objet, will take the position of CEO for the combined company.

The merger will allow the new company to offer two different, yet complementary, 3D printer technologies: Stratasys' fused deposition modeling (FDM) and Objet's PolyJet technology.

MORE ➜ rapidreadytech.com/?p=1094

Shapeways Builds 'Factory of the Future'

Based in New York City, Shapeways provides additive manufacturing (AM) services and also acts as a clearinghouse for AM designs. Customers can log onto the site, browse through hundreds of designs in categories such as art, games and jewelry, and place an order that will be created using AM. Prices range from a couple bucks to thousands of dollars—and business, apparently, is booming.

The company has begun construction on a 25,000-sq.-ft. "Factory of the Future" in Long Island. New York City Mayor Michael Bloomberg was on hand for the ribbon-cutting ceremony for the new

facility, which Shapeways claims will have the capacity to 3D-print 3 million to 5 million custom products a year.

Plans for the building also include a laboratory for "research and development of new materials, post-production techniques, and community experimentation." The Factory of the Future will be open for tours as well, further driving interest in AM technology.

MORE ➜ rapidreadytech.com/?p=2677

A Conversation with Terry Wohlers

Well-known 3D printing guru Terry Wohlers started his company 26 years ago, with a focus on CAD/CAM applications and consulting. He had published a book on AutoCAD (the 20th edition of which was published last August) and, during the course of his ongoing research, he read about the launch of a brand-new company named 3D Systems and a new technology called stereolithography (SLA).

"I read an article about stereolithography and I was intrigued by it," Wohlers recalls. "I contacted the company; they sent me a VHS videotape and a part, a full-scale distributor cap, and I was absolutely astounded by it."

Wohlers began to study the technology and follow developments in the industry ...

MORE ➜ rapidreadytech.com/?p=2766

➜ For more coverage of rapid technologies, visit RapidReadyTech.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.

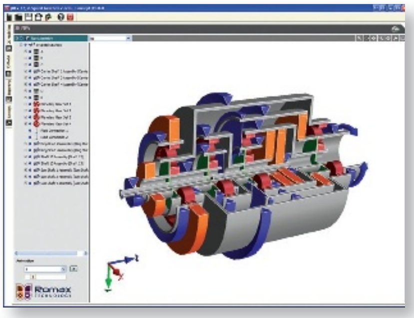

Software for Driveline and Gearbox Design and Optimization

Concept development application allows up-front optimization of gearbox and driveline systems.

Concept from Romax Technology in the UK seems to be a logical way to go about designing driveline and gearbox systems. In a nutshell, Concept is for the upfront design and optimization of such systems used in car transmissions, power tools, and even wind turbine gearboxes.

"Upfront" is the keyword here. Concept's reasoning for being, so to speak, is to provide engineers a comprehensive concept development platform so that you can make better decisions at the beginning of your design process quickly.

MORE → deskeng.com/articles/aabhay.htm

Mcor Technologies Launches the Matrix 300+

3D printer that uses paper for build material is enhanced to reduce build time.

The other day, I came across the Matrix 300+ 3D printer from an Irish outfit called Mcor Technologies. One intent of this device seems to be to break this bottleneck for SMBs, or any outfit wanting to get into 3D printing that has not satisfied their concerns about its total operating costs. Here's why I think that: The Matrix 300+ uses regular old letter-sized paper and a water-based bonding agent to produce eco-friendly models.

Parts are built using a layering technology the company calls Variable Volume Deposition (VVD), which seems to offer a way to make parts both more durable as well as complex.

MORE → deskeng.com/articles/aabgzp.htm

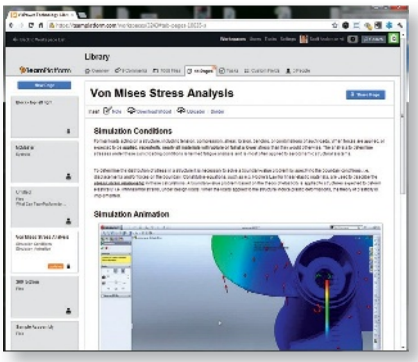

Cloud Collaboration Portal Supports Exporting to Excel

Complimentary 3D model sharing and visualization service also opened.

TeamPlatform is a cloud-based service offering a bevy of CAD project collaboration and management tools. You need a browser to use it. It keeps current on all the stuff that's been driving you nuts, so right off the bat you can eliminate installing updates and forking over dough for add-ons.

It starts with file sharing ,viewing and mark-up. It seems TeamPlatform supports every file type known to CAD, including the majors natively and 2D SVG files. You can do all the neat stuff: explode and rotate views, take measurements, annotate, highlight with arrows, and create discussion threads within the context of your model and/or project.

MORE → deskeng.com/articles/aabgyf.htm

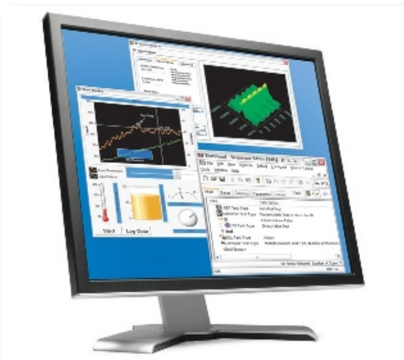

NI TestStand 2012 Increases Automated Test Throughput

Provides flexibility with new modular framework; said to be ideal for parallel tests.

Look, I know that developing any sort of automated test system is a tough business. This is why over the years I have found NI TestStand test management software such an intriguing product. It helps you build automated tests, run tests of your tests, and collect results. It also helps you deploy the tests when that time comes. Let's take a look at it.

Out of the box, NI TestStand provides scalable test development tools. These include things for reporting test results, testing devices in parallel, customization, logging results, and deploying your final version to production.

MORE → deskeng.com/articles/aabgxb.htm

How to Configure the Ultimate Engineering Workstation

Why you need overclocking, SSD caching and dedicated rendering.

Most engineering applications, like SolidWorks, Autodesk Revit, or Autodesk Inventor, are frequency bound (meaning they predominantly use only one processing core), so their ideal hardware platform is a workstation with fewer cores but higher frequency. However, if your workflow also includes rendering and simulation (which use multiple cores simultaneously), you'll need the maximum number of cores to run at peak performance. For most workflows, a quad core Intel® Core™ i7 processor is optimal, especially if it's overclocked like those found in XTREME edition workstations from BOXX.

Overclocking Matters

While brand name computer manufacturers focus most of their attention on mass produced, general use PCs, BOXX offers professional grade workstations custom configured and overclocked for 3D visualization. And with the ability to achieve 4.5 GHz, overclocked 3DBOXX XTREME workstations hold a decided advantage over competitors' top-of-the-line models, which can only manage 3.7 GHz— the speed threshold since 2006.

"It's the frequency plateau," says Tim Lawrence, BOXX Technologies' VP of Engineering. "Improvements to architecture have helped somewhat, but not enough. With processor speeds remaining virtually stagnant for six years, overclocking is the only way to significantly increase core speed and thus, performance."

Faster processes result in an accelerated workflow, greater efficiency, higher productivity, and a better overall user experience.

And if you're concerned about the effects of overclocking on a processor, rest assured knowing BOXX has shipped overclocked systems since 2008 and with thousands of systems in the field, the company has not experienced a processor failure rate any different from that of standard processor systems. And like all BOXX systems, XTREME systems are backed by a three-year warranty.

Critical Components

3DBOXX 4050 XTREME and 4920 XTREME performance is enhanced by the option of Intel® Smart Response Technology, whereby the system automatically learns which files users access frequently and copies them from the hard drive to the solid state drives. The next time the files are requested, they're loaded from the SSDs, rather than the slower hard drive. The result is faster booting, faster application loading, and accelerated performance.

In regard to system memory, you'll need at least 8 to 10 GB. With this amount, if your workstation is responsive and executes tasks quickly, you've made the right choice. If not, you may need to increase your RAM (in many instances) to as much as 16GB.

Although an NVIDIA Quadro 600 card is serviceable, BOXX recommends the NVIDIA Quadro 2000 as an ideal mid-range graphics card for most engineering workflows.

Because rendering is a key aspect of engineering workflows, engineers should consider off-loading it to a dedicated rendering system like BOXX renderPRO. Available with Intel® Xeon® E5-2600 series processors, renderPRO features up to 16 processing cores. It enables users to deliver complex projects within budget and on schedule by drastically reducing rendering time.

Increasing Productivity and Profit

The key to increasing productivity and profit is to accomplish more in less time. Faster turnaround means fewer employee hours invested and more time for new projects and clients. So when you configure the ultimate engineering workstation, consider solutions you won't find anywhere else delivering performance you can't get anywhere else—overclocked 3DBOXX XTREME workstations. DE

INFO ➜ BOXX Technologies:

www.boxxtech.com/solutions/solutions.asp

Highly Recommended

For engineering workflows, BOXX recommends the following workstations:

3DBOXX 4050 XTREME is a liquid-cooled workstation, powered by an overclocked quad core, Intel® Core™ i7 processor running at 4.5 GHz. Available with up to two GPUs (NVIDIA Maximus™ technology) and support for solid state drive (SSD) caching for increased storage performance, 4050 XTREME is the industry's fastest single socket workstation for engineering and product design applications.

3DBOXX 4920 XTREME, another liquid-cooled BOXX workstation, includes an overclocked, six core, Intel® Core™ i7 processor also capable of speeds up to 4.5 GHz. 4920 XTREME is available with up to four GPUs (NVIDIA Maximus™ technology), and support for solid state drive caching for increased storage performance.

Meshing Your Design for Analysis: *Which Path to Take?*

Automation makes the job easier, but flexible options still rule.

BY PAMELA J. WATERMAN

Whether generating ballpark figures or finely detailing non-linear responses, engineers are increasingly using simulation software to analyze mechanical design behavior. A necessary part of the process requires subdividing the design into a mesh of 2D or 3D elements defined by nodes and closely following the geometry. This task can be tackled in one of three ways:

- work within today's CAD-embedded programs;
- use tools native to a separate analysis package; or
- bring in a specialized third-party tool.

Complexity generally increases with the size of the model, the physics involved, or the number of physics you're studying. Keeping these factors in mind should help you evaluate options. *DE* takes a look at a number of these meshing approaches—with their strengths and trade-offs—to put you on an efficient pre-processing path.

From CAD on up

Pointwise has almost three decades of experience in computational fluid dynamics (CFD) pre-processing, so Rick Matus, the company president, has a broad perspective on the topic.

"There is a ladder of increasing control [and decreasing automation] as you move from CAD-embedded CFD to standalone CFD packages, to adding a separate CFD meshing package like Pointwise," Matus explains. "Likewise, as you look at the people using these various CFD approaches, you are moving from someone doing basic fluid flow to someone doing more difficult applications, to someone with difficult applications and demand for accuracy."

Stephen Endersby, SolidWorks senior product manager, sees a similar variation along the spectrum of the tools themselves: "Most CAD packages are good at capturing geometry. If you have an embedded mesher, because you have access to the actual feature history parametrics, you're probably going to capture the geometry very precisely."

Endersby notes that an external mesher may have a harder task, as it has to do some recognition of the geometry. This task can take more time, but that mesh will probably still be of equal worth.

However, he adds, if you also need to capture stress changes and gradients, you need adaptive meshing capabilities that recognize, perhaps, the need for a finer mesh in the area of a fillet or sharp corner. Today's SolidWorks can do a great deal of that, too, and may be sufficient.

The target audience for doing analysis within CAD programs is the generalist mechanical engineer (as well as an analyst), and that person may or may not have also created the geometry. With this group in mind, Christos Katsis, PTC vice president of R&D simulation software, notes his company's Creo Simulate software tries to automate the numerical controls that the analyst has to use.

"By being an integral part of the CAD environment, having an intimate knowledge of the geometry, we can provide powerful ways to automatically clean up or heal defects," says Katsis. "Additionally, not only in Creo Parametric but also in Creo Direct, there are tools to simplify [the model] by suppressing features or modifying geometry."

If you're using a product data management (PDM) system, you can store iterations of the model.

"Given the direct associativity between the simulation and the CAD model, doing edits for 'what if' studies or sensitivity and optimization studies can be performed without additional work," Katsis says. "You may only need to partially remesh."

Have FE/CFD Software, Will Mesh

As soon as you start looking at meshing from the perspective of the finite element analysis (FEA) or CFD analysis software developers, you find many more ways to approach the task. Your choice will depend both on what software you're used to and the level of complexity of your project.

NEi Software now offers NEi in-CAD for SolidWorks (a progression from NEi Fusion), as well as NEi in-CAD for IronCAD, supporting meshing and analysis within a familiar environment. For more analysis-heavy users, the company markets a version of Siemens PLM Femap pre-processing/meshing software optimized for use with NEi Nastran in its various flavors.

Nick Trebilcock, NEi Software product development manager, points out one difference where the latter approach may be preferred, noting, "CAD models are modeled with mathematical accuracy, but may not actually be modeled with real-life accuracy."

Mitch Muncy, NEi's director of engineering, concurs: "For the most accurate results, you go into more advanced tools like Femap, where the goal is to obliterate the geometry and turn it into the perfect mesh. By doing so, you've got the best mathematical model that represents what you originally created in CAD." He says specialized meshing tools take this further, going after that extra 10% improvement.

The CAD-embedded versions of an analysis package offer features many designers prefer. Baskar Rajagopalan, CD-adapco product manager for CAD client integration, says that these include not having to do geometry translation, having the same look and feel as the CAD product, and being able to do more work up front and concurrently, with a tightly coupled workflow for modifying the geometry and rerunning the analysis.

Behind the scenes of CD-adapco's CAD versions are the STAR-CCM+ meshing engine and solver themselves.

"The CAD clients give you the ability to either push just the geometry into STAR-CCM+ for all the simulation that needs to be done, or push both the geometry and the simulation data," says Rajagopalan, offering such examples as boundary conditions or mesh. "Engineers, casual users and CFD analysts can all benefit from using CAD clients."

Altair Engineering's HyperWorks software includes HyperMesh, a powerful pre-processor first developed in 1990 and used extensively in automotive and aerospace applications. While readily handling simple models, it really shines when tackling detailed geometry or mesh in very large models. You can import geometry, other meshes and such information as ply thickness and orientation for working with composites. The software can automatically assemble and recognize spot welds, seam welds and bolted assemblies, and has full support for adhesive modeling.

Users can run HyperMesh in an auto-cleanup mode, where you define such

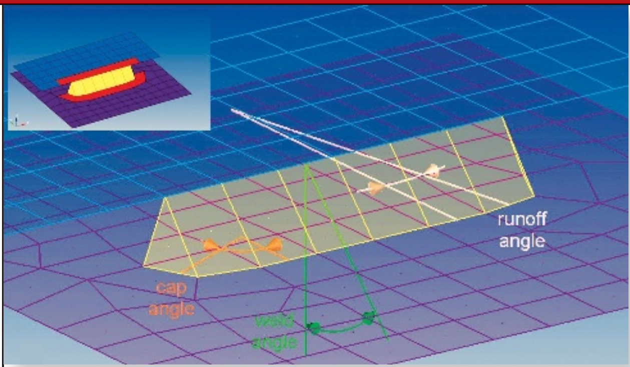

Altair HyperMesh pre-processing software includes connector technology that automatically assembles individual parts, modeling such structures as a seam weld. *Image courtesy of Altair Engineering.*

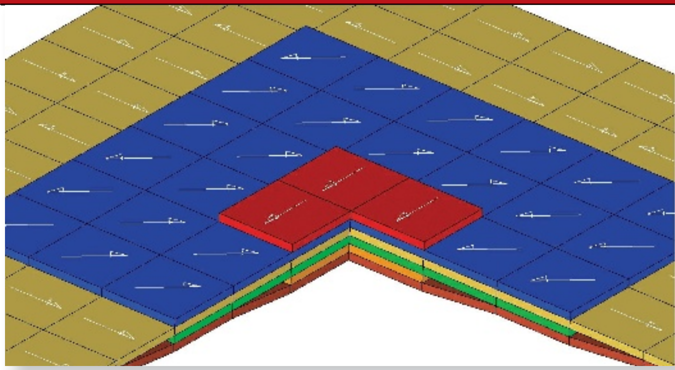

Altair HyperMesh is a high-performance finite element pre-processor that can prepare the largest models, even including composite structures, for analysis runs in different disciplines. *Image courtesy of Altair Engineering.*

criteria as what to do with holes smaller than a certain diameter. A batch mesher option is the most automated process, including geometry cleanup and meshing whose quality can be precisely measured; this is useful if you have, say, a full vehicle, that might consist of many hundreds of parts. You can submit them all at one time—and, by using multiple CPUs, get through a full vehicle quickly.

"Process automation doesn't mean you need to compromise on accuracy," says Uli Gollwitzer, Altair Engineering marketing manager, modeling and visualization "You are able to set up all these control parameters, and you are still going to get the same level of quality."

Even the more specialized analysis packages can still target the engineer who is not a full-time analyst. To do so, Autodesk offers three choices in its Autodesk Simulation line, based on products acquired from Algor, Blue Ridge Numerics and Moldflow, as well as including some meshing inside Inventor Optimization.

Autodesk Product Marketing Manager Luke Mihelcic says customers need to spend time on efforts other than meshing, so the approach provides robust, fully automated steps inside their products that remove the guesswork. However, users can still choose to manually loosen, tighten or refine a mesh as desired.

Mihelcic notes that users get helpful visual results to make more adjustments.

"There are feedback tools that say, 'There are a lot of thin surfaces or thin edges in this model—you may want to consider adjusting your mesh in this location,'" he offers as an example. He adds that you can do mesh optimization studies within Autodesk Simulation CFD, telling it to run, for example, three or five mesh sensitivity studies, based on your direction to look at areas of high temperature or high pressure, etc. The software determines where there is or is not a lot of activity, and refines the mesh automatically.

The whole idea of evaluating meshing independently of the mathematical solvers invites discussion.

"Our view is that meshing is an integral part of the whole simulation process," says Barry Christenson, ANSYS director of product marketing. "The more automated, the more integrated, the better. It also requires a high degree of synchronized development with your solvers. Maybe as you're developing new advancements in automation, you're developing new element formulations—to keep up with changes in element shape and uniformity. Your solver needs to utilize those meshes effectively."

The capability to handle complex physics and their necessary meshes is reflected in COMSOL Multiphysics analysis software. It supports creating flexible, highly detailed versions, as well as importing third-party creations.

A Quick Look at Meshing Options

Editor's Note: Rick Matus, president of Pointwise (CFD meshing software), shared with Pam Waterman a quick pro/con comparison of meshing capabilities:

If you are looking for a highly automated approach, maybe only requiring a single button-push to get results, then the CAD-embedded computational fluid dynamics/finite element analysis (CFD/FEA) solutions are the way to go. Engineers interested just in gross effects of the basic flow will probably be happy with CAD-embedded apps. The disadvantage is it does not always produce accurate results, and it is not always readily apparent when the results are wrong and by what degree.

A standalone CFD/FEA package would typically provide more control over the mesh, solver and physical models. A CAD-embedded app might not include the model of a particular application, so you need to go to the standalone app to get the desired accuracy. Likewise, you may know a finer grid is required in a localized region in order to resolve a flow feature there. Highly automated meshers do not give you the ability to control the mesh to that degree, so this would be another case where the standalone CFD package would be needed.

Sometimes even the standalone CFD/FEA apps do not provide enough mesh control, so you need a separate mesher to get the highest accuracy.

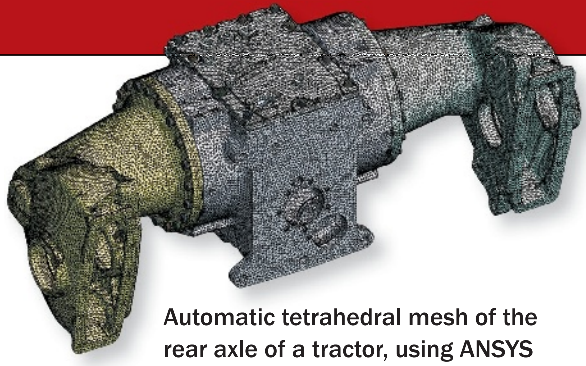

Automatic tetrahedral mesh of the rear axle of a tractor, using ANSYS software. *Image courtesy of ANSYS.*

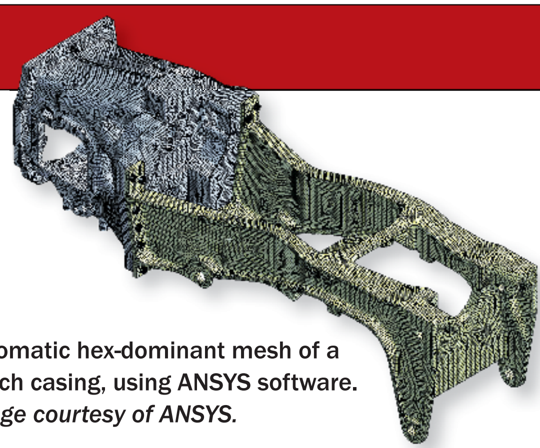

Automatic hex-dominant mesh of a clutch casing, using ANSYS software. *Image courtesy of ANSYS.*

Comparing the two approaches, Lorant Olasz, COMSOL program manager of external interfaces, says the advantage of working solely within a certain package is that the software developers really know their users' needs. Because the mesh may need to be constructed differently for each application, COMSOL Multiphysics offers tools that support optimizing the mesh for CFD, plasma, structural and many other analysis types. Physics-controlled meshing (with nine different preset mesh size settings) and user-controlled meshing each offer refinement options.

For users who rely on legacy systems or have in-house groups that do the meshing, Olasz suggests building a single mesh, then handing it off for designers to use in multiple analysis packages with structural or thermal specializations. With that in mind, COMSOL supports importing NASTRAN format meshes, and allows splitting a mesh into domains based on material data or element type.

Automatic adaptive remeshing is a hot topic, and it's one that figures prominently in SIMULIA's Abaqus/CAE software. The feature allows the program to work iteratively with Abaqus/Standard or Abaqus/Explicit to determine an optimal mesh—again demonstrating an advantage of working within a single environment.

For example, a sequential coupled thermal-structural analysis can progress from a coarse to fine mesh via a series of jobs submitted between the two programs. Each job communicates user-defined error-indicators that control remeshing iterations, balancing computational cost and accuracy. Such an optimized mesh would be difficult to generate manually.

The Abaqus documentation offers several recommendations about meshing decisions. The company even has a dedicated training class on selecting appropri-

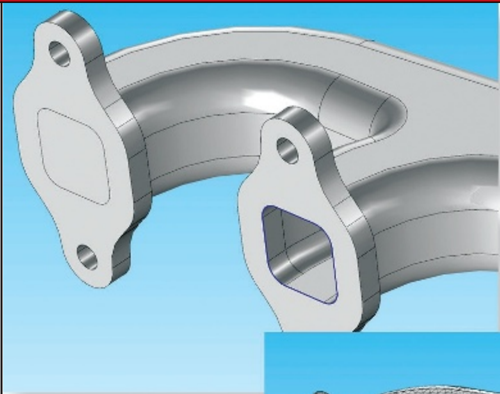

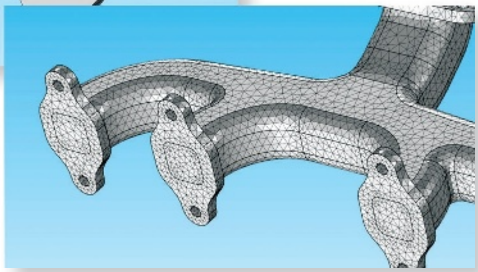

COMSOL Multiphysics' Cap Faces geometry operation is designed to make it easy to cover the ends of fluid channels and then mesh the interior of imported CAD parts. *Image courtesy of COMSOL.*

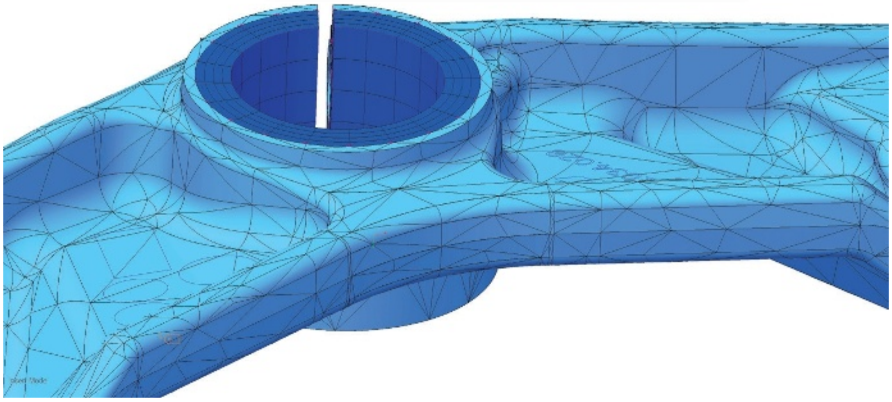

With PTC Creo Simulate, high-order, p-type finite elements follow the model geometry exactly. If part of the model is meshed with bricks for solution efficiency, the brick mesh is automatically transitioned and bonded to an unstructured tetrahedral mesh in the rest of the model. *Image courtesy of PTC.*

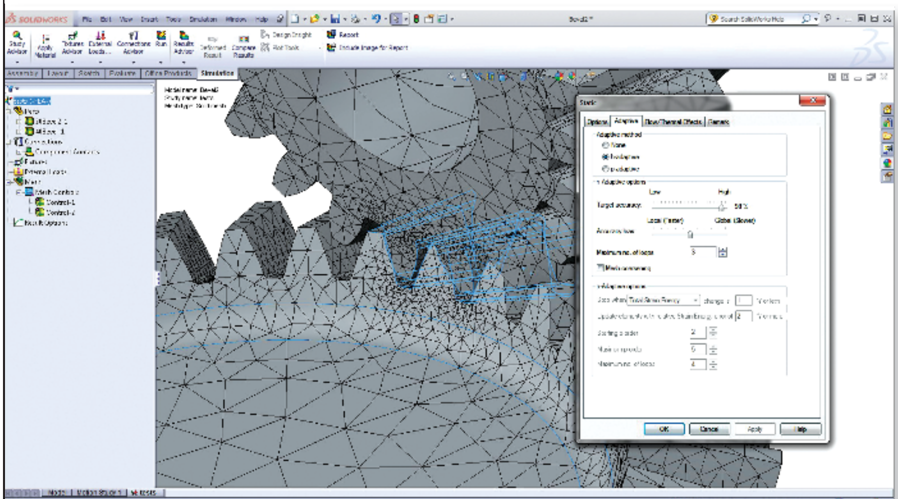

H-adaptive meshing control detail performed within SolidWorks simulation software. *Image courtesy of SolidWorks.*

ate elements that help with special situations, such as bending or incompressible materials or a particular application.

"We've found that most users learn very quickly what are the different element topologies and types to use," says Asif Khan, Abaqus/CAE product manager.

Third-party Power

Pointwise CFD pre-processing software is joined by several other major, standalone packages that support detailed meshing capabilities. As mentioned earlier, Siemens PLM Femap is a well-known tool that people turn to when they are faced with challenging modeling or meshing problems.

"CAD-embedded FEA programs are great for creating push-button solid models from solid CAD data. However, you can very easily create huge models that become difficult or impossible to solve," says Al Robertson, Seimens Femap marketing manager. "Femap provides a variety of capabilities to manage model size by idealizing the structure. It also supports more advanced analysis by creating specialized elements like nonlinear springs, rigid elements, spiders or mass elements."

Robertson says the functionalities in the Femap Meshing Toolbox include tools to edit a mesh interactively once it's been created. It affords the user control over mesh sizing, seeding along edges, and even the ability to move individual nodes while conforming to the underlying geometry.

"In the just-released version, we're adding the ability to create surface geometry from existing shell models, modify or update the geometry, then remesh," he adds. "This capability allows users to more easily make edits and updates to legacy finite element models for which there is no geometry available. "

For Robert Rainsberger, president of XYZ Scientific Applications, there are at least a dozen reasons to use TrueGrid, his company's specialized mesh generating software. He says many of the advantages stem from the fact that TrueGrid generates hexahedral element meshes instead of the automatic tetrahedral meshes commonly bundled with CAD systems

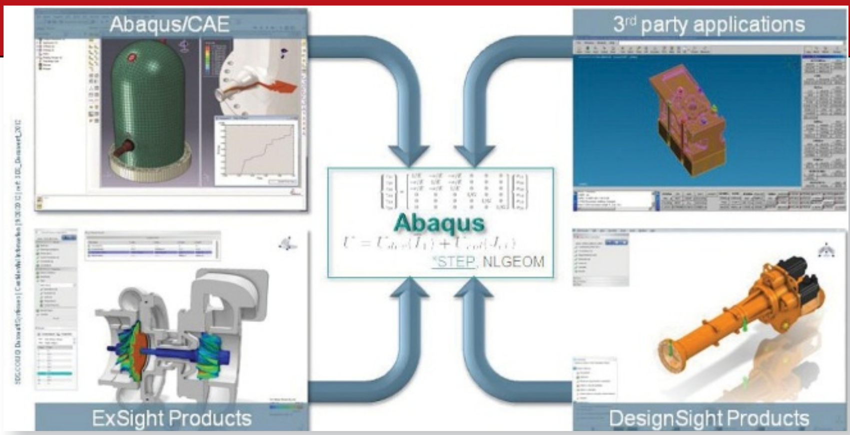

Workflow diagram for using SIMULIA's Abaqus/CAE multiphysics analysis software with optional third-party meshing functions. *Image courtesy of SIMULIA.*

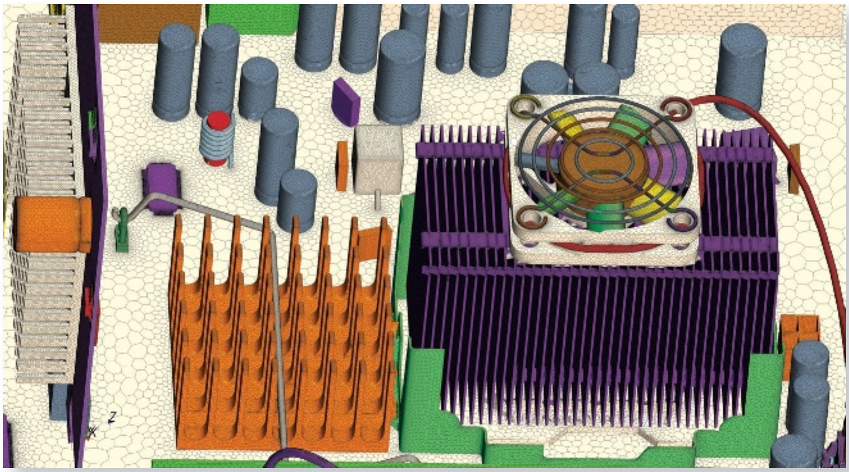

Mesh of an electronics-cooling layout for CFD analysis, performed with CD-adapco STAR-CCM+ software. *Image courtesy of CD-adapco.*

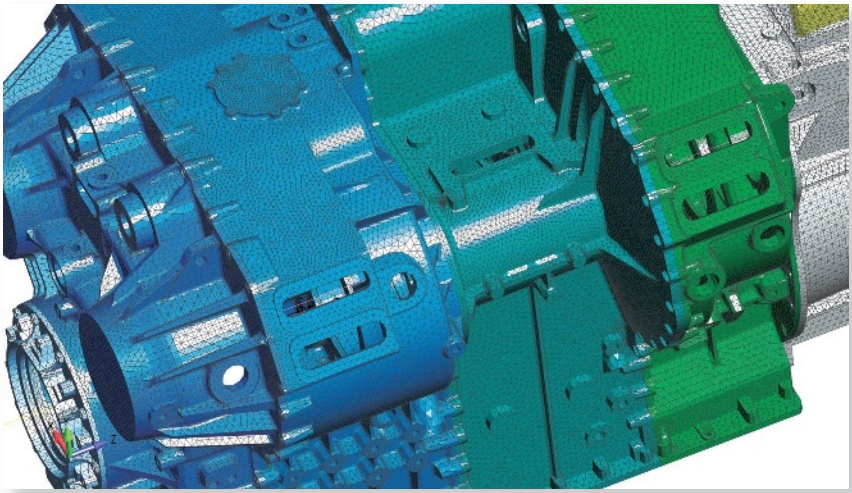

George Laird, principal mechanical engineer and owner, Predictive Engineering, says, "With the new automatic geometry preparation feature in Femap 10.3, our average speed for meshing a typical casting model is nearly 50 times faster. It's is a true game-changer in the world of meshing complex solid models." *Image courtesy of Siemens PLM.*

or simulation codes. His top four include:

1. There is no need to clean up the geometry, because TrueGrid can handle gaps and overlaps.

2. The software can handle complex geometry, with no limit on mesh size. Rainsberger says a 1 billion-node TrueGrid mesh is unusual, but not uncommon.

3. TrueGrid is not restricted to common surface types.

4. Any aspect of the model can be parameterized, letting users try a variations either manually or via an automated search.

The rest of Rainsberger's Top 12 TrueGrid benefits address such features as being able to form boundary layers with high aspect ratios (typical for CFD analysis), dealing with anisotropic materials such as fiber composites, and accurately modeling elastic-plastic material models. He says it pays to use a versatile mesh generator to ensure success, flexibility and desired accuracy every time.

With all these choices, it's a good thing that the company websites offer detailed comparative charts and user webinars. Gather up your models; it's time to go for a meshing test drive. **DE**

Contributing Editor **Pamela Waterman**, *DE's simulation expert, is an electrical engineer and technical writer based in Arizona. Contact her via DE-Editors@deskeng.com.*

INFO ➜ Altair Engineering: AltairHyperWorks.com

➜ **ANSYS:** ANSYS.com

➜ **Autodesk (Algor):** USA.Autodesk.com/simulation-software

➜ **CD-adapco:** CD-adapco.com

➜ **COMSOL:** COMSOL.com

➜ **NEi Software:** NEiSoftware.com

➜ **Pointwise:** Pointwise.com

➜ **PTC:** PTC.com

➜ **Siemens PLM:** Siemens.com/plm

➜ **SIMULIA:** SIMULIA.com

➜ **SolidWorks:** SolidWorks.com

➜ **XYZ Scientific Applications:** TrueGrid.com

Meshing for FEA

Become more productive by mastering this technique.

BY TONY ABBEY

***Editor's Note:** Tony Abbey teaches the NAFEMS FEA Intro class in Stratford UK, Calgary Canada, and Los Angeles over the next three months. Contact tony.abbey@nafems.org for details.*

It is a sobering fact that structures can take up an enormous part of an engineer's time in the meshing phase. A complex fabrication may use very little of available CAD geometry when meshing—or worse, be just a 2D drawing image. Meshing can then hit 80% of project time.

For most structures, a good understanding of meshing is important for productivity. Factors include:

- effective techniques in your meshing tool;
- how analysis objectives affect the mesh;
- correct element types and numbers;
- loading and constraint regions; and
- the level of accuracy needed.

Many other considerations also affect meshing. Let's review meshing within the finite element analysis (FEA) process, and how to get the most out of your software.

3D Idealization

Idealization defines how we simulate structural characteristics of the component. A component may be defined by CAD solid geometry, having all the information needed to calculate stiffness and strength. The geometry can be "cleaned" to remove irrelevant features such as tooling holes, bolt holes, handles, electronic components, etc. Cutouts and fillets that define detail, but are away from areas of high stress concentration, can be ignored. We apply loading and constraints directly to geometry faces—or edges.

The geometry has fully defined everything we need (see Fig. 1). We mesh with 3D elements, in a 3D idealization.

Alternatively, we may have a 3D CAD model of a fabrication (such as an aircraft fuselage; see Fig. 2). We cannot mesh with solid elements, because the level of detail implies a model of the order of a billion (1E9) or more elements. The practical performance limit today is around 5 million (5E6) elements!

A fabricated structure requires a 2D/1D idealization strategy (see Fig. 3):

- 2D thin shells (skins, bulkheads and spars)
- 1D beams or rods (frames and stiffeners)

We may also use shells for the frames and stiffeners. Surface geometry is derived, and then meshed with 2D shells and edges with beams or rods. The derivation of the geometry

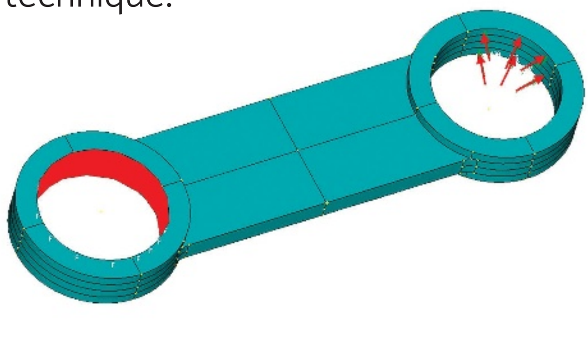

FIGURE 1: Component loaded and constrained at surface geometry.

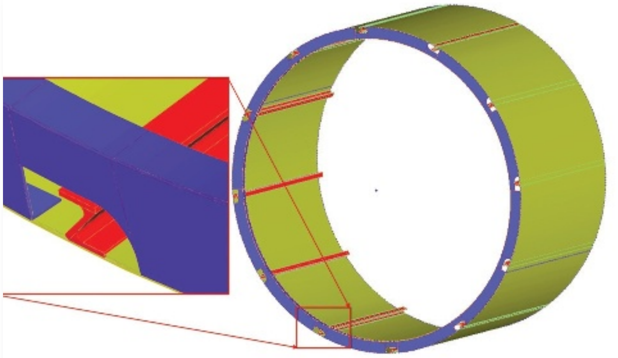

FIGURE 2: Typical fabricated structure with solid geometry.

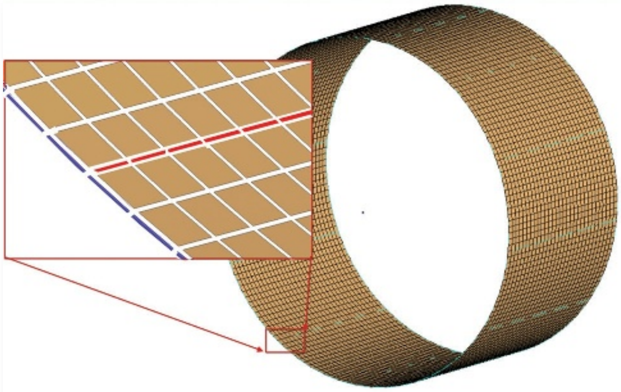

FIGURE 3: Typical fabricated structure and mesh.

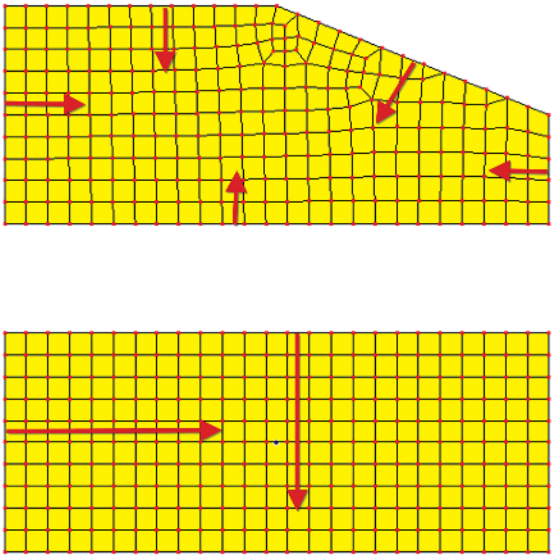

FIGURE 4: Regular mesh (bottom) and "Pac-Man" type mesh (top).

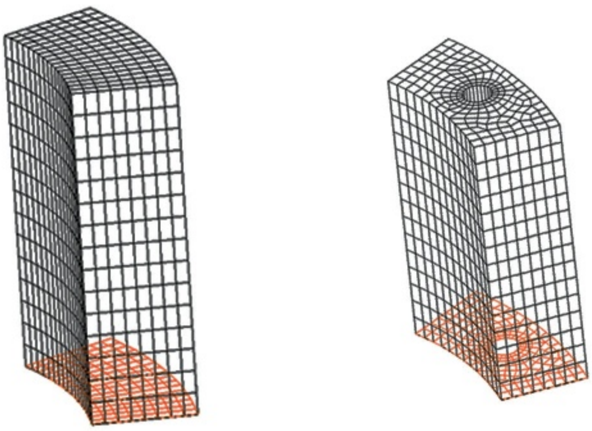

FIGURE 5: 2.5D meshing, resulting in hex elements.

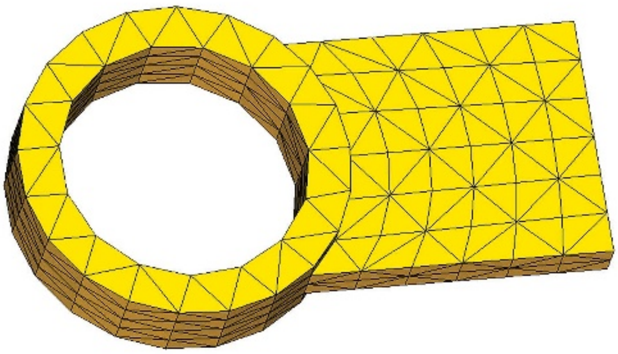

FIGURE 6: Global mesh control.

from the original 3D CAD can be labor intensive, however.

The mesh can be created directly without need for parent geometry; this gives greater flexibility in approach.

Quadrilateral (Quad) elements are preferred, and Triangular (Tria) elements avoided, as they have inferior accuracy. A four-sided geometric surface will always generate smooth "mapped" type meshes. An arbitrary five- or greater-sided surface requires an irregular, "Pac-Man-like" generation, by spawning elements from each side and then smoothing confrontation regions where sides meet. Fig. 4 shows surfaces and mesh of the two types.

Hex and Tet Elements

If a CAD component can be formed by extruding, sweeping or lofting in the CAD model, then we can use 2.5D meshing. A Quad mesh surface forms 3D Brick (Hex) elements by sweeping, or the solid can Hex mesh directly (see Fig. 5). The Hex elements are much more accurate than 3D Tetrahedral (Tet) elements. However, it is difficult to mesh other than in 2.5D or "brick-like" geometry. Many have tried to invent a pure Hex mesher—and so far, none have succeeded.

All 3D geometry can be Hex meshed, given sufficient time and patience. A Rubik's Cube is a good analogy. 3D meshing is a tradeoff—Hex for high accuracy, Tets for speed and ease of meshing. Warning: Do not use four-noded Tet elements. They give appalling results. Only use 10-noded Tets (mid-side nodes present).

Meshing Control

The basic mesh control is global mesh size, applied throughout the geometry (see Fig. 6). Most meshers will calculate this automatically. You can override, but be aware that a reduction from 6mm to 3mm implies eight times more elements.

Mesh control is improved by defining mesh density at vertices, along curves or on individual surfaces. Curvature-based meshing can be used, because curvature increases and mesh density increases (see Fig. 7).

This hierarchy of controls can be extremely effective, but requires practice. I recommend starting with simple shapes and experimenting with controls to see how they interact (see "Achieving Perfection," page 27).

Checking Mesh Quality

Mesh quality is checked via element geometry such as aspect ratio, value of internal angles, etc. Check documentation definitions, because they do vary across solvers. Contour plots and filters are useful in identifying poor elements. The solver should also give warnings of poor elements. I recommend taking simple geometry and introducing mesh distortions to explore this.

Simply put, badly shaped elements affect the accuracy of results. In an extreme case, a Tet element can be so distorted that

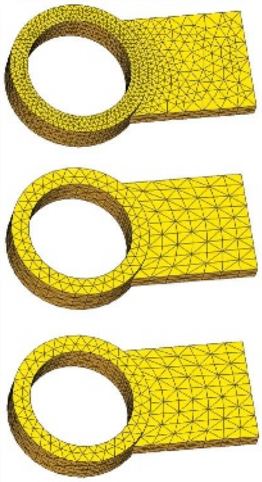

FIGURE 7: Curvature-based mesh refinement.

FIGURE 8: Poor CAD detail, resulting in slivers.

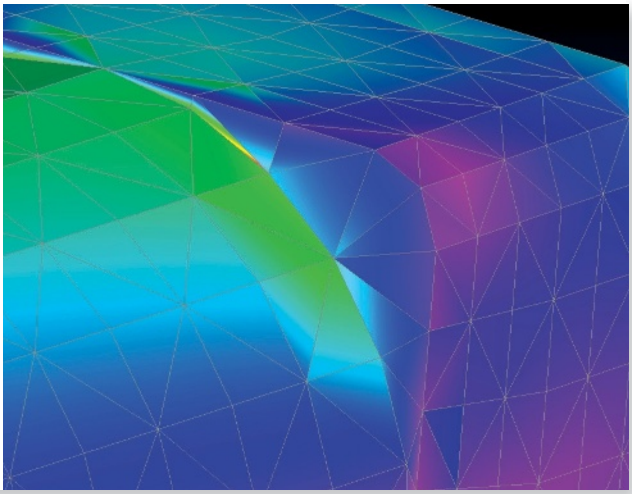

it will "blow up" numerically, causing solver failure. We need to trap this before the analysis starts.

One useful "catchall" check is the Jacobian value. This is based on the mathematics used to evaluate the element stiffness. A good limit value is element- and solver-dependent, thanks to varying mathematics and definitions. I recommend creating a regular mesh, then distorting the mesh and comparing values.

Bad Geometry

CAD geometry can have a variety of problems that directly affect meshing. Small slivers are common at the intersection of curved surfaces, for example (see Fig. 8). The mesh has high aspect ratio Tets in this region, and gives poor results. It is best to simplify or alter the geometry to eliminate the feature, or to add a properly defined fillet.

Other CAD features can result in poor meshing, including tolerance errors resulting in "cracks" in the geometry, negative volumes, etc. CAD geometry robustness may be poor because of interdependence of geometric features. Removing a hole or fillet may destroy the integrity of the CAD model.

While this is a big topic, in summary, I recommend a careful assessment of just how "mesh-ready" the CAD geometry is.

Element Stress Accuracy

FEA solves for displacements, not stresses. The displacements are continuous and smooth if the mesh is reasonably fine. The derived stresses are not continuous, and can jump widely in value at a common node if a poor mesh is present.

Fig. 9 shows a fillet with increasing element density, improving the stress distribution and accuracy. The postprocessor is prevented from artificially averaging stress values at nodes. Averaging hides the stress "jumps" between elements, and can reduce the apparent peak stress values. I recommend keep averaging switched off.

Fig. 10 shows stresses around a loaded plate with hole; the maximum stress value increases as the number of elements increases and converges to an accurate solution.

Switching on stress averaging gives smooth contour plots, but the peak stresses are averaged. This is what I refer to in training courses as the "shell game." Where is the right answer hiding?

Accuracy Levels

An accurate estimate of stress levels in areas of concern is needed, such as around holes, fillets, weld toes and similar stress raisers. FEA loads and constraints may be applied to point or line features, instead of being spread over a surface.

In practice, any load or support is distributed over a region. Using a point or line implies a finite force applied on an infinitely small area, giving an infinite stress—a stress singularity. Unrealistically sharp fillets and welds also cause this.

Fig. 8 shows a fillet de-featured or undefined in the CAD

Meshing Technique

1. Start with known meshing problem areas. Meshing is always iterative; don't be satisfied with the first attempt. The meshers are designed for fast deletion and remesh. Join up the difficult areas once satisfied.
2. Set a time limit. Time can run away when meshing—it is easy to get diverted by what can be quite a therapeutic pastime! I recommend setting a time limit for specific mesh regions. For example, allot 20 minutes to achieve a best compromise mesh distribution.
3. Focus on the big picture. Keep the overall strategy and element count in mind as you progress through the structure.
4. Do a coarse mesh run to get a feel for stress concentrations, then refine accordingly. Avoid a leap in mesh density, which risks tripping over the performance cliff edge.

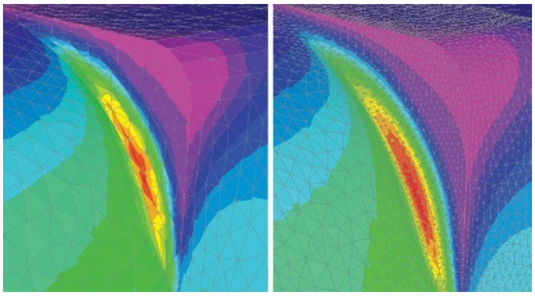

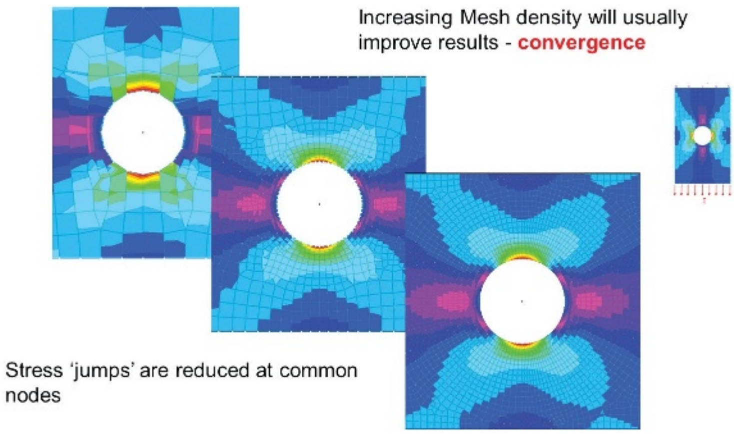

FIGS. 9 (ABOVE) AND 10 (RIGHT): **Improving stress distribution and accuracy.**

model. The mesh follows the right angle corner. We can't manufacture that sharp a corner, theoretically down to molecular level. Our FE model is "dumb" in that sense—we defined an infinitely sharp corner, and mathematically, that's what we get. Increasing the mesh density to improve the distribution of stress and get a "better" answer won't work; the stress just keeps going up. This poorly modeled region can result in an arbitrary stress value.

If we include the fillet, as in Fig. 9 (left image), then we get a finite stress value. But the mesh may be too coarse. The mesh jumps are big, giving low confidence in the values. Increasing the mesh density results in only a small increase in stress and a good distribution (see Fig. 9, right). The solution is nearly converged. A study like this is very important to show accuracy of results.

Just Hit the Meshing Button?

We need to put the elements where they will be most effective—at stress concentrations with high stress values. We use manual controls, or an adaptive mesher, by running trial analyses and refining the mesh.

Fig. 9 (right) shows a local fine mesh at the fillet in a global coarse mesh, giving 208,826 elements. Using a global fine mesh gives 3,806,030 elements. The solution in CPU time is proportional to element number squared, so it is a significant increase.

The solution must fit into available memory. If we overflow, the FE solver writes to virtual memory or disc with a very sharp increase in solution time. This is known as the "cliff edge" effect.

Achieving Perfection

The perfect element shape is:

- Hex: cube
- Quad: square
- Tet and Tria: equal angles

These shapes give the best numerical integration to derive the displacement and stress distributions in the elements.

The amount of output data—stresses and displacements—can reach terabytes. Industry anticipates a rapid increase in FEA data output to petabytes, or 1,000 terabytes. A lot of that is unnecessary, however, and can be avoided by good planning.

Everyone has finite resources, so aim to work within those limits. You can explore your resource limits by building dummy "block" models meshed to death. Compare model size against run time and data storage, and tune your FE model size. For example, if 5.8 million elements trips you out of core memory with a huge increase in runtime, aim for around 5.2 million elements.

Slicing and Dicing

Fig. 1 shows a solid with additional slicing on faces, and through depth to improve the mesh distribution. The extra operations result in a CAD model that is not an efficient definition of the component—although it is a more useful, "FE-friendly" CAD model where we exert better mesh control.

Loading is applied on one quarter of the end lug, so a surface must be created to define this (see Fig. 1).

In summary, if you have a structure that is mesh-intensive, it is going to take some planning. Performance tuning exercises help by giving you a feel for how many elements your system can handle, and accuracy tests will help you understand what the element density and permissible distortions are. DE

Tony Abbey *is a consultant analyst with his own company, FETraining. He also works as training manager for NAFEMS, responsible for developing and implementing training classes, including a wide range of e-learning classes. Send e-mail about this article to DE-Editors@deskeng.com.*

INFO → **FETraining:** FETraining.com

→ NAFEMS: NAFEMS.org/e-learning

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

Thoroughly *Modern* Metrology

Thanks to a host of ease-of-use and accuracy improvements, a new generation of 3D scanners and portable CMMs are improving productivity and reducing costs.

BY BETH STACKPOLE

The combination of ease-of-use improvements, higher precision accuracy, and a narrowing price/performance curve is advancing the cause of 3D scanning as a means of modernizing and bolstering productivity for both quality inspection and reverse engineering applications across a growing range of manufacturers.

The newest generation of 3D scanners and related quality inspection and metrology software promises to deliver clear benefits over mechanical measurement processes, including traditional coordinate measurement machines (CMMs), experts say. Specifically, smoother workflows via integration with 3D CAD systems, ongoing improvements to ease operation, and continually evolving scanning speeds are helping engineers improve efficiency and quality for first article inspection applications—all while shaving hours (and dollars) from product development processes. This is particularly true for reverse engineering objects with complex geometries.

"There is definitely a trend toward more 3D scanning," says David Olson, director of sales and marketing for Verisurf, a provider of a model-based definition metrology software suite that supports a range of metrology hardware, including laser trackers and scanners as well as stationary and portable CMMs. "You can capture many more points in a shorter amount of time, which can be very useful when you have organic shapes and ergonomic forms, like a pistol grip or an interface between a prosthetic device and a human body."

Consumer interest in 3D technology, and specifically, a rising comfort level with products like Microsoft's Kinect gaming interface, which uses light to measure and display 3D virtual environments, is also drawing attention to the category. Metrology experts say that while Kinect and other consumer-driven 3D scanning technologies don't yet provide anywhere near the level of accuracy and precision required by industrial applications, they are increasing exposure to 3D scanning—thus helping to defuse change management issues that have historically impeded widespread adoption.

"The consumerization of 3D is making people much more willing to

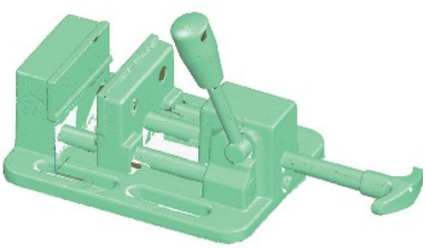

Scan data of the vise seen directly in Geomagic Spark's capture window.

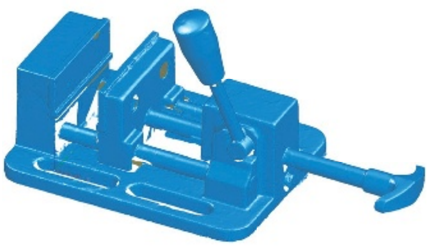

Automated creation of polygon mesh data from the scan data, using Geomagic's tools.

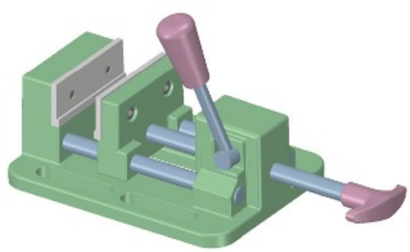

Vice assembly as a solid model, created directly from the mesh data into each related part.

employ it in business," notes Ralph Rio, research director with ARC Advisory Group. "Kinect is improving the knowledge and willingness of consumers to adopt 3D, and definitely helping out in the adoption of laser scanning."

A Contrast in Technologies

Despite what many claim is the faster evolution of 3D scanners, CMMs are still widely entrenched, and are also being made over to handle the demands of modern metrology applications. Traditional CMMs are three-axis machines built on a granite slab for stability. They employ a fixed probe-like device to take a precise, specific set of point-by-point, X, Y, Z measurements.

CMMs' highly trusted approach has long been lauded for its accuracy and repeatability, especially for prismatic and primitive objects such as rectangles, cylinders and cones, experts say. On the down side, CMMs require extensive programming, which demands a specific set of skills and certifications. They are not well suited for measuring complex, organic shapes, and the general consensus is that it can be difficult to communicate results beyond the pool of CMM experts because the output is too cryptic for anyone not versed in metrology science.

The current crop of CMMs is being updated with designs intended to give users more flexibility. Specifically, many CMMs can now accommodate either discrete point probes or 3D laser scanning arms—and the technology is being offered in an array of new portable form factors, including those that can be used directly on the shop floor (as opposed to requiring use only in a quality assurance laboratory).

3D laser scanners, on the other hand, take a wholly different approach. While there are variations on the technology, most 3D laser scanners work by measuring thousands of points per second to calculate the X, Y, Z coordinates for each point, then creating a digital representation of the surface of an object, typically known as a point cloud. 3D scanning's point cloud approach is garnering a lot of interest as a means of delivering a more complete picture of an entire surface of an object, metrology experts say.

In addition, its streamlined integration with 3D CAD software is also viewed as an asset—especially when trying to create a seamless workflow for computer-aided first article inspection and reverse engineering applications. Applications like Geomagic's Studio and Rapidform XOR transform 3D scan data into highly accurate polygon and native CAD models, which can then be integrated into popular MCAD programs for further design and modeling work.

"CMMs are good at telling me the size and location of holes, but if something is going on in between

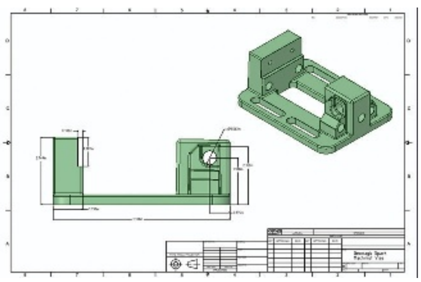

This 2D detail sheet was created from the solid assembly for dimensioning and annotation using Geomagic Spark.

The vise assembly, rendered using the KeyShot third-party add-in available alongside Geomagic Spark.

Geomagic and SpaceClaim Spark 3D Scan Data Designs

The combination of 3D scan data and 3D CAD modeling has long been the recipe for successful reverse engineering; however, the capabilities have traditionally been serviced by separate—and oftentimes, non-integrated—technologies.

Geomagic, a major player in the quality inspection and metrology software arena, is seeking to change that equation with Spark, the newest addition to its product line. Unlike traditional CAD programs, which have modeling capabilities, but lack the tools to process 3D scan data into usable 3D geometry, Geomagic Spark combines a live 3D scanning interface with robust 3D point and mesh editing capabilities, along with a full-blown CAD modeling package.

Specifically, Geomagic Spark integrates Geomagic's scanning technology with SpaceClaim, a popular 3D direct modeling platform. The goal is to allow users to create accurate and manufacturable solid models and assemblies using a combination of 3D scan data and 3D CAD data in the same application. The two companies, which have been collaborating since the launch of the SpaceClaim integration in the 2012 release of Geomagic Studio, saw opportunity in partnering around a fully integrated application that would alter the way users design with 3D scan data.

"It's not intimidating the first time you use it," says Kevin Scofield, senior product manager at Geomagic. "It has simple controls. You don't get lost. We were able to add a new ribbon inside of SpaceClaim that has new 'Capture' tools."

Here's how the integrated package works: Designers scan directly into Geomagic Spark, or load an existing point cloud or polygon mesh. They then choose from a range of automated tools to convert and edit data into a polygon mesh. Once the mesh is created, Geomagic Spark's solid modeling tools—based on SpaceClaim's smart and predictive user interface—allow users to create solid geometry from the mesh, and compare the solid model to the mesh to identify areas of deviation.

Scofield says there were a number of reasons why Geomagic chose to partner with SpaceClaim, including similar company cultures, size and an East Coast location.

"What really sealed the deal was, once they were able to load meshes natively—that was the 'aha' moment," he says. "Once we saw that, we knew if we were able to add some of our tools, we'd have a winning combination. Other tools we looked at can load meshes, but you can't interact with it as much."

The ShapeGrabber Ai810 automated 3D laser scanner quickly scans parts that vary in size and shape, enabling it to support automated inspection applications in minutes vs. hours. ***Image courtesy of ShapeGrabber.***

the points that I've measured, the CMM can't tell me," notes Dan Perreault, president of NeoMatrix Technologies Inc., an engineering services provider specializing in 3D scanning technology for reverse engineering and computer-aided inspection applications. "If a part passes inspection on a hole, it still might not work because a CMM only gathers a few data points—it doesn't give you a complete picture."

Consider a quality inspection application for the inside of a car door. While a CMM probe can serve up an accurate measurement of the hinge location or other holes where things mount, it doesn't have sufficient data to pick up on imperfections in the sheet metal that might impede the door's fit.

"With a white light scanner, you can get 100% coverage, with millions of data points that show hole locations and the shape of the rest of the door," Perreault explains. "You can also take the data from the scanner, overlay it with the original 3D CAD design, and see glaring red spots to indicate out-of-tolerance conditions and deviations."

3D scanning can also be less labor-intensive than traditional CMMs, in terms of the actual process of recording measurements. That's according to ARC's Rio, who is projecting the worldwide market for 3D laser scanners to grow at an 8.8% annual growth rate (CAGR) from 2011 to 2016.

With a CMM, an operator has to move the probe from location to location to physically capture each point, or at least be knowledgeable enough to program the CMM to

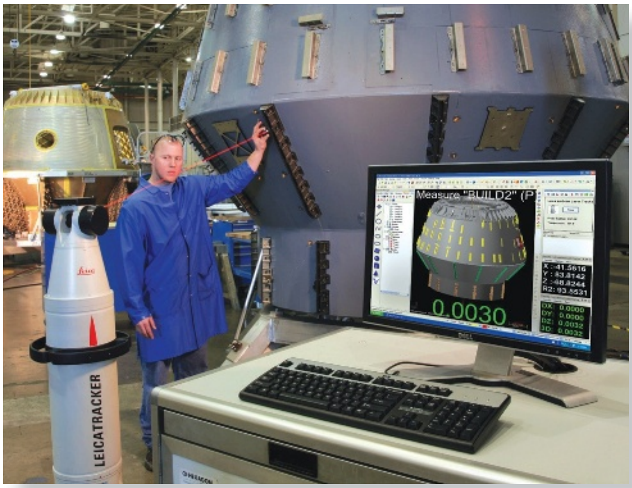

Computer monitor shows Verisurf Software controlling the Laser Tracker portable CMM used to inspect the SpaceX Dragon capsule. The software displays the in-tolerance deviation between the capsule and its respective CAD model in real-time to determine whether the capsule is assembled within specifications. *Image courtesy of SpaceX.*

Optical CMM scanners like Creaform's arm-free, handheld MetraSCAN tackle metrology or large-scale reverse engineering applications by promising to extend measurement volumes without loss of accuracy. *Image courtesy of Creaform.*

automate data capture. Not so with a 3D laser scanner, Rio says. It not only scans a large area—the hood of a car, for example—but can do so without having to be as accurate as to where you physically place the laser.

"The ability to read tens of thousands of points in a few seconds is a huge change in labor," he explains. "You can go from a couple of hours to a few minutes for inspection."

Simplicity and Speed

The various ease-of-use and automation features associated with 3D scanning technology are one of the primary drivers for increased adoption. Many 3D scanners are now relatively easy to operate without any programming, with some even going as far as to deliver fully automated quality inspection environments. Consider ShapeGrabber's Ai810 3D laser scanner, for example, which is designed to quickly scan parts that vary in size and shape, and perform automated inspection duties in minutes vs. hours, notes Pierre Aubrey, ShapeGrabber's CEO.

In addition to the right scanner, 3D inspection and metrology software like Geomagic Qualify is another crucial element in automating quality inspection processes, according to Bill Greene, CEO of Level 3 Inspection LLC, a provider of white light scanning and computer-aided inspection services. Level 3, which caters to manufacturers of precision products and tooling, has wholeheartedly embraced 3D scanners and Geomagic's 3D software because the combination transforms computer-aided inspection into an automated, batch process.

"Automation takes human error out of the equation—

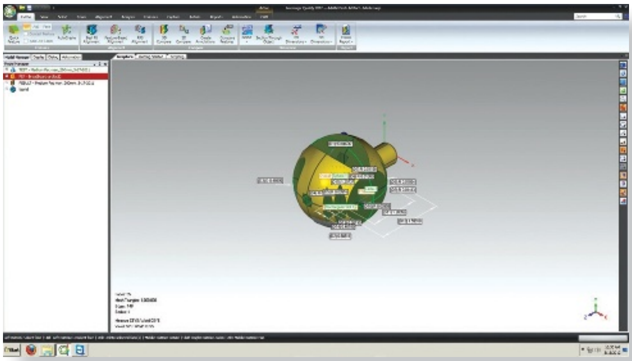

Geomagic's Qualify metrology and 3D inspection software serves as the centerpiece for Level 3 inspection's computer-aided inspection services, particularly its scanning-to-measurement process. *Image courtesy of Level 3 Inspection.*

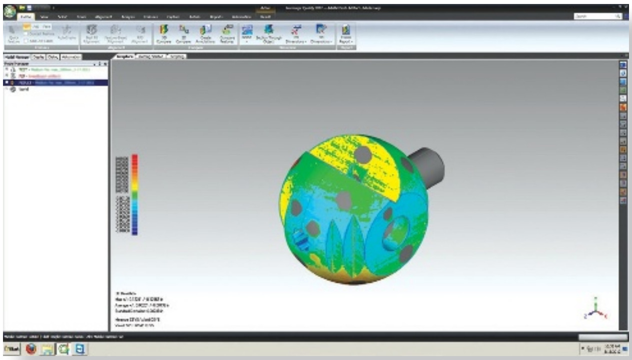

Level 3 Inspection creates a Geomagic Qualify color plot to showcase measurement deviations of the scan data from the reference data. *Image courtesy of Level 3 Inspection.*

and it's fast," Greene says. "You can do computer-aided inspections in 12 to 15 minutes that use to take a week or two."

For reverse engineering applications, portability and ease-of-operation are crucial considerations. That's because it's often design engineers or researchers doing the scanning—and they are often not certified metrologists, notes Daniel Brown, product manager for 3D scanners for Creaform, a provider of portable 3D measurement solutions, including portable 3D scanners and portable CMMs.

"For the metrologist who does quality control, it's a job—and they are less concerned about time than quality and the reliability of the device," Brown says. "For the person doing reverse engineering, they want to extract information and use it for something else, so the issue is around speed. They want a simpler, faster, more portable device."

For many shops, however, it's not an either-or choice, but rather a combination of 3D scanning and CMM technology. That's certainly the case at Schneider Electric, a global manufacturer of energy transmission and automation products, which employs a two-stage process for its quality inspection processes, according to Rus Emerick, the company's global process owner for 3D imaging. Schneider Electric employs white-light 3D laser from Laser Design initially to determine whether a part is a geometric match to what was initially designed in a 3D CAD model, Emerick explains. Once that determination is made, the process is turned over to a CMM, which automatically checks the production lot of parts to ensure quality.

It's this two-stage process that sets Schneider Electric's quality inspection operation apart, Emerick maintains: "Most companies still go directly to measuring things; we have become more efficient by eliminating the need to measure parts that aren't correct. It's more efficient to use 3D scanning and Geomagic [software] to extract specific measurements out of a data set, and use CMMs for repetitive linear measurement." **DE**

Beth Stackpole *is a contributing editor to* DE. *You can reach her at beth@deskeng.com.*

INFO → **ARC Advisory Group:** ARCWeb.com

→ **Creaform:** Creaform.com

→ **Geomagic:** Geomagic.com

→ **Laser Design:** 3DScanCo.com

→ **Level3 Inspection LLC:** Level3Inspection.com

→ **Microsoft:** Microsoft.com/en-us/kinectforwindows

→ **NeoMatrix Technologies:** 3DScanningServices.net

→ **Rapidform:** Rapidform.com

→ **Schneider Electric:** Schneider-Electric.com

→ **ShapeGrabber:** ShapeGrabber.com

→ **SpaceClaim Corp.:** SpaceClaim.com

→ **Verisurf:** Verisurf.com

VIDEO LINKS

→ Creaform's MetraSCAN handheld CMM: YouTube.com/watch?v=P3EDhSGszGY&feature=plcp

→ Laser Design's Auto Gage 3D Scanning System: YouTube.com/watch?v=ZKt9RRLjIak

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

Tools of the Month

1 Tektronix Introduces TBS1000 Oscilloscopes

Tektronix's (tek.com) TBS1000 Series of entry-level oscilloscopes are targeted at general-purpose electronics test applications for electrical engineers, educators and hobbyists. The new oscilloscopes deliver 1 GS/s or 500 MS/s sample rates and 3% DC gain accuracy down to 2 mV. A new waveform data-logging feature allows users to save up to eight hours of triggered waveforms to a USB flash drive. Waveform limit testing allows users to monitor signal changes and output pass/fail results. The TBS1000 is based on the same proprietary technology as the TDS2000 Series, and includes the same triggering, automatic measurements, USB connectivity, and FFT capability.

Geomagic Unveils New Suite of Solutions

In addition to Geomagic Spark (see page 30), Geomagic (geomagic.com) has released 2013 versions of its Freeform, Qualify and Qualify Probe solutions. Freeform 2013 3D modeling software includes subdivision modeling, and allows designers to combine four different modeling representations (voxels, NURBS, subdivision surfaces and polygons) in the same model. It is integrated with the company's Sensable Phantom haptic 3D input devices. The 2013 releases of Geomagic Qualify and Geomagic Qualify Probe deliver improved automation, GD&T and hard-probing functions.

2 Omega Introduces DC Current Data Logger

The Omega (omega.com) OM-CP-PROCESS101A is a low-cost logger with a 10-year battery life, 4Hz reading rate, multiple start/stop function, ultra high speed download rate, 1 million reading storage capacity, memory wrap, battery life indicator, optional password protection and programmable high and low alarms. Two current input ranges are available: 20mA and 160mA. There are free firmware upgrades for the life of the product so that data loggers already deployed in the field can support new technological developments.

Mantis Vision Launches F5-Short Range 3D Imager

Mantis Vision's (mantis-vision.com) F5-Short Range 3D imager features a wide depth of field for small-object, close-range capture. According to the company, the handheld 3D scanner combines an ergonomic design with speed and ease-of-use to capture small, intricately detailed objects. Designed to work alone, or with Mantis Vision's F5 imager, the F5-SR can complete hard-to-reach detail in larger areas, according to the company.

FARO Unveils 3D App Center for Focus3D Users

FARO Technologies (faro.com) launched the FARO 3D App Center, the company's new online exchange for Focus3D software applications. Through the 3D App Center, Focus3D users and software developers can download and share purpose-built apps to meet specific scanning challenges, the company says. According to FARO, the 3D App Center reduces costs and simplifies workflows by replacing expensive software packages that perform tasks users may never need. Instead, customers can access and download FARO software apps built specifically to resolve their particular scanning challenges, such as surface or volume calculations.

VPG Unveils New Data Acquisition Scanner

Vishay Precision Group (vishay-pg.com) announced that its Micro-Measurements brand has introduced a new data acquisition system featuring a scanner with eight software-configurable input channels with RJ-45 connectors. The System 8000 StrainSmart Data Acquisition System can accept signals from strain gages or strain-gage-based transducers, thermocouples, or high-level voltage sensors. Strain gage channels support full-, half-, or quarter-bridge configurations. The System 8000 is intended for static and dynamic test and measurement applications, the company says. System scanners may operate independently, or up to 16 scanners can be used concurrently for a maximum of 128 channels.

3 Test and Measure Quality

Icona Solutions' (icona-solutions.com) aesthetica version 4 'virtualizes' the final manufactured product by simulating and visualizing the effects that dimensional variation and part flexibility will have on its aesthetics. The company says it encourages a continuous, cross-discipline focus on perceived quality, from concept design and styling, through engineering development to manufacturing and final assembly. New features include an interactive 3D rendering system that supports more photo-realistic visualization through the use of high dynamic range (HDR) environments, global illumination and shadows; an extended material library; a more accurate measurement system; and a new graphical user interface. It is designed to enable those who are primarily concerned with the visual quality of a product and those who are involved in the definition and control of manufacturing tolerances to engage in meaningful discussions about the perceived quality of a product under development. DE

Scope Out the Latest Technologies

Oscilloscope innovations lead to faster, more accurate testing.

BY JAMIE J. GOOCH

Oscilloscopes have come a long way since The Heath Co. packaged a bunch of surplus World War II electronics together in 1947 and sold it as the O1 oscilloscope kit for $50, making it the first scope to gain wide acceptance. Since then, manufacturers have responded to the engineer's need to test different aspects of increasingly complex products.

Tektronix says its THS3000 series of portable oscilloscopes have a seven-hour battery life.

"If you look at an engineer's test bench, the scope is the centerpiece of their bench," says Richard Markley, product planner for Agilent's InfiniiVision oscilloscopes. "The oscilloscope is what they use most often, so we tried to focus on making a phenomenal scope. But engineers want to use them for other things."

Those "other things," such as mixed signal analysis, digital voltmeters, logic and protocol analyzers, math functions, and measurements, tended to slow down the primary function of the scope: processing and displaying signals. But just like the devices they test, oscilloscopes are getting faster, smaller and more advanced. Modern oscilloscopes are judged on the speed of their sample and update rates, as well as their ease of use and flexibility in offering the number of analog and digital channels, bandwidths and advanced functions that make them multi-use pieces of test equipment.

Fast and Easy Functionality

The speed at which waveforms are sampled can be the difference between seeing a rare anomaly and missing it while the data is being processed. When a signal is sampled, an oscilloscope needs time to preprocess it, write it to memory, and display it. While it's doing that in real time, it's not acquiring a signal.

For example, Aglilent says its latest scope, the InfiniiVision 4000 X-Series, samples 1 million waveforms per second (WFM/s), even with digital channels turned on. That brings it in line with the R&S RTO line of oscilloscopes from Rhode & Schwarz, which also hits 1 million WFM/s.

With all of an oscilloscope's potential for advanced, specialized functions, manufacturers are also challenged by keeping the user interface simple and efficient. With the popularity of smartphones and tablet devices, it was a small leap to imagine touchscreen user interfaces for oscilloscopes.

"With where the touchscreen consumer devices are going, there's clearly an opportunity to capitalize on that, and make things easier for our customers who are used to that," says Markley.

To test the limits of that opportunity, Agilent spent more than two years researching and developing the InfiniiVision 4000 X-Series. During that time, the company put its new capacitive touchscreen interface in front of end users. The new touchscreen interface is more akin to a modern tablet than previous resistive (pressure-based) touchscreen technology, in that it allows fast zooming and scrolling.

"What we found was that if customers owned a tablet or a smartphone, they usually loved the touchscreen interface," says Markley. "About half of the customers love it—and the other half's first question was 'Where is the off button for the touchscreen?' It was not so much a break along different ages like you might expect. It really depended on whether they owned a touchscreen product."

Released in November, the 4000 X-Series includes a 12-in. touchscreen user interface that allows users to draw a box on the screen around an anomaly and trigger on it, and an onscreen touchpad that can be used instead of the standard knobs. However, those knobs are still there for users who aren't comfortable with the touchscreen interface. The 4000 X-Series also accepts mouse-based input.

Pico Technology's PicoScope is an example of a PC-based oscilloscope.

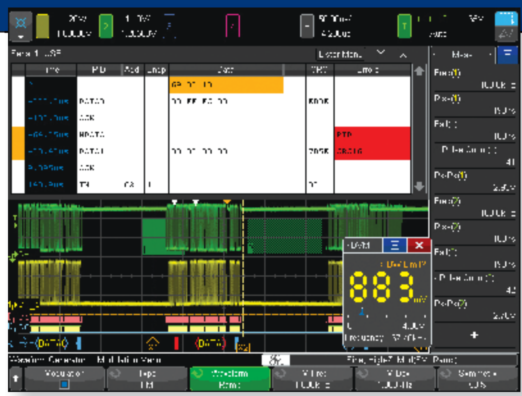

The Agilent InfiniiVision 4000 X-Series oscilloscope features a tablet-like touchscreen.

Flexible and Expandable Options

In addition to the user interface options, Markley says the InfiniiVision 4000 X-Series was designed to be flexible and upgradeable. The series includes 200-, 350- and 500-MHz, and 1- and 1.5-GHz models. It also includes the capabilities of five instruments on one unit: oscilloscope, digital channels, protocol analysis, digital voltmeter and a dual-channel WaveGen function/arbitrary waveform generator.

"We've been hearing a lot about how engineers' needs are changing so fast," Markley says. "Project needs and requirements for test equipment are moving quickly. Our customers were getting frustrated with having to buy a new scope every few years to keep up with those requirements."

To counter that frustration, Markley says most features of the 4000 X-series are upgradeable, except for having to choose between either two or four analog channels at the time of purchase.

A Diverse Market

Other manufacturers have taken different approaches when it comes to offering flexibility and upgradeability in an oscilloscope by tying it and its software to a workstation.

For instance, Pico Technology offers PC-based oscilloscope software and hardware that it says can turn your workstation into a test and measurement lab. Its PicoScope 6 software, which is provided free with purchase of one of the company's PicoScope USB products, allows the company to provide the features of an oscilloscope, spectrum analyzer and signal/waveform generator, all of which can be updated via free downloads.

National Instruments calls its PC-based oscilloscopes digitizers, because they can be defined with software to perform the functions of oscilloscopes or other digital measurement instruments. They can be programmed using NI's free NI-SCOPE driver software, which comes with more than 50 prewritten example programs. The NI-SCOPE Soft Front Panel provides an interface that looks like a traditional, dedicated oscilloscope.

While PC-based oscilloscopes can be plugged into a laptop and taken out into the field, they're not the only option for engineers who need to analyze on the go. A number of manufacturers have seen the need for smaller form factor oscilloscopes in today's mobile workplace.

Tektronix's answer to the portability problem is its THS3000 series of oscilloscopes. With a weight of 4.8 lbs., seven hours of battery life and four isolated channels, the THS3000 series is built for use in demanding conditions, according to the company. Models are available with either 100MHz and 2.5 GS/s maximum sampling rate, or 200MHz and 5 GS/s sample rate performance. The 10,000-point record length on all four channels enables the capture of extra waveform details at high sample rates. Scopes in the series include 21 automated measurements, built-in limit testing, and waveform math and data logging capabilities. They also include a USB port for adding additional flash storage or downloading data and screenshots to a PC.

Extech Instruments' MS6000 series of portable digital oscilloscopes have 5.6-in. color LCDs. Each of the three scopes in the series—the 60MHz MS6060, the 100MHz MS6100 and the 200MHz MS6200—feature 23 automatic measurements and a waveform recorder to capture/replay up to 1,000 waveform screens to their installed 2GB microSD cards. According to the company, the series has a battery life of three hours.

Yet another take on expanding oscilloscope functionality can be seen in Yokogawa's ScopeCorders. These multi-channel test and measuring solutions combine oscilloscopes with paper chart recorders. In addition to displaying signals on screen, the ScopeCorders can record to paper and/or memory for 30 days or more.

We've just exposed the tip of the iceberg when it comes to advancements in oscilloscopes. When comparing oscilloscopes, be sure to do your research by visiting the manufacturers' websites, and check out distributors, such as Saelig and Mouser Electronics, for a variety of oscilloscope solutions. DE

Jamie Gooch *is* DE's *managing editor. Contact him via jgooch@deskeng.com.*

INFO → **Agilent:** Agilent.com

→ **Extech Instruments:** Extech.com

→ **Mouser Electronics:** Mouser.com

→ **National Instruments:** NI.com

→ **Pico Technology:** PicoTech.com

→ **Rhode & Schwarz:** Rohde-Schwarz.com

→ **Saelig Co. Inc.:** Saelig.com

→ **Tektronix:** Tek.com

→ **Yokogawa Electric Corp.:** Yokogawa.com

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

Hardware-Software Design Optimization

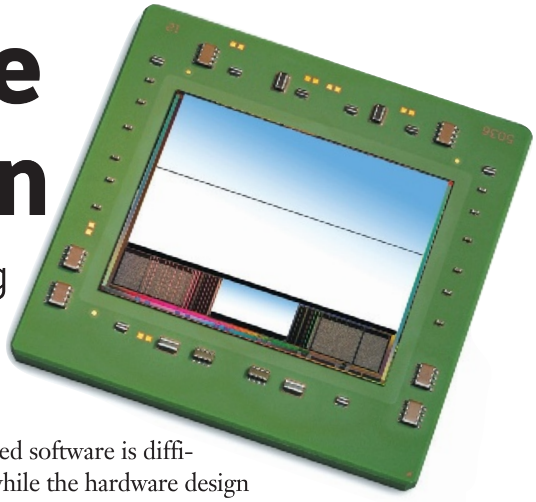

New tools for electronic system prototyping are speeding design testing.

BY PETER VARHOL

Engineers designing electronic systems likely have experience working with field-programmable gate arrays (FPGAs), programmable hardware chips that can be overlaid with logic to create electronic components—and even entire systems. These chips provide the ability to quickly configure and customize electronic devices after manufacturing, in the field.

In many cases, FPGAs and similar logic chips are also used to prototype complete electronic systems that are used in system-on-a-chip designs that deliver full processing and input-output functionality for applications like mobile phones and automotive electronics. These prototypes are then cast into application-specific integrated circuits (ASICs) for performance and manufacturability, and used in complex hardware-software systems.

But engineers can have a difficult time working with FPGAs, and deeply embedded systems in general. FPGAs and other logic devices have hundreds of thousands, or even as many as tens of millions of logic gates, and programming them using traditional hardware definition languages can be a challenge. Most designers prefer working with pre-defined logic components, referred to as intellectual property, or IP. These are hardware description language (HDL) programs that define a particular function, such as a processor, I/O interface like USB, or even graphics.

Further, these systems are almost always customized with software. An automotive braking system, for example, uses a processor, I/O for data inputs and instructions, and storage. But it is also coded with explicit instructions on how to interpret and respond to data inputs, and how to handle data and processing errors.

This embedded software is difficult to develop while the hardware design and implementation are still in flux. Interfaces, timing and even performance aren't usually fully defined until the hardware system is designed and the components selected. That puts software development behind the hardware, typically at the very end of the project. Putting software last often has undesirable consequences, such as poor performance, bad design, incomplete or missing features, and poor testing.

Building Virtual Prototypes in Logic Devices

In prototyping, electronic designers partition the FPGA logic chips and position the IP logic onto the chip, then design interfaces so that they work together in providing the needed functionality. While this may not be the final design, work on the software can begin as soon as there is an initial prototype of the hardware on the FPGA. The resulting circuit is then tested for functionality in the larger system, or within a testing harness.

Once hardware is finalized, it can be placed on the FPGA and the software can be completed. The device can be used as-is, if the end product is custom or for a small manufacturing run. For large runs, the FPGA can be produced as an ASIC or a system on a chip (SOC), and the costs spread over a larger product base.

This prototyping process sounds straightforward, but it often isn't. IP blocks can be difficult to work with, especially if multiple blocks are being used on a single chip. Often, the position, layout and interfaces of these blocks can influence both the capacity and

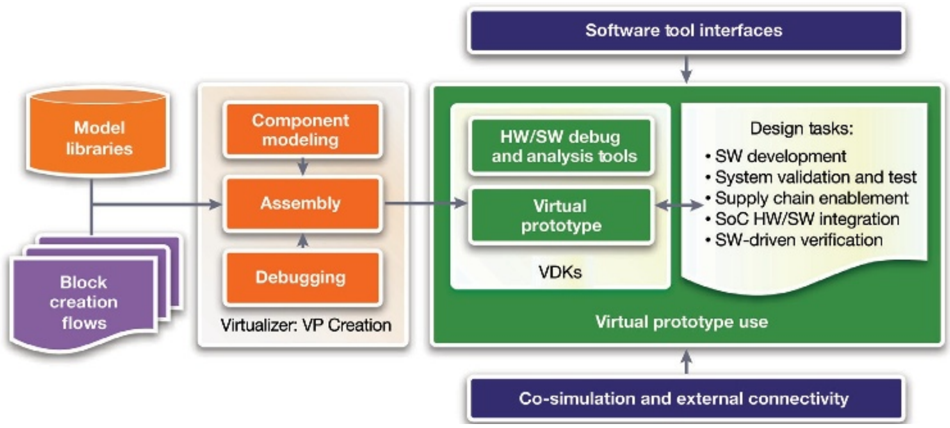

A typical design flow for embedded software optimization of electronic hardware-software systems, using TLM Central system models and the Virtual Development Kit.

performance of the resulting circuit. Sloppy design can easily result in not meeting product specs or performance. As embedded electronic systems become more complex, the costs of bad design and placement become more significant.

Vendors such as Xilinx and Altera make the FPGAs and other programmable logic devices that are used in prototyping these embedded software systems. The hardware description languages are standard across the industry. The competitive advantage in creating prototypes comes in the placement and mapping software that determines the optimal places for IP blocks on a chip. The best approaches to prototyping can save significant time and chip expenses over manual methods.

Hardware-Software Co-development

An initiative by the programmable logic industry is making it possible to do more embedded software development as a part of the hardware prototyping process. This initiative, called TLM Central, consists of a portal, a set of high-level descriptions and models of standard hardware components in the SystemC language, and comprehensive information on how to create and manage virtual prototypes.

These virtual prototyping systems are integrated with an intelligent prototyping software environment that enables faster partitioning, and automates the creation and debugging of prototypes for a range of designs—from individual IP blocks and processor sub-systems to complete SOC, easing the path from concept to operational prototype. Engineers can use the on-chip virtual prototype as a platform for loading an operating system, coding control and management software unique to the product.

The important part of the process is that the FPGA-based prototype is used for software development while the hardware is still being defined. The TLM Central solution includes a set of freely available software development components integrated into the engineer's software development environment.

According to Tom De Schutter, senior product marketing manager, System Level Solutions at founding member Synopsys, the more engineers that use TLM Central, the more valuable it becomes: "With over 1,000 hardware component models, engineers can build virtual simulations of a wide variety of hardware-based embedded systems in many different industries."

The accompanying software environment, called a Virtual Development Kit, incorporates popular software debuggers, system control mechanisms, and debugging information. To work with a virtual prototype, engineers install and start using—there's no need to wait for a hardware prototype to develop software.

Possibly the biggest advantage in this approach is a reduction of project risk. Any major issues with the software can be worked out more easily, either through redesigning interfaces or substituting other hardware components.

Engineers engaged in complex hardware-software system design in mobile communication, automotive, aerospace and similar industries are often faced with a rush to complete software development and integration once the hardware is in place. Using a virtual prototype of an FPGA and TLM Central-based system models, designers can begin work on the software systems before the hardware components are fully defined. Once final silicon is ready, either on an FPGA or ASIC, product designers can finalize the software much more quickly. Most important, there is no need to wait until final silicon to begin software development.

It's not a panacea, of course. Any software written before final silicon will have to be tested—and likely undergo some changes once it is able to run on actual hardware. And while the two can be developed more or less in parallel, the software will always take longer. But using a virtual prototype can reduce risk at the same time it accelerates time to market for an optimized design. **DE**

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 → **Altera:** Altera.com

→ **Synopsys:** Synopsys.com

→ **TLM Central:** TLMCentral.com

→ **Xilinx:** Xilinx.com

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

Power to the Engineer

SC12 showed how engineering computing is moving to the next performance level.

BY PETER VARHOL

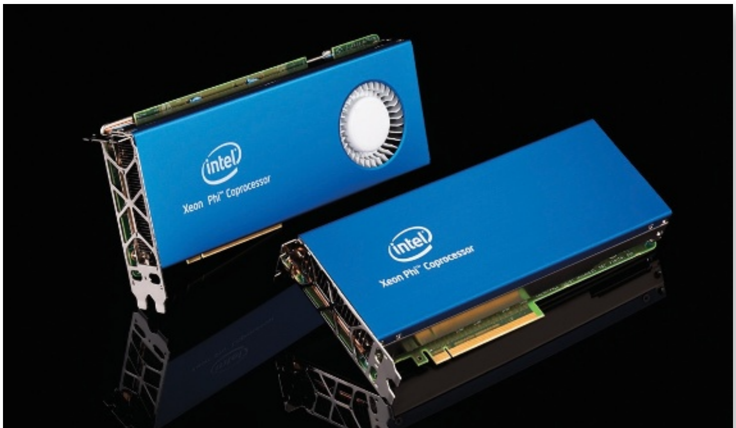

Intel made a splash at SC12 with the introduction of its Xeon Phi.

The Supercomputing 2012 conference featured computers, processors and other subsystems that are placing the fastest computing power within the reach of more engineering professionals. With high-powered alternatives in abundance, and rapidly shrinking prices, it won't be long before virtually all engineers either have the equivalent of a supercomputer at their desks, or in their laptop bags.

Supercomputing 2012 was also about other revolutionary changes in engineering. Cloud computing alternatives abounded. Some were simply hosting solutions, with a combination of CPUs and GPUs, for those who are able to upload an entire virtual machine and execute independently. Others, such as Microsoft's Azure, offer a more comprehensive execution environment—with specific language and system frameworks, and sometimes even specific engineering applications.

Mobile computing is also gaining a toehold in high-performance computing (HPC)—as a way to not only access computing jobs in the cloud, but to make designs more mobile and sharable. NVIDIA's Tegra processor adds both processing and graphics power, across four cores, for mobile devices.

On the software side, the focus was on getting the most out of parallel computing. Domain languages such as MATLAB, Mathematica and Maple are enabling engineers to run specialized computations on multiple cores, and in the cloud.

Big Announcements and Trends

Perhaps the most significant announcement of the week was the Intel Xeon Phi coprocessor. The Intel Xeon Phi is based on the vendor's many integrated cores (MIC) architecture, which uses Intel's industry-standard architecture and instruction set across multiple high-performance cores. The principle advantage of Xeon Phi and MIC is the ability to run existing software without recompile or other modification.

Many of the system and board vendors demonstrated systems that included a Xeon Phi slot, making them ready to go as soon as the coprocessor was available.

But the Xeon Phi wasn't the only high-performance processor on display. AMD is also in the game, and with a unique approach. During the conference, AMD launched the AMD FirePro S10000 server graphics card, designed for a combination of HPC workloads and graphics-intensive applications. Unlike Intel and NVIDIA, which have separate alternatives for computation and graphics, AMD combines graphics and computational GPU cores onto a single die. This provides 5.91 teraflops of peak single-precision and 1.48 teraflops of double-precision in floating-point calculations. With two GPUs in one dual-slot card, the AMD FirePro S10000 enables a combined compute-rendering GPU solution while increasing overall processing performance.

Hardware systems are undergoing a significant change to accommodate faster parallel computing processors. One of the most significant architectural bottlenecks in computing is the single, relatively narrow bus between the processor and main memory, and between main memory and secondary storage. Fusion-io offered a partial solution to that bottleneck with a flash memory storage device that was software-configured to act as either storage or an extension of main memory, with separate buses into the processor space. The advantage here is that program instructions and data can be stored on these flash memory devices, which then have direct access to main memory and processor space. The single bus isn't such a bottleneck any more.

Dell also demonstrated an innovative approach to managing engineering workstations in a secure yet flexible environment. Using technology from its recent Wyse acquisition, the vendor showed a display and keyboard remotely connected to a system located in the data center, or perhaps a vault. The workstation and I/O were connected by a box that transferred a high-speed protocol, passing graphics and keyboard data in real time across Gigabit Ethernet connections on the server and box. The end result is a separation among monitor, keyboard and physical computer—but with no lag in performance.

Why would an engineering group want to do this? One reason is security; the actual servers can be in a locked server room or vault, physically protecting designs on the system. But it also provides a way for engineers to work remotely, while accessing the applications and files on their workstation.

Software Still Makes an Engineering Solution Possible

The $64,000 question continues to be how effectively engineering software can take advantage of multiple cores to bring the full processing power to bear on a problem or computation. Many engineering vendors are rewriting their commercial software packages to parallelize computations, but this can be an expensive and time-consuming process.

Some languages, such as MATLAB and Maple, offer a limited ability to assign jobs to multiple processors and cores. But a lot of more general-purpose software may never be ported to execute in parallel. That's where Advanced Cluster Systems comes in. This vendor provides a software solution called SET that enables vendors and engineering groups with proprietary source code to easily parallelize that code. In some cases, if the application is constructed appropriately, source code may not even be required.

From talking to dozens of exhibitors and attendees at the conference, it was clear that engineers are using these advances in hardware and software to do more simulation of design components and even entire designs. The ongoing story was one of refining and optimizing simulations, and perhaps building one or two physical prototypes before going to manufacture. **DE**

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 → **Advanced Cluster Systems:** AdvancedClusterSystems.com

→ **Advanced Micro Devices:** AMD.com

→ **Dell:** Dell.com

→ **Fusion-io:** FusionIO.com

→ **Intel:** Intel.com

→ **Maplesoft:** Maplesoft.com

→ **MathWorks:** MathWorks.com/products/matlab

→ **Microsoft:** WindowsAzure.com

→ **NVIDIA:** NVIDIA.com/object/tegra.html

→ **Supercomputing 2012:** SC12.supercomputing.org

→ **Wolfram Research:** Wolfram.com/mathematica

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

AMD's Latest FirePro Boards

A new GPU powers the FirePro W-series.

BY DAVID COHN

We got our first look at AMD's four new FirePro graphics boards at a technology day held at the company's headquarters in Sunnyvale, CA, and brought samples back to review—our first comparison of AMD FirePro boards in more than a year (see *DE*, August 2011).

While AMD has long been known as a competitor to Intel in the CPU market, the company moved into graphics in a big way in 2006 with its acquisition of ATI. In 2011, AMD sold more than 132 million discrete graphics cards. While most were Radeon boards aimed at the consumer market, more than 23% of revenue came from sales of professional graphics products. According to AMD managers we spoke with, the company is focusing on key vertical markets and expanded ISV support as the new FirePro boards become part of a stronger product mix.

Next GPU Technology: New Graphics Core

At the Sunnyvale event, AMD unveiled four new FirePro boards—the mid-range W5000, the high-end W7000 and W8000, and the ultra-high-end W9000. All four are based on the company's fourth-generation Graphics Core Next (GCN) graphics processing unit (GPU), which features improved graphics performance, increased compute density (higher performance per square mm of board real estate), multi-tasking, and improved power efficiency. According to AMD, the new boards are optimized for heterogeneous computing and offer enhanced scalability. In previous generations of graphics boards, if you wanted to use the GPU for computing, you had to turn off the graphics until computations were completed—or use two graphics cards, one for

The mid-range AMD FirePro W5000. *Images courtesy of AMD.*

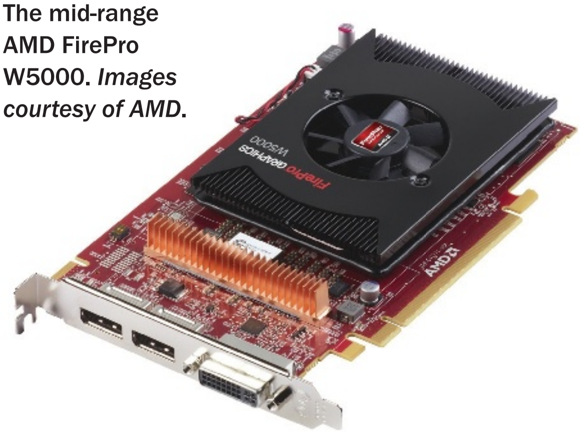

The high-end AMD FirePro W7000.

The high-end AMD FirePro W8000.

The ultra-high-end AMD FirePro W9000.

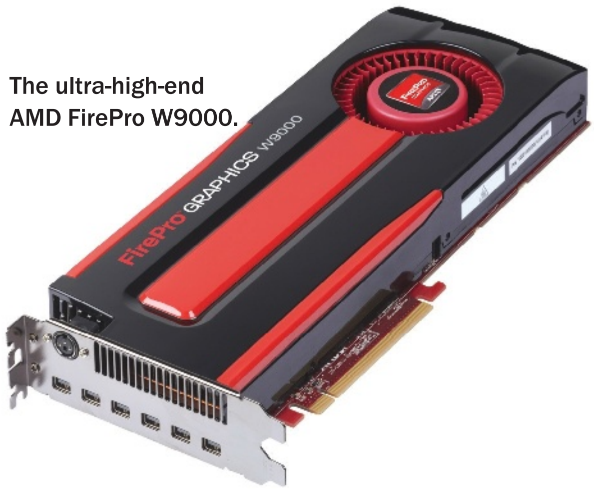

computing and one for graphics. AMD's GCN technology enables execution of a graphics thread and up to two compute threads per clock cycle at the same time on the GPU.

AMD has also implemented partially resident textures (virtual texturing), textures that have only portions of the texture stored in GPU memory at runtime. PowerTune technology monitors the boards' power consumption to deliver faster GPU clock speed while remaining within thermal design limits. If an application is not making the most power available to the GPU, PowerTune can improve that application's performance by raising the GPU's clock speed by up to 30% automatically. And to reduce power consumption, AMD ZeroCore Power technology reduces idle power by up to 95% by shutting down the GPU when the system is idle.

All four of the new boards support PCI Express 3.0, whereas the previous generation only supported PCIe 2.1. The new FirePro W-series boards also allow a maximum DisplayPort 1.2 resolution of 4096x2160 at 60Hz (compared to 2560x1600 in the previous generation).

The New AMD FirePro W Family

For the mid-range, the AMD FirePro W5000 essentially replaces last year's V5900. The new board carries a manufacturer's suggested retail price (MSRP) of $599, the same as its predecessor. Its average street price of $450 also matches that of the V5900 when it was first released.

The board's 2GB of GDDR5 memory and 256-bit interface also match that of its predecessor, but bandwidth improves to 102.4GB/second (from 65GB/s in the V5900). Based on a Pitcairn GL-PRO GPU with 768 stream processors and an 825MHz engine clock, AMD says that the W5000 can process up to 1.65 billion triangles per second. GeometryBoost allows the GPU to process geometry data at a rate of twice per clock cycle, resulting in a doubling in the rate of primitive and vertex processing.

Like its predecessor, the W5000 consumes a single slot and draws a maximum of 75 watts. And like the V5900 before it, the FirePro W5000 provides one dual-link DVI output and a pair of DisplayPort connections.

The W5000 supports CrossFire Pro, which enables users to harness the power of two GPUs by linking two identical FirePro cards together with the supplied connection cable. The board also comes with a DisplayPort-to-single-link-DVI adapter. But the new board also supports stereo 3D via an expansion bracket, a feature not found on the V5900.

At the high end, the AMD FirePro W7000 replaces last year's FirePro V7900. But the new board's MSRP of $899 is $100 less than its predecessor. With an average street price of $759, however, the W7000 costs essentially the same as the V7900 did when it was first released. The new Pitcairn GL-XT GPU has 1,280 stream processors and a 950MHz engine clock. It retains the same 256-bit interface as its predecessor, although the W7000's memory bandwidth drops a bit, to 154GB/second. But the W7000 comes with 4GB of GDDR5 memory—double that of last year's board.

The company claims floating point compute performance of 2.4 teraflops for single precision and 152 gigaflops for double precision. Power consumption of 150 watts matches that of the older board, so although the W7000 uses just a single PCIe X16 slot, a 6-pin auxiliary power connector is required.

Like its predecessor, the W7000 eschews DVI ports in favor of four DisplayPort connectors. You can use a DisplayPort-to-single-link-active-DVI connector should you want to attach an older monitor, and the W7000 comes with two of these adapters. And like the V7900, the W7000 supports stereo 3D via an additional bracket. The board also supports framelock/genlock via an optional FirePro S400 synchronization module—and also supports CrossFire Pro.

The new W8000 replaces the FirePro V8800 as AMD's category leader for multitasking and application performance across up to four displays. With an MSRP of $1,599 ($1,325 average street price), the W8000 is $100 more expensive than its predecessor. That extra cost seems justified, however, considering the W8000 comes with 4GB of GDDR5 memory—double that of the previous board, and on the W8000 it's error correcting code (ECC) memory. Both boards sport the same 256-bit interface, but the W8000 raises the memory bandwidth to 176GB/second.

The W8000 is based on AMD's new Tahiti GL-PRO GPU, with 1,792 stream processors and a 900MHz engine clock. AMD claims floating point compute performance of 3.23 teraflops for single precision and 806 gigaflops for double precision.

Of course, the power requirements of the new board are also significant: 225 watts compared to 208 watts for its predecessor. Its two required 6-pin auxiliary power connections mean that the W8000 can only be installed in a workstation that can supply the necessary juice. The W8000 is also so thick that it covers the adjacent expansion slot. Like the V8800, the W8000 provides four DisplayPorts and a 3-pin Stereo 3D connector, and also accommodates the S400 synchronization module for framelock/genlock support. The W8000 also features two DisplayPort-to-DVI connectors and a CrossFire Pro connector.

At the ultra-high-end, the AMD FirePro W9000 comes with 6GB of ECC GDDR5 memory—compared to 4GB of non-ECC memory in the V9800 that it replaces—and an MSRP of $3,999. Based on the top-of-the-line Tahiti GL-XT GPU, with

SPECviewperf Benchmark Results for Current AMD FirePro Graphic Cards Reviewed

	AMD FirePro W9000 NEW!	AMD FirePro V9800	AMD FirePro W8000 NEW!	AMD FirePro V8800	AMD FirePro W7000 NEW!	AMD FirePro V7900	AMD FirePro W5000 NEW!	AMD FirePro V5900
Manufacturer's price	$3,999	$3,499	$1,599	$1,499	$899	$999	$599	$599
Average street price	$3,400	$2,899	$1,325	$1,186	$759	$759	$450	$459
SPECviewperf 11.0								
catia-02	20.42	n/a	22.01	20.71	22.02	23.65	20.83	20.27
ensight-03	56.80	n/a	52.68	26.93	55.53	41.85	48.02	32.68
lightwave-01	69.77	n/a	70.34	56.64	69.42	59.13	69.93	58.79
Maya-03	52.64	n/a	47.52	39.79	46.53	43.69	50.04	47.12
proe-05	5.57	n/a	5.49	5.08	5.56	5.35	5.55	5.29
sw-02	58.98	n/a	59.17	46.00	58.13	49.01	58.54	47.12
tcvis-02	20.21	n/a	47.11	20.56	18.46	19.10	18.39	16.50
snx-01	43.51	n/a	47.11	27.27	48.27	38.93	44.14	31.58
Specifications								
Bus architecture	PCI Express	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 (2) (6)	Yes (2) (6)	Yes (2)	Yes (2)	Yes	Yes	No	No
Form factor	4.38"x11.00"	4.38"x10.5"	4.38"x11.00"	4.38"x10.5"	4.38"x9.63"	4.38"x11.00"	4.38"x7.31"	4.38"x9.00"
Slots used	2	2	2	2	1	1	1	1
Max power (watts)	274W	225W	225W	208W	150W	150W	150W	75W
PCIe version	3.0	2.1	3.0	2.1	3.0	2.0	3.0	2.0
Length	full-length	full-length	full-length	full-length	full-length	full-length	half-length	full-length
Processors	2048	1600	1792	1600	1280	1280	768	512
Memory configuration	6GB (GDDR5) ECC	4GB (GDDR5)	4GB (GDDR5) ECC	2GB (DDR5)	4GB (GDDR5)	2GB (GDDR5)	2GB (GDDR)	2GB (GDDR5)
Memory interface	384-bit	256-bit	256-bit	256-bit	256-bit	256-bit	256-bit	256-bit
Memory bandwidth	264 GB/sec	147.2 GB/sec	176 GB/sec	147.2 GB/sec	154 GB/sec	160.0 GB/sec	102.4 GB/sec	64.0 GB/sec
Dual-link DVI outputs	0	0	0	0	0	0	1	1
Display port outputs	6 (mini DP)	6	4	4	4	4	2	2
Stereo 3D connector (3-pin)	Yes	Yes	Yes	Yes	Yes (2)	Yes (2)	Yes (2)	No
Framelock/genlock	Yes (1)	Yes (1)	Yes (1)	Yes (1)	Yes (1)	Yes (1)	No	No
ATI CrossFire Pro	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
OpenGL version	4.2	4.1	4.2	4.1	4.2	4.1	4.2	4.1
DirectX/shader model	11/5.0	11/5.0	11/5.0	11/5.0	11/5.0	11/5.0	11/5.0	11/5.0
Maximum resolution support (@ 60Hz)	4096x2160 (3) 2560x1600 (4)	2560x1600 (4)	4096x2160 (3) 2560x1600 (4)	2560x1600 (4)	4096x2160 (3) 2560x1600 (4)	2560x1600 (4)	4096x2160 (3) 2560x1600 (4x5)	2560x1600 (4)

Notes:
1. Requires optional Synchronization module
2. With included expansion bracket
3. Resolution with DisplayPort 1.2
4. Resolution with DisplayPort 1.1
5. DVI resolution
6. Two 6-pin or one 8-pin auxilliary power

an incredible 2,048 stream processors and a 975MHz engine clock, the W9000 has a 384-bit memory interface and achieves a massive memory bandwidth of 264GB/second.

AMD claims floating point compute performance of 4.0 teraflops single precision and 1.0 teraflops double precision. At 274 watts, however, its power consumption is also massive—approximately 50 watts higher than its predecessor. To support this board, you'll need both a 6-pin and an 8-pin auxiliary power cable, and the extra-thick W9000 covers the adjacent expansion slot.

As one would expect, the W9000 has a 3-pin mini-DIN for stereo 3D, and offers optional framelock/genlock support, but while both the V9800 and W9000 support up to six displays, the ports on the new W9000 are mini DisplayPorts. The W9000 comes with two mini-DisplayPort-to-single-link-DVI adapters and a CrossFire Pro connector.

Benchmarking the Boards

We tested the new AMD FirePro W5000, W7000, W8000 and W9000 boards in an HP Z820 workstation equipped with a pair of Intel Xeon E5-2687W 3.1GHz eight-core CPUs and 32GB of memory, running the 64-bit version of Windows 7. That system supports PCIe 3.0, and its 1125-watt power supply has more than enough capacity to meet the demands of the ultra-high-end boards, although we needed a special 6-to-8-pin power adapter to handle the W9000.

We ran version 11 of the SPECviewperf video benchmark. We also retested the previous generation FirePro V-series boards in the HP Z820, using the latest FirePro video driver.

Based on our results, the newest AMD FirePro boards surpass the performance of the older boards, with the new mid-range W5000 actually outpacing the older, high-end V7900 and the new W7000 outperforming the V8800. Because we didn't test the V9800, we were not able to make any comparisons to that board.

We also tested several new- and previous-generation NVIDIA graphics boards in the same HP Z820 workstation, so that we could see how AMD stacks up against its competition (see page 44 for a quick price/performance comparison).

While we certainly saw improvements from one generation to the next, the differences on the SPECviewperf benchmark were simply not as great as we had expected, particularly for the W9000. Remember, however, that benchmarks are synthetic tests that don't necessarily reflect real-world performance. We suspect that the benchmarks simply never really pushed the W9000 board to where its 6GB of memory and 2,048 processors would show off its true power.

Considering its price, this board is not for everyone. But if you work with very large models or run GPU-based computations, the W9000 is worth considering.

The new AMD FirePro boards are fully certified with most CAD and DCC applications, and all of the boards in the FirePro line use the same unified video driver. Drivers are available for most 32- and 64-bit operating systems, including Windows 8, Windows 7, Windows Vista, Windows XP, and Linux. All of the new AMD FirePro boards carry a three-year limited product repair/replacement warranty, and direct toll-free phone and email tech support.

AMD continues to sell the previous-generation FirePro V-series boards, so customers have choices at every price point. With AMD's increased focus on the pro workstation market, competition with rival NVIDIA is clearing heating up. **DE**

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

INFO ➜ AMD: AMD.com

AMD Fire Pro W9000 **MSRP:** $3,999 • **Avg. Street Price:** $3,400

AMD Fire Pro W8000 **MSRP:** $1,599 • **Avg. Street Price:** $1,325

AMD Fire Pro W7000 **MSRP:** $999 • **Avg. Street Price:** $759

AMD Fire Pro V5000 **MSRP:** $599 • **Avg. Street Price:** $450

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

NVIDIA Delivers on Price vs. Performance

The company's new Kepler-based Quadro K5000 turns in a stellar performance.

BY DAVID COHN

It's been more than a year since we last looked at workstation-class graphics accelerators from NVIDIA (see *DE*, July 2011). At that time, we reviewed four boards—ranging from the entry-level Quadro 600 to the high-end Quadro 4000—all based on the company's Fermi architecture. We had also previously reviewed the ultra-high-end Quadro 5000, also a Fermi-based board.

In August 2012, NVIDIA announced its new Quadro K5000, the first Kepler-based GPU built as part of its NVIDIA Maximus platform. Maximus, which was first introduced in November 2011, enables workstation users to simultaneously perform complex analysis and visualization on a single machine. While the Quadro K5000 delivers state-of-the-art graphics performance, it can also be combined with the new NVIDIA Tesla K20 CPU computing accelerator, freeing up the new NVIDIA Quadro K5000 GPU to handle the graphics functions.

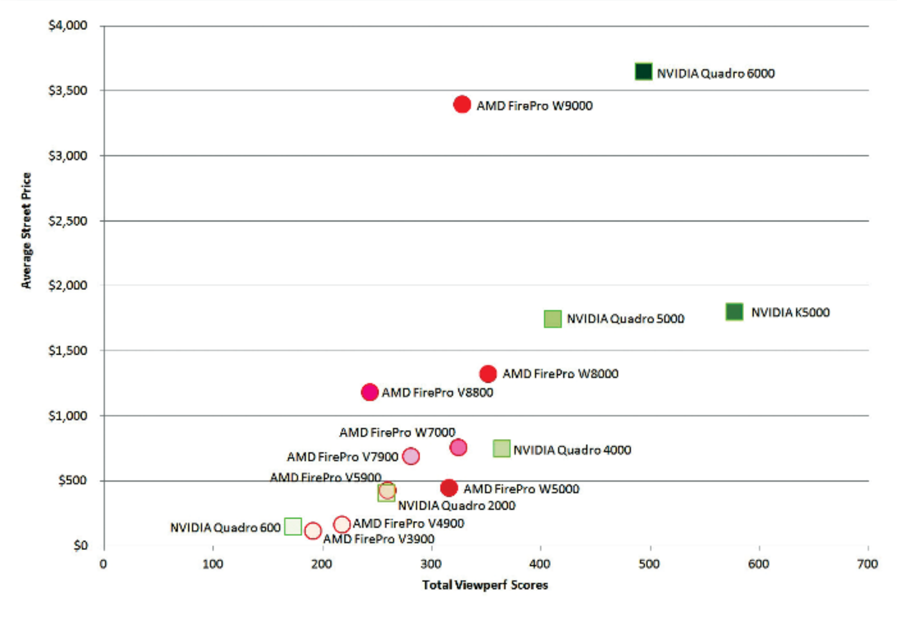

Price/performance of current NVIDIA Quadro boards and comparable AMD FirePro boards.

We ran the new Quadro K5000 through its paces, and also looked at the company's other ultra-high-end Quadro 6000 board, the company's top-of-the-line Fermi-based board.

NVIDIA Quadro K5000 Hits Price Sweet Spot

The NVIDIA Quadro K5000 is the company's new high-end graphics solution, yet with a manufacturer's suggested retail price (MSRP) of $2,249 and an average street price around $1,800, the K5000 delivers a lot of bang for the buck. The Quadro K5000 features 1536 compute unified device architecture (CUDA) parallel processing cores and includes 4GB of GDDR5 error correcting code (ECC) memory.

The board's next-generation streaming multiprocessor design, called SMX, offers several important architectural changes. These include substantial increases in per-clock throughput of key graphic operations that combine to deliver unprecedented performance and power efficiency. NVIDIA claims floating point performance of 2.15 teraflops single precision and 90 gigaflops double precision. With a 256-bit memory interface and a memory bandwidth of 173GB/second, the Quadro K5000 delivers 1.8 billion triangles per second.

The new NVIDIA Kepler architecture also introduces the concept of bindless textures, which enables the GPU to reference textures directly in memory, effectively eliminating the limit on the number of unique textures that can be used to render a scene, and reducing the CPU overhead to deliver improved performance.

The Quadro K5000 provides a dual-link DVI connector, as well as two DisplayPort 1.2 connectors, which support resolutions of up to 3840x2160 at 60Hz. You can also add stereo 3D with the addition of an optional 3-pin mini-DIN. The board also accepts an optional Quadro Sync card for framelock/genlock, and the board supports SLI and SDI.

Although the K5000's thickness means that it takes up the space of the adjacent expansion slot, this NVIDIA board has a maximum power consumption of only 122 watts, so it requires just one additional 6-pin auxiliary power connector.

NVIDIA Quadro 6000 Delivers Power

With a manufacturer's suggested retail price of $4,999 and an average street price still holding at $3,650—more than a year after its introduction—the NVIDIA Quadro 6000 is certainly not for everyone. But for those who need a huge frame buffer coupled with lots of compute power, the Quadro 6000 certainly delivers.

The Quadro 6000 features 448 CUDA parallel processing cores and includes 6GB of GDDR5 memory. The use of ECC memory offers protection of data in memory to enhance application data integrity.

The board features a 384-bit memory interface and a memory bandwidth of 144GB/second. NVIDIA claims floating point performance of 1.03 teraflops single precision and 515 gigaflops double precision. NVIDIA's Scalable Geometry Engine technology enables the Quadro 6000 to deliver 1.3 billion triangles per second.

Output options include a dual-link DVI connector and a pair of DisplayPorts. There's also a built-in 3-pin stereo connector, and the Quadro 6000 also accepts an optional Quadro G-Sync card for framelock/genlock. Although all of this power results in an extra-thick board that takes up the space of the adjacent expansion slot, the Quadro 6000 has a maximum power consumption of 204 watts, so it requires either two 6-pin or one 8-pin auxiliary power connectors.

This board's maximum digital display resolution reaches only 2560x1600 at 60Hz, but the Quadro 6000 does incorporate a Parallel DataCache with on-chip shared memory to accelerate real-time ray tracing, physics processing, and texture filtering. It also features Dual Copy Engines, which allow simultaneous transfer of data between the GPU and the host computer—further accelerating operations such as ray tracing and physical simulations.

Benchmarking the Boards

We tested both the new NVIDIA Quadro K5000 and the Quadro 6000 in an HP Z820 workstation equipped with a pair of Intel Xeon E5-2687W 3.1GHz eight-core CPUs and 32GB of memory, running the 64-bit version of Windows 7. That system supports PCIe 3.0, and its 1,125-watt power supply has more than enough capacity to meet the demands of even ultra-high-end boards.

SPECviewperf Benchmark Results for Current NVIDIA Quadro Series Graphic Cards Reviewed

	NVIDIA Quadro 6000 NEW!	NVIDIA K5000 NEW!	NVIDIA Quadro 5000	NVIDIA Quadro 4000	NVIDIA Quadro 2000	NVIDIA Quadro 600
Manufacturer's price	$4,999	$2,249	$2,249	$1,199	$599	$199
Average street price	$3,650	$1,800	$1,749	$750	$410	$150
SPECviewperf 11.0						
catia-02	69.52	77.21	50.86	47.00	32.59	17.68
ensight-03	61.60	81.76	43.18	31.90	21.12	10.89
lightwave-01	66.69	72.65	59.48	68.69	65.12	51.79
Maya-03	109.27	115.24	99.73	81.26	37.55	23.89
proe-05	11.11	15.30	11.37	11.79	11.08	9.37
sw-02	60.82	60.80	57.05	52.88	42.24	30.27
tcvis-02	52.26	71.49	44.11	36.92	25.92	15.91
snx-01	62.65	82.10	44.66	33.95	22.69	13.30
Specifications						
Bus architecture	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16	PCI Express X16
Extra power req'd	Yes (2) (6)	Yes	Yes	Yes	No	No
Form factor	4.38"x9.75"	4.38"x10.50"	4.38"x9.75"	4.38"x9.75"	4.38"x7.0"	2.73"x6.6"
Slots used	2	2	2	1	1	1
Max power (watts)	204W	122W	152W	142W	62W	40W
PCIe version	2.0	2.0	2.0	2.0	2.0	2.0
Length	3/4-length	Full-length	3/4-length	3/4-length	2/3-length	1/2-length
Processors	448	1536	352	256	192	96
Memory configuration	6GB (GDDR5) ECC	4GB (GDDR5)	2.5GB (GDDR5)	2GB (GDDR5)	1GB (GDDR5)	1GB (DDR3)
Memory interface	384-bit	256-bit	320-bit	256-bit	128-bit	128-bit
Memory bandwidth	144 GB/sec	173 GB/sec	120 GB/sec	89.6 GB/sec	41.6GB/sec	25.6 GB/sec
Number of dual-link DVI outputs	1	1	1	1	1	1
Number of display port outputs	2	2	2	2	2	1
Stereo 3D connector (3-pin)	Yes	Yes (2)	Yes	Yes (2)	No	No
SDI-enabled	Yes	Yes	Yes	Yes	No	No
Framelock/genlock	Yes (1)	Yes (1)	Yes	No	No	No
SLI	Yes	Yes	Yes (2)	No	No	No
OpenGL version	4.3	4.3	4.0	4.1	4.0	4.0
DirectX/shader model	11/5.0	11/5.0	11/5.0	11/5.0	11/5.0	11/5.0
Maximum resolution support (@ 60Hz)	2560x1600 (1)	3840x2160 (3) 2560x1600 (4)	2560x1600 (4)	2560x1600 (4)	2560x1600 (4)	2560x1600 (4)

Notes:

1. Requires optional Quadro G-Sync or Quadro Sync card
2. With included expansion bracket
3. Resolution with DisplayPort 1.2
4. Resolution with DisplayPort 1.1
5. DVI resolution
6. Two 6-pin or one 8-pin auxilliary power

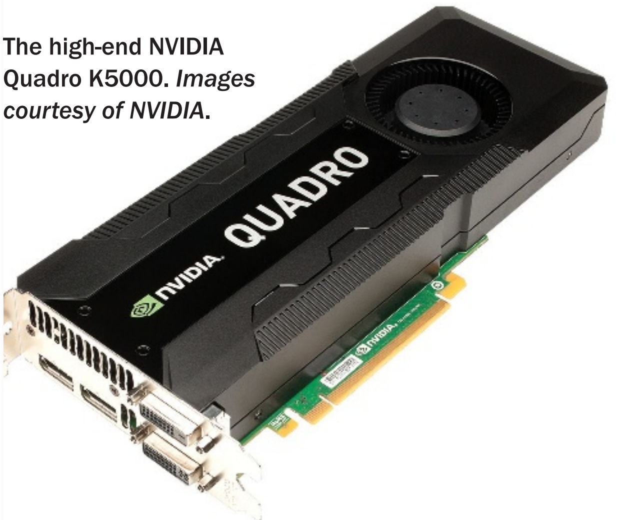

The high-end NVIDIA Quadro K5000. *Images courtesy of NVIDIA.*

The ultra-high-end NVIDIA Quadro 6000.

We ran version 11 of the SPECviewperf video benchmark. We also used the HP Z820 to retest all of the other NVIDIA Quadro boards we reviewed last year, using the latest version of the NVIDIA video driver, so that we'd have a direct comparison of all of the boards.

Based on our results, the new Quadro K5000 clearly outperformed every other board we've ever reviewed, including the Quadro 6000. We also tested new- and previous-generation AMD FirePro graphics boards in the same HP Z820 workstation so that we could see how NVIDIA stacks up against its competition. You can make some quick comparisons in our Price/Performance chart (see page 44).

Of course, like all other NVIDIA Quadro boards, the Quadro 6000 and K5000 are fully certified with most CAD and DCC applications. All of the boards in the Quadro line use the same unified video driver. Drivers are available for most 32- and 64-bit operating systems, including Windows 8, Windows 7, Windows Vista, Windows XP, Linux, Solaris and FreeBSD. NVIDIA also released a version of the Quadro K5000 for Apple Macintosh, and provides drivers for OSX Leopard (10.5.7), Snow Leopard (10.6.8), Lion (10.7.5), and Mountain Lion (10.8.2).

Clearly, NVIDIA has delivered a significant new level of performance with its new Kepler-based Quadro K5000. It will be extremely interesting to see if and when the company's latest GPU technology trickles down to the company's midrange and entry-level boards. **DE**

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

INFO → **NVIDIA Corp.:** NVIDIA.com

NVIDIA Quadro K5000 **MSRP:** $2,249 • **Avg. Street Price:** $1,800

NVIDIA Quadro 6000 **MSRP:** $4,000 • **Avg. Street Price:** $3,650

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

Advertising Index

Commentary

Bernt Nilsson

The True Power of Simulation

When we launched the first version of our multiphysics (MP) simulation software, we knew it would fill a void in the CAE market that historically had been limited to single-physics codes only. Fueled by the ingenuity of its early users, the software was applied to solve just about any design challenge that involved coupled physics effects. While addressing the need for coupling physics affecting the same application, it became clear that MP modeling was a vital component for delivering a unified simulation platform.

Today, COMSOL Multiphysics has become an integral tool throughout the product development process. It encompasses a wide range of applications, including electrical, mechanical, fluid and chemical simulations. Companies who have adopted this simulation power are designing new products in less time and with lower costs than ever before. It all comes down to the ability to build accurate simulation models.

Accurate simulation software boosts productivity for the design of electrical, mechanical, fluid and chemical applications.

Accurate Model Characteristics

To start, the model needs to replicate the characteristics of a design by including all relevant physics. In addition, the contributing physical effects need to be coupled the same way as they are in nature. This is accomplished by solving a tightly integrated system of physics, each governed by the laws of nature, in one and the same simulation. The objective is to find a self-consistent solution that satisfies all physics. This is the MP approach that enables reliable simulation results.

Using accurate models from the very start of product development is the key to increased productivity. This was a recurring theme at the 2012 COMSOL Conferences held this past fall. More than 2,000 researchers and product designers gathered to present and discuss the groundbreaking MP advances in their workflow. Presentations by leaders in mathematical modeling and simulation spanned a variety of industries, such as automotive, life sciences, energy, consumer electronics, and many more.

Case in Point

Where does MP fit in the product design workflow? A great example can be found throughout the conceptual phase of design, where new ideas are explored. Making the right design decision at this stage is crucial in bringing the right product to market on time and within budget. Any simplifications applied during the early design phase, such as considering single physics only, risks jeopardizing the project down the road—with expensive change requests and even product recalls.

Here is where the adaptability of MP simulation makes a big difference. It supports the designer by building realistic models that include all relevant physical effects influencing the product. These simulation results can be used to gain understanding of the combined effects, and to drive key design decisions.

Design optimization is another area in which COMSOL users are reporting excellent results. With a true representation of the real-world product, the case is made for running large-scale optimizations. Today, the MP community enjoys a solid infrastructure of tools to get the job done. For example, seamless integration of CAD modeling with simulation streamlines the iterations in optimization of design. For cases where the simulation requires more computational power than what is available on the premises, users are turning to the cloud. Loads of RAM, high-end hardware, and cluster computing are offered through Amazon's Elastic Compute Cloud.

Verification and testing are other areas where productivity gains have been significant. A proven, reliable model simplifies the verification of design ideas and product alterations. In addition, MP simulation reduces the need for costly and time-consuming testing of physical prototypes. Reports from the COMSOL Conference indicate simulation results are accurate to such a high degree that, in many areas, physical testing can be significantly reduced.

MP simulation simply offers too great of a competitive advantage to be overlooked. Companies from a broad range of industries are successfully adopting these tools for conceptual design, optimization, testing and verification. Without a doubt, MP modeling is driving innovation and productivity gains to new levels. That is the true power of simulation. **DE**

With more than 20 years of experience from promoting analysis and simulation software, **Bernt Nilsson** *is the senior vice president of marketing of COMSOL Inc., Burlington, MA. Send e-mail about this article to DE-Editors@deskeng.com.*