November 2015 / deskeng.com

DE

Technology for Optimal Design Engineering

The Latest Processors P.16

3D Printing Thinks Big P.44

Workstation Reviews P.28

FOCUS ON: Engineering Computing P.12

Simulation Keeps Quantum Chip Cool

AUTOMOTIVE NVH TESTING P.42

HPC DEPLOYMENT TIPS P.20

CAE CLOUD OPTIONS P.23

ANSYS 16.2

NOW AVAILABLE

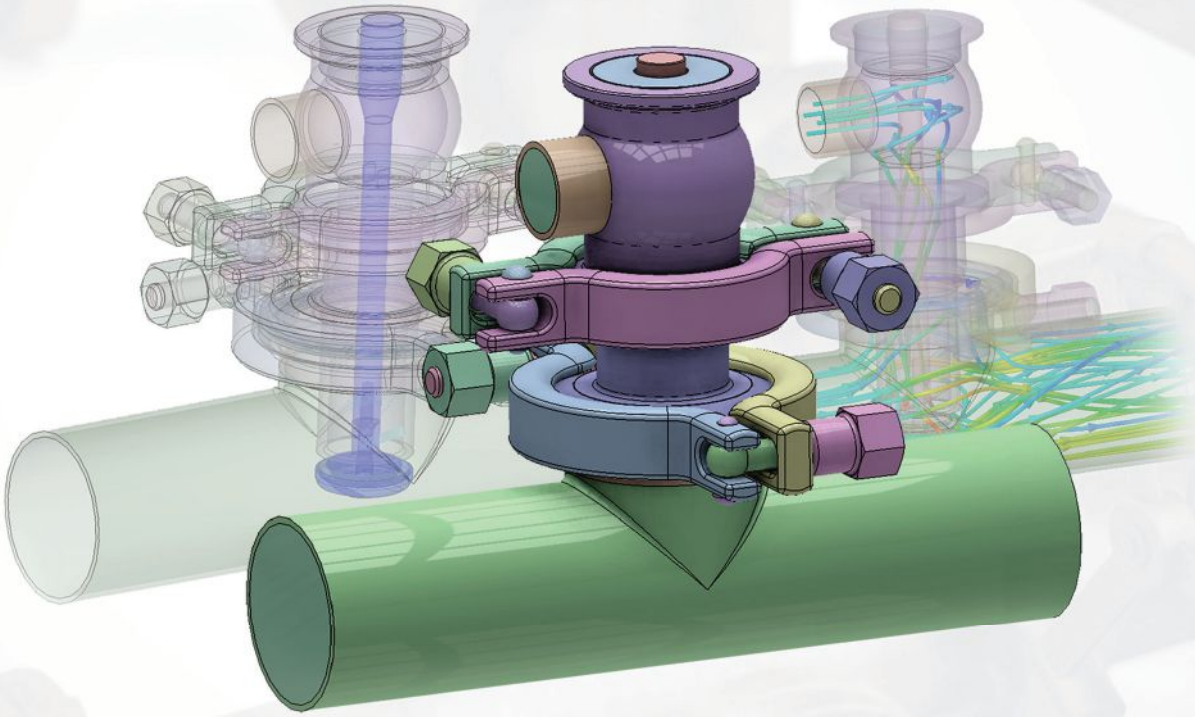

With ANSYS 16.2, you can simulate virtual prototypes of complete systems - ensuring you can Realize Your Product Promise®.

ANSYS®

Learn more at ANSYS.COM/16.2

The World's Smallest, Overclocked, Liquid-Cooled Workstation

APEXX 1

5" wide , 8.5" high, 9" deep

Expandable and Stackable

Available with up to 18 cores

APEXX Workstations

The world's fastest design, visualization, and motion media systems, APEXX workstations feature dual or single Intel processors, safe overclocking, groundbreaking multi-GPU editions, and state-of-the-art liquid-cooling.

intel CORE i7 inside

intel XEON inside

MADE IN THE USA

BOXX TECHNOLOGIES

For more information on APEXX 1 and the full APEXX line visit:
www.boxxtech.com/apexx1 OR call **888-984-7589**

by Jamie J. Gooch

HPC Is Not Yet For All

Upon seeing the massive and tragic flooding of South Carolina in the news last month, my teenage daughter asked why all that excess water couldn't be transported to drought-stricken California, where it's needed. Why not just fill up a convoy of tanker trucks and head west? I remember having a similar thought when I was about her age and learned the majority of the Earth's fresh water is frozen in the polar ice caps. Why not just tow an iceberg to Ethiopia?

I began to to explain the complicated supply chain logistics involved in such a project, the significant engineering challenges and the slim chance for a return on investment ... "So it's just too hard and costs too much?" my daughter asked, right to the point.

Just because enough of a resource exists doesn't mean everyone can access it. That brings me to the current state of high-performance computing (HPC).

Massive computing power is available, but elusive for some.

The Democratization of HPC

We've all heard of Jeopardy!-winning Watson, IBM's cognitive computing effort that can read millions of unstructured documents in seconds and finds patterns in them that can help make decisions. We see the lists of top supercomputers every year, many owned by government departments and universities. According to the TOP500 list of the world's most powerful supercomputers, the top performer in June, Tianhe-2, hit 33.86 petaflop/s (quadrillions of calculations per second) on the organization's benchmarks. Maybe you've seen the colorful images of Google's data centers, with thousands of feet of color-coded pipes for cooling the racks (if you haven't, see google.com/about/datacenters). It's clear that massive computing power is available out there. And yet, one of the top complaints we still hear from engineers is how long it takes to run simulations or even update large model assemblies.

Why aren't more engineers tapping into the power of high-performance computing? The short answer for many: It's just too hard and costs too much.

Cloud computing was hyped as a great equalizer for small businesses to compete against the big guys. For some it is. Small companies that don't have dedicated IT staffing or hardware resources can turn to the cloud for everything from e-mail to online CAD to simulation runs. As we note in this issue's focus on engineering computing, not all cloud-computing resources are created equal (see page 20). Still, many third-party cloud service providers and software vendors are addressing the ease-of-use and licesnsing issues (see page 23).

For large companies that have a hefty investment in computing infratructure and staff, there's good news as well. There are a plethora of options available to them. Even supercomputers are not out of reach, as prices have fallen dramatically. For engineers in large companies who want to access more computing power, the biggest hurdles are often internal policies and red tape.

Finding the Missing Middle

Mid-sized companies have options as well, but for them the choice is not as clear. Pay-as-you-go models don't always make sense for companies who use advanced simulation regularly. Investing in additional on-premise resources may not be attractive when it seems the whole world is moving to the cloud. When you're working with very large files, moving to the cloud isn't as easy. There's a significant investment in time just to get your file uploaded; so much so that it might not be worth the speed boost you get in using off-site resources to solve the simulation.

The issue of getting large amounts of data into the cloud is so common that Amazon launched AWS Import/Export Snowball last month. The company calls it "a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of AWS (Amazon Web Services)." The appliance looks like a big, hard plastic desktop tower case. It has a Kindle on the outside and drives to hold up to 50TB of data on the inside, along with its own power supply and 10GB networking port. Users plug the appliance into their local network, install its client and then transfer their files to it. The client encrypts the files as they are transferred. Once the transfer is complete, the appliance can be sent back to Amazon via FedEx. The Kindle on the side acts as an E Ink shipping label. While this may seem like a step backward to the days of sneakernet, when it was faster to walk a file over to the next cubicle than it was to email it, Amazon points out that it can take months and cost thousands of dollars to transfer 100TB of data over a dedicated 100 Mbps connection.

It's easy to forget we have the power of yesterday's supercomputers in today's desktop workstations. That sets a high bar. We now expect the same ease of use when it comes to tapping into HPC resources. We're not there yet, but we've begun addressing the logistics, engineering challenges and ROI needed to make computing resources accessible to all. DE

Jamie Gooch *is the editorial director of* Desktop Engineering. *Contact him at de-editors@deskeng.com.*

HyperWorks Unlimited™

CAE Cloud Appliance for Infinite Design Exploration

Increasing the pace of innovation with unlimited design exploration

Intel® Cluster Ready

Physical Appliance:

On-premise, plug-and-play private cloud solution

powered by amazon web services

Virtual Appliance:

Off premise, hosted on AWS

Learn more at altair.com/cloud

Altair

INNOVATION INTELLIGENCE®

DE November 2015 VOLUME 21/ISSUE 2

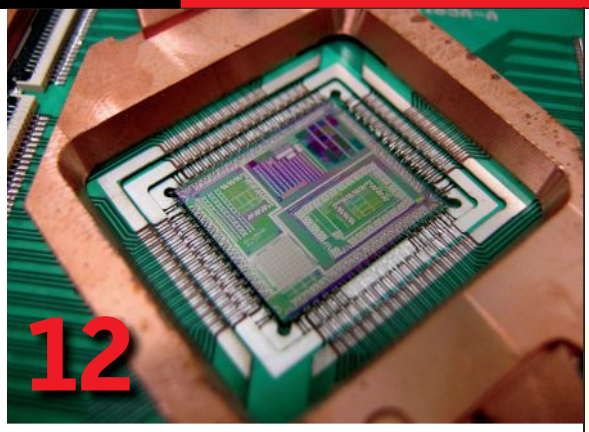

Extreme Computer Engineering at D-Wave

Engineers at D-Wave Systems explain how computational simulation and modeling help them get the job done.

By Michael Belfiore

WORKFLOW

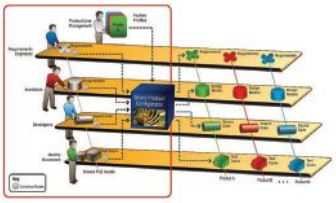

38 PLE FYI

What product line engineering means to the design engineer.

By Frank J. Ohlhorst

ENGINEERING SERVICES

40 Optimizing and Compromising Lead to Winning Design

Altair ProductDesign helps DECKED design, simulate and test an innovative truck bed storage system.

By Jamie J. Gooch

TEST

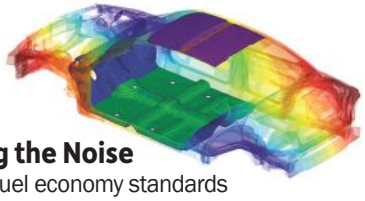

42 (Don't) Bring the Noise

Lightweighting and fuel economy standards have created new noise and vibration challenges for engineers.

By Brian Albright

ON THE COVER: D-Wave Systems used ANSYS software to model and simulate the environment around the chip inside its quantum computer. *Image courtesy of D-Wave Systems.*

PROTOTYPE

44 3D Printing Thinks Really Big

Additive manufacturing systems offer amazing build volumes.

By Pamela J. Waterman

FOCUS ON ENGINEERING COMPUTING

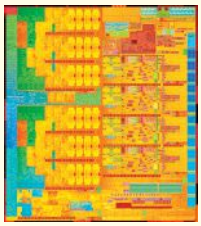

16 Processor Progress

Intel's new CPU with integrated graphics moves into territories traditionally dominated by the GPU.

By Kenneth Wong

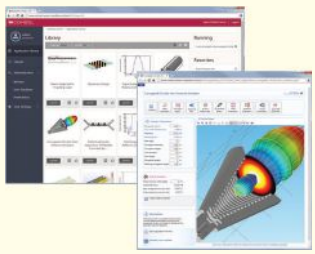

20 Help for HPC in the Cloud

Engineering organizations offloading simulation workflows to cloud-based HPC will need a helping hand from specialists and service providers to optimize their deployments.

By Beth Stackpole

23 Simulate Anytime, Anywhere

SaaS, IaaS and PaaS providers let engineers use CAE in the cloud exactly when they need it.

By Jessica Lulka

25 The Horizon Beyond the Desk

The workstation takes on a new life through virtual desktop infrastructure.

By Kenneth Wong

28 Mobile Workstation Review: Still Sleek

The updated Dell Precision M3800 G2 remains slim, fast and affordable.

By David Cohn

31 Workstation Review: Very Fast Indeed

Computer Direct Business delivers the speedy Volta Pro VP1 engineering workstation.

By David Cohn

DEPARTMENTS

PUBLISHER
Tom Cooney

EDITORIAL
Jamie J. Gooch | Editorial Director
Kenneth Wong | Senior Editor
Anthony J. Lockwood | Editor at Large
Jess Lulka | Associate Editor

CONTRIBUTING EDITORS
Tony Abbey, Brian Albright, Mark Clarkson, David S. Cohn, John Newman, Frank Ohlhorst, Beth Stackpole, Peter Varhol, Pamela J. Waterman

ADVERTISING SALES
Chuck Curley | East Coast Sales • 508-663-1500 (x498)
Chris Casey | Midwest/West Coast Sales • 847-274-5476

ART & PRODUCTION
Darlene Sweeney | Director • darlene@deskeng.com

A PEERLESS MEDIA, LLC PUBLICATION
Brian Ceraolo | President and Group Publisher

ADVERTISING, BUSINESS, & EDITORIAL OFFICES
***Desktop Engineering*® magazine**
Peerless Media, LLC
111 Speen St., Suite 200, Framingham, MA 01701
Phone: 508-663-1500
E-mail: DE-Editors@deskeng.com
www.deskeng.com

Peerless MEDIA, LLC
www.peerlessmedia.com

Kenneth Moyes | President and CEO, EH Publishing, Inc.

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

***Desktop Engineering*®** (ISSN 1085-0422) is published monthly by Peerless Media, LLC, a division of EH Publishing, Inc. 111 Speen St., Suite 200 Framingham, MA 01701. Periodicals postage paid at Framingham, MA and 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.

Send all subscription inquiries to *Desktop Engineering* 111 Speen St. Suite 200, Framingham, MA 01701

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

Postmaster: Send all address changes to *Desktop Engineering*®, PO Box 1496 Framingham, MA 01701-1496.

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

BPA WORLDWIDE

Simulation-Powered Innovation

Why you need to take a holistic approach to design experiences, not just products.

The "Experience Economy" may sound like just another term used to justify shorter product development deadlines and longer lists of requirements, but it's firmly rooted in market demands. Customers today expect more than a product when they do business with a company — they want a great experience that starts before they even purchase the product and extends well beyond unboxing. Companies that can't provide those great experiences will lose out to those that can.

To create innovative experiences, an organization has to leverage agility, iterative approaches, organization-wide creativity and knowledge. It requires a product design process in which simulation is ubiquitous, flexible and intuitive, yet precise and meaningful. Simulation must be infused throughout the entire development cycle while facilitating the sharing of data, expertise and insights across an organization, not just within the engineering department. Traditional workflows that restrict simulation to a handful of experts need to be transformed. New processes and tools are needed to make product requirements more visible early in the development process. They should also open simulation tasks to a wider audience while capturing inputs and sharing results related to the holistic experience. In addition, simulation practices need to be built on a foundation that supports sustainable innovation, enabling organizations to continuously evolve the experience as opposed to focusing on one-off product designs.

Making the Case for Simulation in the Age of Experience

This new paper produced by *Desktop Engineering* on behalf of Dassault Systèmes makes the case for adopting a holistic approach to product development that enables the entire enterprise to take advantage of simulation-powered innovation. Multiple case studies and additional resources are provided.

Download the white paper for free at deskeng.com/de/experience.

Simulation as a standalone expert practice, or even as a simplified toolset for designers, won't hit the mark on its own to fuel next-generation innovation. What's required is a broad approach that delivers the collaboration, visibility, knowledge capture and computational power necessary for simulating a product's real-world performance and creating optimal product experiences that help companies stand out in a crowd.

Did You Know?

The SIMULIA simulation applications from Dassault Systèmes might be best known for its flagship Abaqus Finite Element Analysis (FEA) technology. However, in recent years the company has expanded its portfolio to include:

- Abaqus for FEA and multiphysics analysis
- Isight for process integration and design exploration
- fe-safe for fatigue and durability
- Tosca for structural and fluid optimization
- Simpoe-Mold for plastic simulation and injection molding analysis
- SIMPACK for multibody dynamic simulation

SIMULIA is actively integrating these robust technologies into the **3D**EXPERIENCE platform, which provides breakthrough capabilities to improve efficiency, throughput and collaboration for all simulation users including; high-performance visualization, batch and rule-based meshing, collaborative assembly, results analytics and process and data management.

To learn more about Dassault Systèmes, visit 3ds.com.

A Holistic, Industry Specific Approach

Dassault Systèmes provides world-leading solutions that transform the way products are designed, produced and supported. The group brings value to over 190,000 customers across all industries by providing collaborative solutions that foster social innovation which makes it possible for discoveries in the virtual world to be used to improve the real world.

The **3D**EXPERIENCE® platform from Dassault Systèmes has been architected from the ground up to make simulation an integral part of the cross-disciplinary workflow of specific industries. Simulation models, workflows and results are maintained as part of the same single source of truth for product and experience-related requirements. They are also accessible across disciplines so that simulation data is just as available to executive management and marketers as it is to design engineers and analysis experts. Because design changes are automatically reflected in the system of record, companies are able ensure that each engineer is 100% aware of consumer demands, performance requirements and their colleagues' work at each stage of development.

Increasing system complexity and more technology being used for co-simulation and integrated simulation workflows make it even more critical for design and engineering teams to leverage one source of truth to collaborate on meeting performance requirements.

The **3D**EXPERIENCE platform enables simultaneous collaboration on simulation-powered innovation, allowing, for example, engineers to validate and test designs against a number of parameters while manufacturing engineers evaluate whether a particular iteration can be produced efficiently at optimal cost. Early simulation encourages greater design exploration, allowing multidisciplinary teams to collaborate on more possibilities and zero in on alternatives that may have never been considered with traditional simulation practices.

Leveraging a holistic approach helps companies not only deliver the best performing product with the latest technological features, but also a market-winning experience that wows customers and gives companies a competitive edge.

For more details, download "Making the Case for Simulation in the Age of Experience," for free at deskeng.com/de/experience.

DASSAULT SYSTEMES | The **3D**EXPERIENCE® Company

A New Spin on Innovation

Imagine a washing machine that calculates and automatically balances load weight. It runs so quiet that it can be housed anywhere. And because it's connected, the consumer can control its operation with a mobile device. The **3D**EXPERIENCE platform allows innovators to design, simulate, and collaborate on evaluating thousands of variables including system performance, stress, durability, noise, and damage due to accidental drop or impact. Below are two examples of how realistic simulation can be used to bring this market-winning experience to life.

THE DESIRED EXPERIENCE: Reduce the 50-75 decibel noise levels common to most washing machines.

BENEFITS OF SIMULATION: Understand and improve frequency response using noise emission studies.

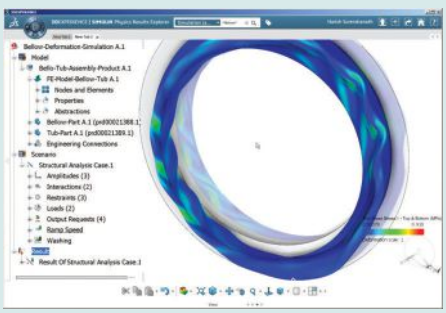

THE DESIRED EXPERIENCE: Reduce maintenance calls over the expected 10-year lifetime of a washing machine.

BENEFITS OF SIMULATION: Simulate potential component failure due to heavy and unbalanced loads over an extended period of time.

Autodesk, NI, and UC Berkeley to Jointly Develop Control System Design Workflow

At NIWeek 2015, Autodesk's vice president of Corporate Strategy Jon Pittman, National Instruments' (NI) Academic Programs Director David Wilson, and University of California, Berkeley's mechanical engineering Professor Emeritus David Auslander joined hands onstage during NI's annual user conference. Their handshake represents Autodesk, NI and UC Berkeley's shared interest to improve the tool chain linking mechanical modeling and system simulation.

The aim of the partnership is to make it possible to "take the output of Autodesk Inventor, put it together with Modelica [an object-oriented system modeling language], and feed it into the simulation and control design module in [NI's] LabVIEW," says Wilson

An Interconnected System

Usually, 3D mechanical assembly modeling and system modeling occur in separate silos. The former occurs in 3D MCAD programs like Autodesk Inventor, SOLIDWORKS and Solid Edge. The latter is considered the domain of system modelers like LabVIEW. But the growing use of software-driven functions and embedded systems — a phenomenon of the IoT era — pushes the two disciplines closer. (For more on this trend, read "CAD vs. System Modeling," deskeng.com/de/?p=25005).

"Control design and mechanical engineering are becoming much more interconnected," says Lance Grow, director of Inventor, Autodesk. "We don't believe there's a good multidisciplinary solution in the market today, and we think it'll become crucial to the development of connected products."

Berkeley's Auslander agrees there's a need for a solution. "In mechanical systems, the mathematics can be very messy. They're full of constraints and non-linearity, and can lead to complex algebraic equations that are difficult to handle — especially in robotics or in products that involve mechanical friction," he says.

Prevented by this daunting task, many who could benefit from simulation may have chosen to skip the step altogether. "Engineers could have designed the control system at the same time if they had taken the time to do dynamic simulation. And once the physical system has been built, it's very costly and time-consuming to refine or modify it," Auslander says.

Dave Wilson, NI's academic marketing director, delivering a keynote at NIWeek 2015. ***Image courtesy of NI.***

A round-trip workflow between Autodesk Inventor and NI LabVIEW will go a long way to connect mechanical system design to dynamic system simulation and control system design.

— K. Wong

VCs Pour $80 Million into Onshape

This year, Christmas came early for the team at Onshape. On September 24, the company announced it has secured $80 million in new equity funding from a number of VCs.

Andreessen Horowit, the California-based VC firm that led the latest round of financing, is cofounded by Marc Andreessen, remembered for his role in launching Netscape. The firm is known for taking chances on startups and concepts that others might shy away from.

The firm's investment in Onshape is a sign of the VC community's confidence in cloud-hosted CAD. Andreessen Horowit's general partner Peter Levine said, "We've been closely following the CAD industry for a long time. The Onshape team immediately stood out, having built a robust CAD system using a cloud-based technical architecture that most had considered impossible. The Onshape product is perfectly suited for an increasingly mobile workforce, which now can create and edit 3D models anywhere, on any device."

Onshape offers its parametric mechanical modeler as subscription software, running from a standard browser. This is a departure from traditional CAD software, which typically runs on a professional workstation and is distributed through resellers. Critics often cite poor bandwidth connection and CAD software's heavy computing demand as the reason cloud-hosted CAD isn't practical, but Onshape is proving its product can overcome both.

— K. Wong

HP, Dell Roll Out New Workstations

As September ended and October began, HP and Dell both unveiled new products targeting the professional design and engineering markets.

HP is launching the HP Z240 entry-level workstation in small form factor (SFF) and tower editions. As an upgrade to the previous Z230 model, the HP Z240 is "HP's most affordable workstation … ideal for customers in the video editing, MCAD/AEC, education, public sector and image viewing industries," according to HP.

The spec sheet indicates the Z240 will be available with "Intel Xeon processor E3-1200 v5 product families, Intel Core or Intel Pentium processors and two ultrafast HP Z Turbo Drive G2s." According to Intel, the Xeon E3-1200 V5 chip sets are "manufactured on the latest 14 nm technology" and "offer dramatically higher CPU and graphics performance as compared to the previous generation, a broad range of power options, and new advanced features to boost edge-to-cloud Internet of Things (IoT) designs." Buyers of the Z240 can choose from Windows 7, Windows 10 or Linux OS.

HP removed the legacy PCI slot, which the company estimates is used by less than 2% of customers. The design change leaves room for an integrated M.2 slot for expansion cards and connectors. The Z240 also offers optional dust filters, which HP says can reduce up to 47% dust ingress. The company revealed it has "re-engineered the hard drive cage for the Z240 SFF and engineered a custom air duct around the processor, allowing more efficient airflow, enhanced acoustics and a cleaner layout of internal cables."

U.S. pricing for the Z240 begins at $879. It's expected to be available by the end of this month.

The new HP Z240 entry-level workstation, shown here in small form factor and regular desktop editions. ***Image courtesy of HP.***

New Mobile and Tower Workstations from Dell

Dell's new lineup features thin and light workstations and desktop workstations from the Dell Precision M series. Entry-level SFF and tower workstations in the Dell Precision T series will also be released.

The new thin and light mobile units are Precision 3510 and 5510. The new high-performance mobile machines are Precision 7510 and 7710. They're all available with Intel Skylake H processors, the sixth-generation chips.

In the Precision M3510, the use of Polymer battery, carbon fiber chassis and other enhancements result in a mobile machine that's 27% thinner and 11% lighter, according to Dell's spec sheet. The M3510 features DDR4 memory with error correction, PCIe M.2 solid-state drives (SSDs), and SATA M.2 SSDs.

The Precision M5510, M7510, and M7710 from the thin and light series also benefit from the use of carbon fiber materials. The memory improvements published by Dell for its mobile workstations range from 33% in the M3510 to 43% in the M5000 and M7000 models.

The new tower units T3420 (SFF) and T3620 (regular tower) made their debut in October. The T3420 is 7% smaller than is predecessor T1700 SFF. They come with discrete GPUs (graphics processing units) from AMD and NVIDIA; and they also benefit from the new Intel CPU's integrated graphics (see page 16 for more information).

The 15-inch mobile workstation M3510 is priced beginning at $999. The 15-inch thin and light M5000 is priced beginning at $1,399. The M7510 and M7710 power workstations begin at $1,199 and $1,699, respectively. The desktop SFF T3420 and tower T3620 are priced beginning at $679 and $729 respectively. They're expected to become available in Q4 2015.

— K. Wong

by Kenneth Wong & Beth Stackpole

Elixir Takes Off With Cloud-Based Design Tools

Where does an aerospace startup find a platform to take its two-seater aircraft design to new heights? To the cloud, of course, where it can channel all the promised flexibility and ease of collaboration to help get its small luxury plane ready for takeoff.

Elixir Aircraft, a French aviation startup, has selected Dassault Systèmes' 3DEXPERIENCE platform, specifically its "Engineered to Fly" industry solution, as the primary design platform for its forthcoming plane. The Elixir aircraft, a two-seater plane aimed at enthusiasts, is leveraging lightweight composites, a unique low-wing profile and a single structure design in the wing and the fuselage to make an aircraft that is not only high-performance, but is safer to fly, says Arthur Leopold-Leger, CEO of the firm.

With no legacy systems in place and with no formal allegiance to any one design package, cloud-based CAD and PLM (product lifecycle management) were a natural fit for Elixir, providing flexibility for deployment considering the startup was light on budget and lacked dedicated IT resources. "We don't have an IT department and we didn't need any big computer installation," Leopold-Leger says. "From a money standpoint, [cloud-based software] was very interesting for us — we could just take what we needed when we needed it and not have to take care of anything."

Elixir's design incorporates a unique wing structure to achieve gains in payload, safety and costs. ***Image courtesy of Elixir Aircraft.***

Take simulation, for example. Both CFD (computational fluid dynamics) and FEA (finite element anaylsis) is critical to testing everything from aerodynamics to structural integrity. However, Elixir had about five on-staff employees and didn't need to invest the tens of thousands of dollars in dedicated CAE software licenses it would only use at very specific intervals. " We didn't want to have to buy all the software for a lifetime," Leopold-Leger says. "If we need simulation software for three or six months, we liked the idea of just being able to take it for the time we needed it."

Dassault Systèmes' 3DEXPERIENCE "Engineered to Fly" platform is also tailored to small- and mid-sized companies in this space, providing cloud-based access to the latest engineering technologies to manage the design workflow as well as the documentation process. Elixir was also sold on the cloud's ability to easily allow colleagues and suppliers to collaborate in a secure environment.

— B. Stackpole

SOLIDWORKS 2016 Unveiled

The SOLIDWORKS 2016 release sports a smorgasbord of new capabilities — most in response to customer requests — all intended to get the CAD system out of the way to increase productivity.

"We're not about buzzwords or fancy features, we're about listening to what customers want," says Craig Therrien, SOLIDWORKS product portfolio manager, explaining that as with prior years, SOLIDWORKS 2016's enhancements fall into four categories: Design, Validate, Collaborate and Build.

Design: The user interface has been overhauled to be more intuitive and offers easier access to commands. Breadcrumbs are a major new addition in this area and one of the more popular features of the beta program, allowing users to quickly return to their previous features and location in the assembly.

Validate: Most of the validation enhancements are directly related to simulation, and boosting performance of analysis is the primary goal. Noteworthy to the release is improved control over operation sequencing, part movements and mesh quality to make detailed design simulations easier.

Collaborate: One new collaboration feature is the Mate Controller, which mimics a game controller to provide an intuitive way of creating and animating complicated assembly motions — especially those with many degrees of freedom. eDrawings has also been enhanced for improved communication of designs.

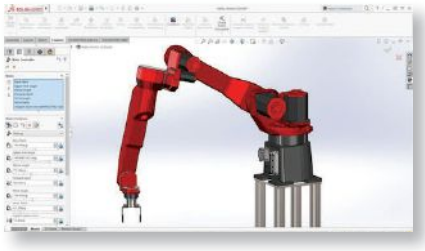

Mate Controller provides an intuitive way of creating and animating complicated assembly motions. ***Image courtesy of SOLIDWORKS.***

Build: Photorealism and rendering are front and center in SOLIDWORKS 2016 as Dassault Systèmes has integrated the Bunkspeed visualization product, which is now called SOLIDWORKS Visualize.

Read more at deskeng.com/virtual_desktop/?p=10957.

— B. Stackpole

Rapid Manufacturing with a Polite Disregard for Tradition

Tech-driven injection molding, CNC machining and 3D printing for those who need parts *tomorrow*

Proto Labs uses proprietary software and a massive compute cluster to accelerate manufacturing of prototypes and production parts for every industry.

Got a project? Get 1 to 10,000+ plastic, metal or liquid silicone rubber parts in as fast as 1 day.

2016 COOL PARTS CALENDAR
Request your free calendar at
go.protolabs.com/DT5GC

ISO 9001:2008 Certified | ITAR Registered
Major Credit Cards Accepted |

ADDITIVE MANUFACTURING
CNC MACHINING
INJECTION MOLDING

proto labs®
Real Parts. Really Fast.™

Extreme Computer Engineering at D-Wave

Engineers at D-Wave Systems explain how computational simulation and modeling help them to get the job done.

BY MICHAEL BELFIORE

It's not often that you hear the phrase "superconducting integrated circuits" dropped casually in conversation, but Jeremy Hilton, vice president of Processor Development at D-Wave Systems Inc. manages to make it sound almost commonplace. The super-cooled, zero-resistance electrical circuits are his stock in trade.

D-Wave is at the cutting edge of computer technology, producing computers that run on the principles of quantum mechanics. In D-Wave's machines, which reportedly cost more than $10 million a piece (D-Wave doesn't disclose pricing details), each quantum bit, or qubit, can register as both a zero and a one at the same time. This potentially allows for exponentially faster calculations than conventional computers.

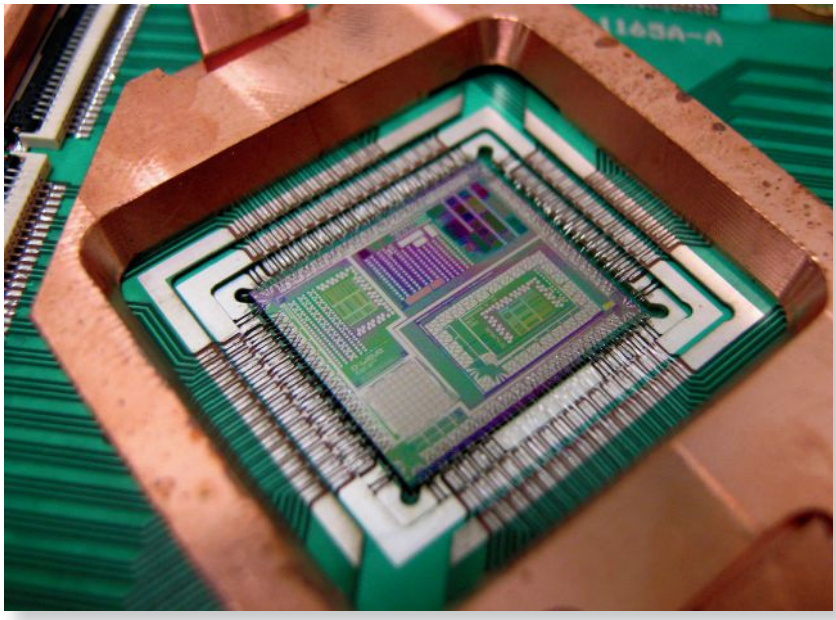

Close-up view of D-Wave's quantum computer chip. ***Photo courtesy of D-Wave Systems Inc.***

For now, D-Wave's machines are used for specialized artificial intelligence research by the likes of Google and NASA, and are still far from commonplace. That's partly because of the extreme environment that D-Wave's quantum computer chips must operate within, which adds to their expense. And engineering that environment is a major part of the challenge facing D-Wave's engineers as they seek to improve their machines. Fortunately, they have help, in the form of computational tools from ANSYS, as well as help from ANSYS software engineers.

"For us, this is a unique application," says Julius Saitz, Ph.D., lead application engineer at ANSYS. He says he hasn't seen anything like it in the 10 years that he has been with his company.

Extreme Cold, Extreme Shielding

The qubits of D-Wave's computers are formed from niobium, and each encodes a so-called qubit of information as a magnetic field.

The field can be externally manipulated to create the input and output that forms the basis of calculations. The big difference between the ordinary bits of conventional computers and qubits is that qubits don't have to choose between the two states that form the ones and zeros of computer data. Instead, they can exist in a state of superposition — that is to say — both states at once.

But the slightest disturbance can spoil the calculation. And "slightest" means even the tiniest level of electromagnetic radiation or heat. That's why the single, 4x4mm superconducting integrated circuit at the heart of each machine must be cooled to 10 millikelvins — colder than the coldest object in interstellar space — and electromagnetically shielded so that it experiences magnetic fields of no more than one nanotesla, about 50,000 times less than that which exists in a typical room, says Hilton.

"Temperature," says Hilton, "you can think of as noise. In these systems, the colder you can operate the processor, the lower the noise is, and the bigger role that the quantum mechanical effects will play." The circuits at the heart of the D-Wave chip, says Hilton, become superconducting at about 9 kelvins, but even colder is better for measuring those minute magnetic fields in the qubits through any background noise.

"The binary information that we're representing and then processing you can think of as little magnetic dipoles. They're like tiny little magnets," says Hilton. But magnets are sensitive to magnetic fields around them, driving the requirement for electromagnetic isolation.

Hilton says that best practices already existed for extreme cooling and electromagnetic isolation before D-Wave came along in 1999, but D-Wave's operating environment is larger and more complex than is required in typical laboratory settings. That thumbnail-size chip, for example, is considered large for a superconducting environment.

Complexity is also introduced by the D-Wave quantum chip's need to be controlled by conventional computers that can be programmed outside of the D-Wave refrigerator. The output of the chip is also read and displayed by the conventional computers. The interface between the classical and quantum worlds takes the form of control lines running from the quantum processor to the conventional electronics. "Those control lines are coming from room temperature," says Hilton, "where room temperature electronics are generating the signals that we need to get on to the [quantum] processor to control and process information."

Engineering that highly specialized environment becomes ever more challenging with each generation of chips. The denser and larger the chips become, the more sensitive they are to electromagnetic fields. And as complexity of both the chips and their operating environment increases, so too does the importance of precisely quantifying the environment. "The more we can understand what that volume looks like, the more flexibility it gives us from an engineering perspective to optimize the packaging that the payload [i.e., the chip] is going to go into," says Hilton.

Increased flexibility could lead to more chips in the same cooled and shielded volume, for example, or more options for integrating the control lines that come in from outside the environment.

Modeling a Unique Environment

To refine the operating environment further for the new chips, the D-Wave team relies on software modeling and simulation tools from ANSYS along with experimentation.

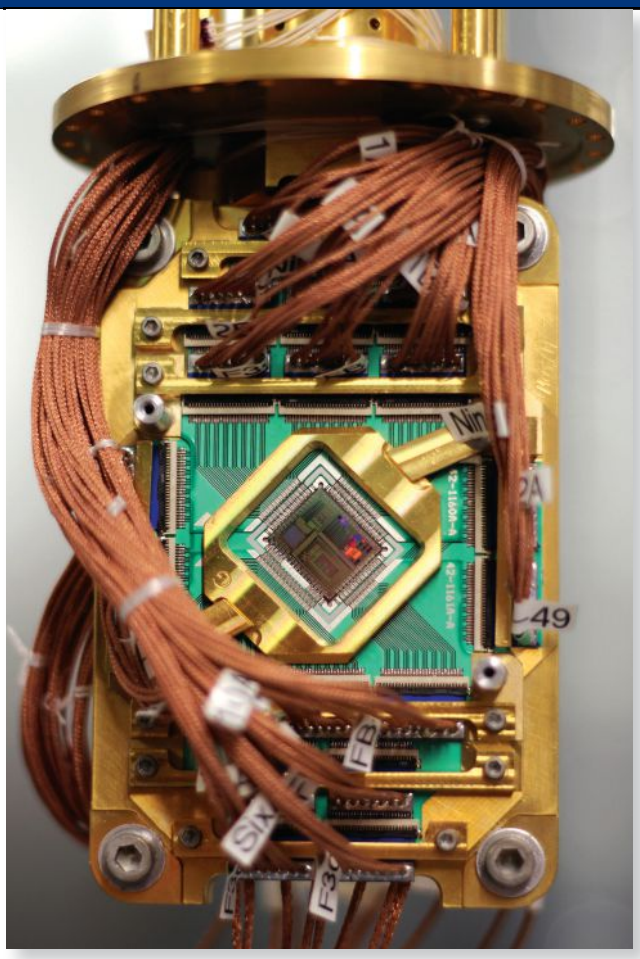

D-Wave's quantum computer chip in the housing in which it will be cooled and electromagnetically shielded. *Photo courtesy of D-Wave Systems Inc.*

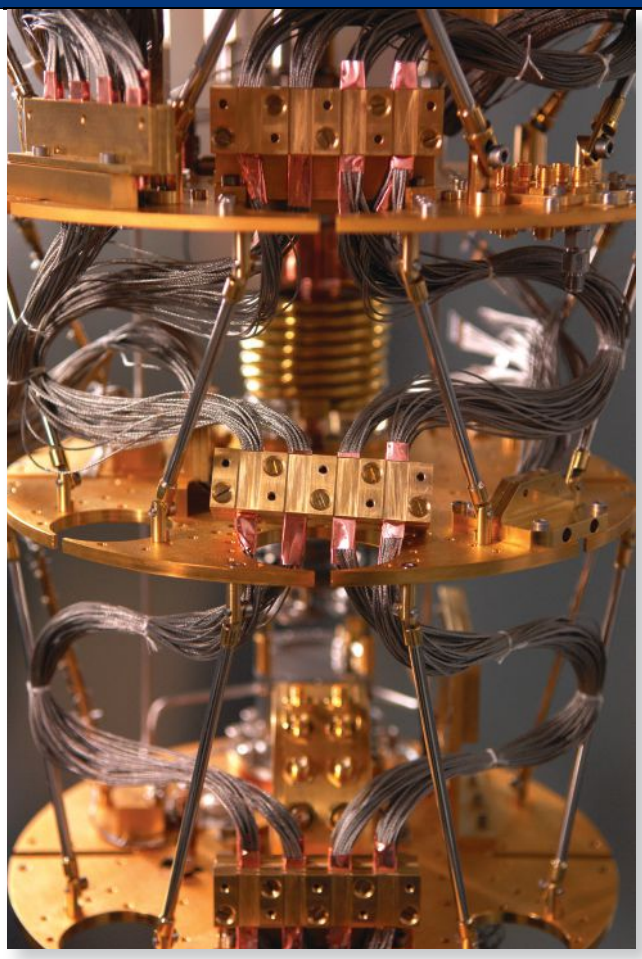

Detail of the refrigeration system used to cool the single chip at the heart of D-Wave's quantum computer. *Photo courtesy of D-Wave Systems Inc.*

Until recently, the team has leaned more on experimentation than simulation, but experimentation gets expensive. "We've done a lot of experimentation," says Hilton. "Now we need to really engineer this more seriously, and so it's time to partner with ANSYS and their products and see what we can learn using some powerful software."

In February, Hilton and his team reached out to ANSYS for help in modeling the environment inside the D-Wave machine. The ANSYS Maxwell product became their weapon of choice for conquering the challenge of precisely modeling the magnetic fields at play around the D-Wave chip, and for informing the design of the electromagnetic shielding around it.

Saitz, who has been working with the D-Wave engineers, notes that ANSYS Mechanical is also an important tool for this application. "The modeling that they requested and we demonstrated was a combination of Maxwell and ANSYS Mechanical," he says.

The two packages work together to characterize the D-Wave chip's operating environment as the circuits within the chip become superconducting in response to the temperature dropping in the refrigerator that houses it. ANSYS Mechanical simulates the changing properties of the chip over time and passes that data to Maxwell, which models the changing electromagnetic field. "When the material properties change, then the magnetic field results are different," says Saitz.

At the same time, the changes in the magnetic field also affect temperature. "If there is a change in magnetic field, that will induce currents in any conductive material," says Saitz. The currents, in turn, can result in the generation of heat, the effects of which the D-Wave engineers want to minimize as much as possible. Simulating these interactions in detail helps the engineers reduce the number of costly experiments they have to make."

The D-Wave team is also using the ANSYS HFSS package for high-frequency modeling of the processor. "We work with very high frequencies in our readouts so we can extract information from the processor as quickly as possible," Hilton says.

Saitz cites the tight coordination between the software packages as one advantage to using a single software developer for all of them. "The whole simulation process is streamlined under an environment we call Workbench," he says.

Hilton and his team run these memory-intensive simulations on Linux-based servers with 512GB of RAM onboard and that are dedicated to the task at D-Wave's headquarters in Burnaby, British Columbia. Big jobs typically run overnight, with some occasionally running for two days.

In the near future, the team also plans to use ANSYS software to model the thermodynamics of the refrigerator structure and determine how the metal in the refrigerator generates thermal currents as it cools.

All of which should help D-Wave scientists and engineers build more complex chips at a faster pace. D-Wave's first-generation chip used 128 qubits to process data. The current generation chip uses 512 qubits, and the next generation, chip, due on the market within the next year, will be a 1,000-qubit processor. Hilton and his team are now working on a fourth-generation chip with a whopping 2,000 qubits that he hopes will help revolutionize the field of machine learning.

"The whole field of machine learning is very hot right now. It's moving rapidly, but because of the limits of classical computing, in order to scale and do more, the algorithms are tending more towards keeping the underlying models as simple as possible," says Hilton. "With D-Wave, what we're saying is: Let's go back to that core model and add some complexity to that model that classical algorithms are having to throw out. And now that we can solve those problems, what does that mean for learning?"

There's only one way to find out. DE

External housing for the D-Wave computer. The refrigerator used to cool the quantum computer chip comprises most of the volume. ***Photo courtesy of D-Wave Systems Inc.***

Michael Belfiore's *book* The Department of Mad Scientists *is the first to go behind the scenes at DARPA, the government agency that gave us the Internet. He writes about disruptive innovation for a variety of publications. Reach him via michaelbelfiore.com.*

INFO ➜ ANSYS: ANSYS.com

➜ **D-Wave Systems Inc.:** DWaveSys.com

Processor Progress

Intel's new CPU with integrated graphics moves into territories traditionally dominated by the GPU.

BY KENNETH WONG

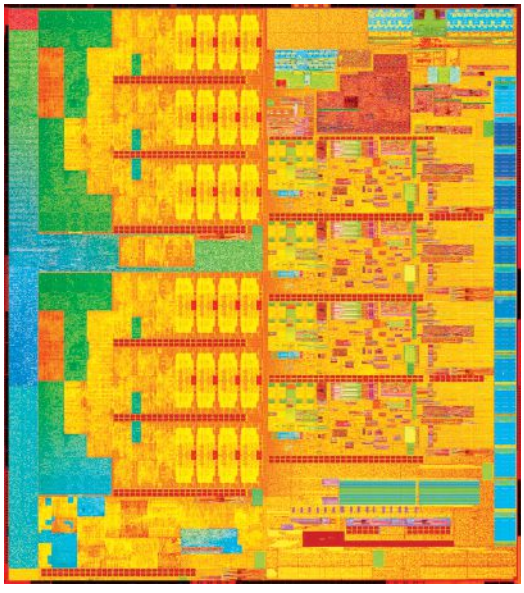

Die shot of Intel Xeon E3-1200 V4 with integrated graphics. *Images courtesy of Intel.*

For a long time, NVIDIA was content to be the gaming graphics company, the one supplying hardware that makes the cliffs and waterfalls in "Tomb Raider" more realistic and the explosions in "Battlefield" more spectacular. But about five years ago, NVIDIA began redefining itself by pitching its products as co-processors for general-purpose computing. The strategy allows the GPU (graphics processing unit) maker to compete in the high-performance computing (HPC) market, what was traditionally the CPU's territory. NVIDIA's R&D efforts to build its new persona as a legitimate HPC vendor are not trivial. It reworked its GPU architecture so researchers could deploy GPU-powered clusters in massively parallel operations, from nuclear activity simulation to climate analysis and human genome mapping.

Now it seems the CPU maker Intel is laying the groundwork to redefine itself as more than a CPU maker. The launch of new Intel CPUs with integrated Iris Pro graphics at Computex in September in Taipei, Taiwan marked Intel's strategy to edge into what has traditionally been the GPU's market. "Intel is beyond CPU performance. We've been working on graphics and imaging over the last few years," says Kirk Skaugen, senior vice president and general manger, Client Computing, Intel. Jim Blakley, general manager of Visual Cloud Computing at Intel's Data Center Group, describes the new Intel Xeon E3-1200 V4 processor as "a CPU and GPU on the same chip." The graphics processing or GPU component in Intel's new chip is the company's Iris Pro integrated graphics technology.

Integrated Graphics vs. Add-on GPU

Intel HD Graphics, Iris and Iris Pro graphics are all designed to provide a visual computing boost in gaming and video — straight from the CPU's architecture. "Integrated graphics have a lower bill of materials, smaller footprint and lower power envelope," says Skaugen.

In its product sheet for Xeon E3-1200 V4, Intel writes: "The tight integration of compute and graphics in the same processor increases performance and density, while reducing power consumption and data movement." The implication is, because the graphics component is built into the chip itself, Intel's integrated graphics technology is a better alternative to a discrete GPU, which needs to communicate with the CPU through a connector.

NVIDIA offers the Quadro product line for workstation users running mainstream CAD software, and the higher-end Tesla cards for HPC and simulation software users. In its current incarnation, Intel's Iris Pro might go head-to-head with NVIDIA's Quadro product line in some applications (for example, video editing, 3D content viewing, and 3D modeling); however, it's not reaching into the HPC and simulation software user markets served by NVIDIA Tesla. "Iris Pro graphics is not meant to compete in the high-end GPU market. The performance of Iris Pro is roughly equivalent of a mainstream graphics card," Blakley says.

Most professional workstations — even entry-level brands — come equipped with discrete graphics cards; therefore, once Intel CPUs with integrated graphics become the norm in processor choices, workstation users would have two graphics options: through the CPU's integrated graphics and through the discrete GPU. Workstation users could therefore put both to use at the same time. For example, they may use the CPU's integrated graphics for video editing and assign the discrete GPU to take on the more demanding CAD visualization. But re-

lying on the CPU's integrated graphics to the exclusion of the discrete GPU may not be the best course of action.

"Workstation users running CAD applications tend to want the best graphics, best visualization, and that's what they get with the discrete graphics card. You can't usually get the same level of graphics performance with integrated graphics," says Scott Hamilton, vertical marketing strategist, Dell. Hamilton pointed out that integrated graphics tend to use less energy; therefore, it could be the preferred option for mobile workstation users doing light to medium modeling work on the road and need to conserve battery life.

"Discrete GPUs in the professional space are here to stay. Discrete GPU makers AMD and NVIDIA tweak their drivers all the times to make sure the ISVs (independent software vendors) get the best performance from their products. Iris Pro could replace the discrete graphics card in some entry-level applications. But for integrated graphics to become a viable alternative to discrete graphics, Intel would have to make the same investment in ISV certifications, recruiting ISVs to qualify their drivers, adding performance modes and plug-ins, and so on," says Rahul Tikoo, general manager for Workstation End User Computing, Dell.

Tikoo pointed out Dell does offer some desktop tower models without a discrete GPU. "Our suspicion is, most of those buyers already have a discrete graphics card or have a way to acquire it on their own. So it doesn't mean they're running their applications solely with integrated graphics. Most likely, they're adding the GPU to the machine on their own," he says.

Supporting Integrated Graphics

The appeal of the GPU as a coprocessor expanded when NVIDIA developed CUDA (compute unified device architecture), the company's proprietary GPU-based programming language for writing massively parallel applications. Nurturing the CUDA community took time. About five years after its debut, CUDA is finally finding a home among some simulation and visualization software vendors. This NVIDIA-ISV partnership was crucial to the company's strategy to redefine the discrete graphics card as a number crunching, parallel-processing engine. The GPU's value proposition also increases as more mainstream 3D software vendors began to implement GPU-accelerated visualization.

Intel argues supporting integrated graphics will be much more straightforward. "If the application is written to use OpenGL or DirectX, then Iris Pro works without any additional programming effort. NVIDIA also supports OpenGL and DirectX.

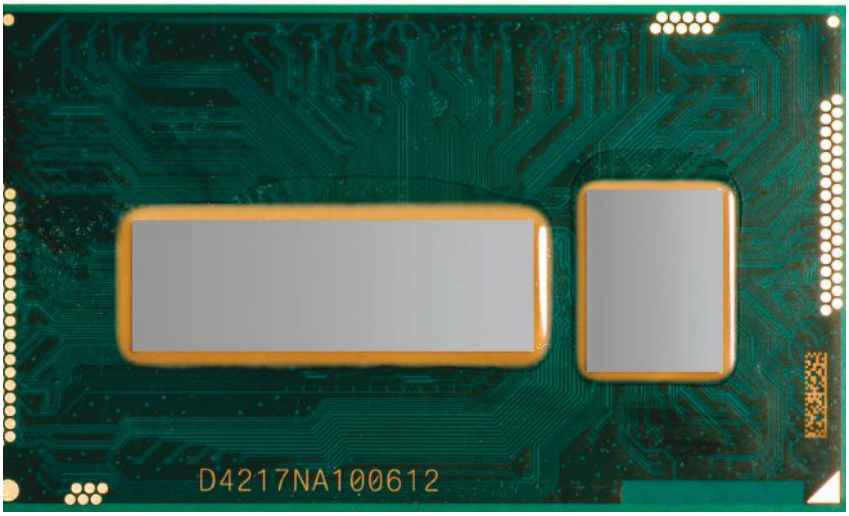

The integrated graphics component for fifth-generation CPUs is pictured.

IBM Research Advances Post-Silicon Future

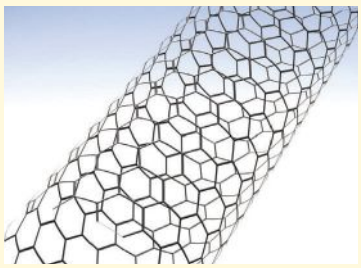

IN what it calls a "major engineering breakthrough," IBM Research scientists demonstrated a new way to shrink transistor contacts without reducing performance of carbon nanotube devices, opening a pathway to dramatically faster, smaller and more powerful computer chips beyond the capabilities of traditional semiconductors, according to a press release. IBM's breakthrough overcomes a hurdle that silicon and semiconductor transistor technologies face when scaling down. In any transistor, two things scale: the channel and its two contacts. As devices become smaller, increased contact resistance for carbon nanotubes has hindered performance gains. These research results could overcome contact resistance challenges all the way to the 1.8 nanometer node – four technology generations away.

Carbon nanotube chips could greatly improve the capabilities of high performance computers, according to the company. Silicon transistors, tiny switches that carry information on a chip, have been made smaller year after year, but they are approaching a point of physical limitation. Shrinking the size of the transistor — including the channels and contacts — without compromising performance has been a challenge troubling researchers for decades.

IBM has previously shown that carbon nanotube transistors can operate as switches at channel dimensions of less than 10 nanometers — the equivalent to 10,000 times thinner than a strand of human hair and less than half the size of today's leading silicon technology. IBM says its new contact approach overcomes the other major hurdle in incorporating carbon nanotubes into semiconductor devices, which could result in smaller chips with greater performance and lower power consumption.

For the full research results, see the October 2, 2015 issue of *Science* (DOI: 10.1126/science.aac8006).

— Source: IBM Research press release

These are graphics programming APIs. If you want to do things that are more compute-oriented, then we support OpenCL programming on the GPU. This is the direct comparison with CUDA. Whereas CUDA is proprietary NVIDIA programming model, OpenCL is industry standard," says Blakley.

In an Intel-produced podcast titled Chip Chat, Blakley said the Xeon E3-1200 V4 offers "substantial performance over previous generations." He says: "We get up to almost 2x in graphics improvement overall, almost 2x in media processing and video transcoding" (Chip Chat, June 10, 2015).

Though Intel has never confirmed it officially, the chipmaker is believed to follow the tick-tock rhythm in its technology releases — tick is a refinement; tock is a major architecture overhaul. The Xeon E3-1200 V4 falls into the tick slot, which may explain the ISVs' subdued reaction. Software makers tend to devote more R&D efforts to align their products with the chipmaker's tock chip releases.

The anticipated tock for the near future is the new chip architecture codenamed Skylake. These new processors are expected to deliver, in Intel's estimates, 1.6x better productivity, 6x faster video conversion, 2x faster media editing, and 13x better graphics (compared to five-year-old PCs). Designed with IoT (Internet of Things) devices and remote computing in mind, Skylake processors promise superior user experiences in all-in-one PCs and mobile devices.

Redefining the CPU, the Xeon Brand

In early September, Intel began sharing details about its upcoming 6th generation Intel Core processor family, codenamed Skylake. It's described as "the company's best processors ever." Intel writes: "The launch marks a turning point in people's relationship with computers."

In the same announcement, Intel reveals its roadmap to bolster its integrated graphics. "In the coming months, Intel plans to deliver more than 48 processors in the 6th Gen Intel Core processor family, featuring Intel Iris and Iris Pro graphics, as well as Intel Xeon E3-1500M processor family for mobile workstations and 6th Gen Intel vPro processors for business and enterprises," the company states.

"Going forward, we'll add substantially better improvement to Iris Pro. We'll be developing more and more GPU features and continuing to increase its performance," says Blakley.

A month earlier, Intel also wrote in a blog post that they have "plans to bring the power of Intel Xeon processors to notebook PCs for the first time." Xeon is usually associated with professional desktop workstations, but the new Intel Xeon Processor E3-1500M v5 product family, according to Intel: "[Is] based on the next-gen Skylake architecture and they will deliver high precision computing horsepower in notebook form factors … Intel Xeon-based mobile workstations will have key features such as error-correcting code memory that automatically detects and repairs errors on-the-fly that cause data corruption and system crashes for peace-of-mind reliability."

In its early incarnation, Intel's integrated graphics competes with the discrete graphics card in narrowly defined entry-level fields. Currently, the CPU with integrated graphics is more a complementary technology than a rival to the GPU. The combo offers users the chance to split his or her graphics workload between two computing resources. But if Intel ramps up its R&D efforts to bring the performance of integrated graphics to a level that's comparable to a discrete GPU, and if users begin exploring its potential to do general-purpose computing, the chipmaker can seriously disrupt the established market dynamics in visual computing and HPC markets. DE

Kenneth Wong *is* Desktop Engineering's *resident blogger and senior editor. E-mail him at kennethwong@deskeng.com or share your thoughts on this article at deskeng.com/facebook.*

INFO ➜ **Dell:** Dell.com

➜ **Intel:** Intel.com

➜ **NVIDIA:** NVIDIA.com

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

Help for HPC in the Cloud

Offloading simulation to the cloud requires help from specialists.

BY BETH STACKPOLE

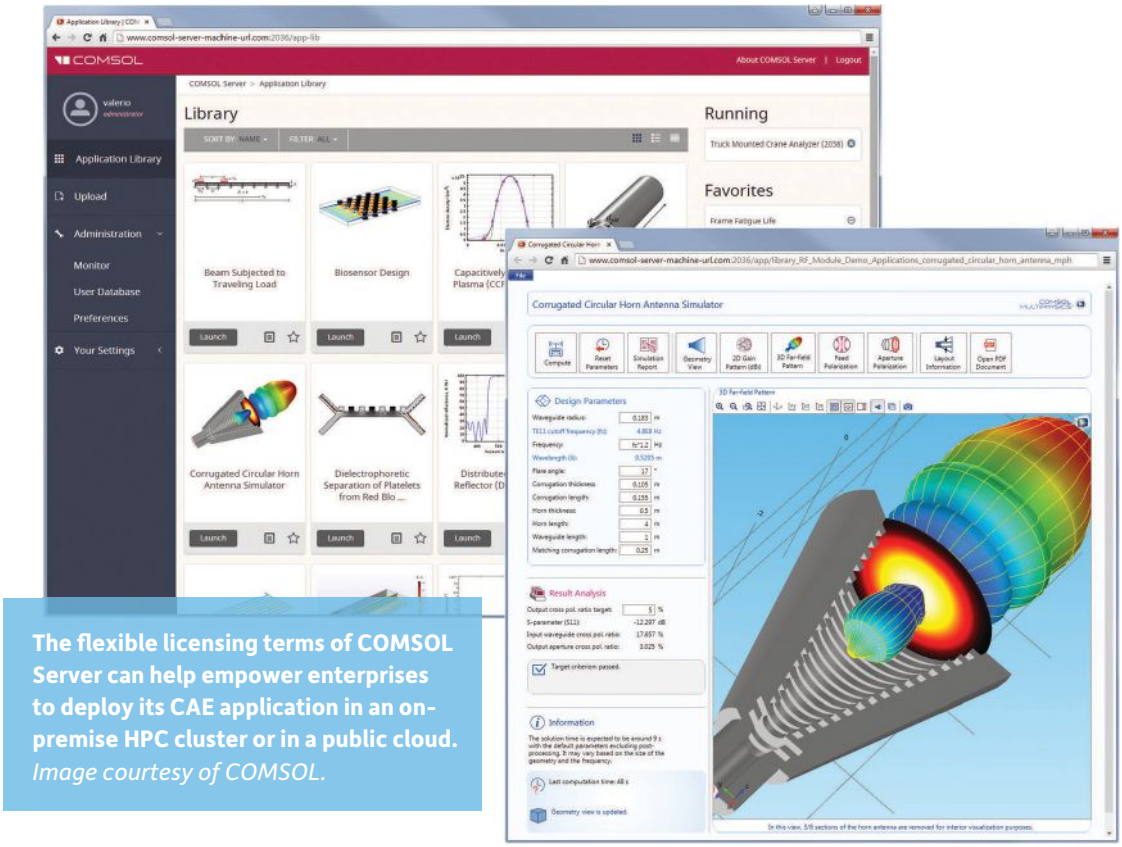

The flexible licensing terms of COMSOL Server can help empower enterprises to deploy its CAE application in an on-premise HPC cluster or in a public cloud. *Image courtesy of COMSOL.*

While the cloud is quickly becoming the platform of choice for many enterprise applications, it's not a slam-dunk for deploying high-performance computing (HPC) environments, especially for organizations light on internal cloud and HPC expertise.

For general business tools like sales or human-resource related applications or even mission critical back-office platforms such as ERP (enterprise resource management), the cloud is increasingly the go-to approach as companies look to capitalize on pay-as-you-go pricing schemes and the ability to be up and running quickly without a lot of IT handholding. According to *Computerworld's* 2015 Forecast, 42% of IT decision makers planned to increase spending on cloud computing in 2015, the greatest spike among enterprises with over 1,000 employees.

All flavors of cloud deployment are feeling the love from the mainstream enterprise. The Cisco Global Cloud Index found that of the total cloud workload, 28% is projected to be Infrastructure-as-a-Service (IaaS) workloads, 13% will be Platform-as-a-Service (PaaS), and the most popular deployment model will be Software-as-Service (SaaS), owning 59% of the workloads by 2018.

Yet when it comes to spinning up HPC resources in the cloud, the offerings are far less mature and the customer uptake much slower. Public cloud platforms like Amazon Web Services (AWS), Microsoft Azure and IBM BlueMix are far less optimized for the more demanding and specialized requirements of HPC environments tuned to run engineering applications such as simulation and optimization. Moreover, some of the barriers slowing down mainstream cloud deployments — things like security concerns, complications integrating legacy systems and unexpected costs — are also proving to be hurdles to effectively establishing cloud HPC as a go-to resource for engineering organizations.

"If you look at offerings like Ama-

zon or Azure, you see most of the cloud layers do have a solid infrastructure for business computing," says Srikanth "Sam" Mahalingam, chief technology officer for HPC and Cloud Solutions at Altair, makers of CAE software. "However, when it comes to technical computing, there are other more demanding needs that are still not addressed."

The most significant gap in current public cloud offerings' ability to support robust HPC environments is the lack of high-speed interconnect fabrics. Traditional HPC clusters have integrated technologies like InfiniBand and Intel's Omni-Path architecture to ensure high levels of communication within the system as well as between systems. This kind of low-latency communications is critical for running large-scale simulations and optimizations in a cloud HPC environment, and to date, few providers have built out these capabilities as part of their public cloud platforms, says Dominic Daninger, vice president of Engineering for Nor-Tech, a systems integrator specializing in HPC implementations.

"Many computational fluid dynamics (CFD) or heat modeling applications that have a lot of intermodal chat going on really benefit from the very low latency fabrics like InfiniBand," Daninger says. "Very few cloud providers can offer that and if they do, there's a substantial cost to it."

Security is another big obstacle slowing cloud-based HPC adoption for many organizations, particularly those who see their HPC environments — tightly coupled with advanced simulation software — as a strategic, competitive edge. "There's a reluctance within many engineering groups [to go to the cloud] because they know their IP (intellectual property) is the company's gold and they don't want to put it outside of the company," he says.

Knowledge is HPC Power

Unlike business systems, which users can easily spin up in no time just by renting compute cycles and storage on AWS or Azure clouds, there's a lot more complexity involved in configuring an HPC environment — even in the cloud. "It's not as easy as it looks," Daninger says. "You have to look deeper and be knowledgeable enough to know the types of jobs you're looking to run in the cloud."

When considering a cloud platform for HPC, organizations will need to have as much in-house expertise as they would if they were configuring an on-premise HPC cluster. For example, they will still need to determine the right operating system, cluster management software, post processing visualization capabilities, even what data is moving back and forth to get the job done. Beyond HPC domain expertise, they need to be versed in all of the small details related to the specific cloud offering and cloud deployment — for example, knowing the costs involved in moving data up to the cloud and more importantly, back down to on-premise engineering systems.

"Many companies don't have the expertise to spin this up on their own — they need specialists like us to help them do that," Daninger contends.

Given the complexities and the minefield of potential hidden costs, experts say cloud-based HPC is not universally the best value proposition at this time, despite on-going reports to the contrary. Larger firms that have already invested in data centers and have HPC expertise on-site are still better served keeping HPC infrastructure on site, says Altair's Mahalingam. However, the story changes for small and mid-sized companies.

"It's really about a price/performance tradeoff," Mahalingam says. "Since smaller players don't have the HPC infrastructure, they're better off going to a public cloud with their jobs running a little slower than outlaying a capital expense for systems with high-speed interconnects," he says.

Cloud-based HPC also makes sense

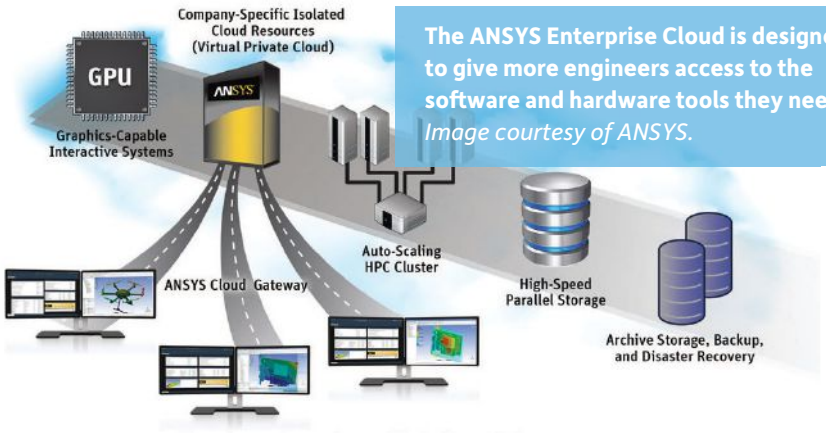

The ANSYS Enterprise Cloud is designed to give more engineers access to the software and hardware tools they need. *Image courtesy of ANSYS.*

for both large and small companies who have periods of peak compute requirements to accommodate simulation cycles for a particular project or a critical stage of the development cycle, according to Bill Mannel, vice president and general manager, HPC and Big Data Solutions for HP Servers at Hewlett-Packard. However, again, it's not as straightforward to flip a switch on cloud-based capabilities for a short-term HPC burst as it is to turn up the processing volume for a CRM (customer relationship management) or mainstream business application.

That Was Then, This is Now

The limitations surrounding cloud HPC may be short lived, however. Cloud providers, HPC specialists and CAE software vendors recognizing the opportunity and are working hard to make cloud-based HPC capabilities more palatable for a larger swath of users. Microsoft, for example, just recently announced it will offer professional graphics applications and accelerated computing capabilities on Azure with NVIDIA GPUs (graphics processing units), staking claim to being the first cloud platform to provide NVIDIA GRID 2.0 virtualized graphics. By deploying the latest version of NVIDIA GRID in its N-Series virtual machine offering, Azure delivers on the ability for engineers and designers to visualize complex, data-intensive designs and simulations from anywhere, Microsoft officials say.

For its part, HP offers what it bills as a "self-service" HPC package, but it's designed for organizations that want an appliance-like approach to HPC for an on-premise private cloud, not for the public cloud. Mannel also says HP doesn't offer turnkey HPC solutions for the cloud or otherwise; rather its approach is to leverage reference architectures and consulting services to create the optimal HPC environment for individual customers, most of which are larger companies. "You may not make the right choices without some level of consultation about what works best for your workload," he says.

Software Vendors Enter the Ring

Many of the simulation ISVs (independent software vendors) are also taking a stab at making cloud-based HPC more accessible, particularly for their specific CAE offerings. At COMSOL, cloud-based HPC isn't currently in top demand among its customers, but the company believes that will change over time, according to Phil Kinnane, COMSOL's vice president of Business Development. As part of its effort to support the shift, COMSOL now offers a more liberal floating licensing policy with its COMSOL Server offering, which now makes it possible to launch an instance on AWS. "You can use the software on a workstation and distribute it over two cores or use it in the cloud and distribute it on those resources," he says.

The ANSYS Enterprise Cloud, launched in May, empowers the transition of simulation workloads to the cloud as a turnkey process. Powered by the ANSYS Cloud Gateway and the ANSYS reference architecture, the ANSYS Enterprise Cloud is delivered as a single-tenant solution through a dedicated corporate account on the AWS public cloud. The platform, which can be managed either by internal IT experts or by ANSYS certified service partners, offers secure storage and data management.

Altair is now offering its HPC solutions on AWS and Microsoft Azure, and is even promoting a free trial of CAE on the cloud as part of a Cloud Challenge with AWS and Intel to showcase the benefits of this kind of deployment. In addition to its fully configured physical HPC/CAE HyperWorks Unlimited appliance that can be used in a private cloud, Altair has recreated a similar environment for use on the public cloud, accessible by engineers in small companies through a Web browser.

Having all of the components packaged up in a virtual or physical appliance is still critical given the complexity of the HPC environment. "For example, explicit solvers need computational cores more than memory, but implicit solvers need more memory than cores," Mahalingam says. "If you look at the appliance stack, there are a lot of layers and there is a right configuration of nodes for different kinds of jobs. We have optimized the whole stack, eliminated complexity, and can also serve as the single point of contact for support." DE

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

INFO
- **Altair:** Altair.com
- **Amazon Web Services:** AWS.Amazon.com
- **ANSYS:** ANSYS.com
- **COMSOL:** COMSOL.com
- **HP:** HP.com
- **IBM:** IBM.com
- **Microsoft Azure:** Azure.Microsoft.com
- **Nor-Tech:** Nor-Tech.com

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

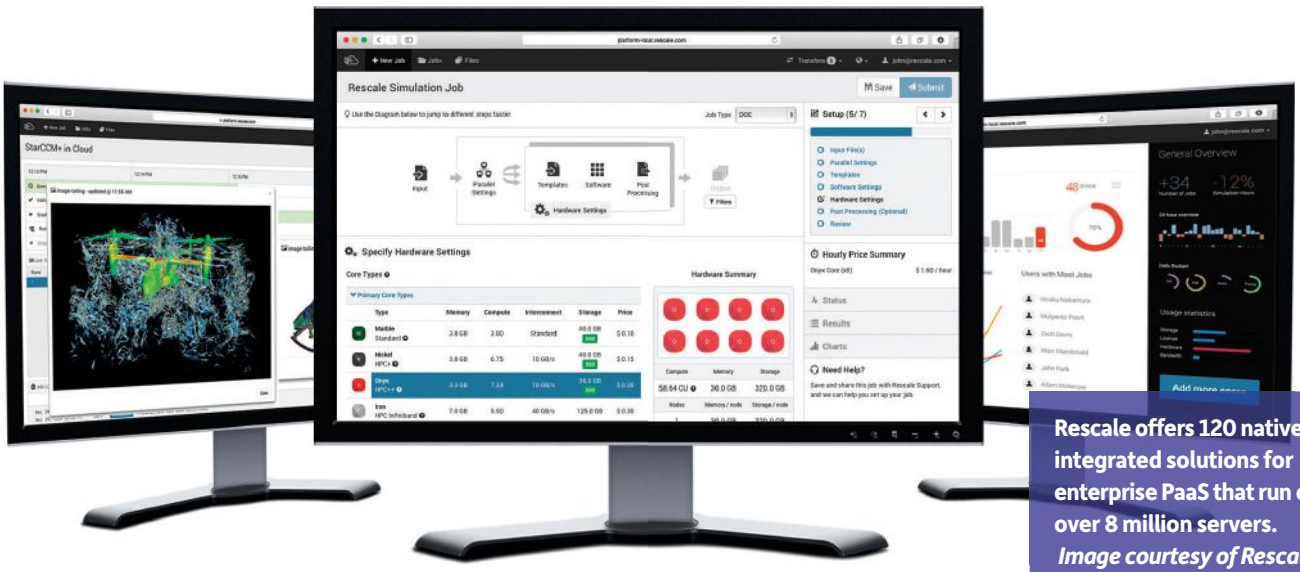

Rescale offers 120 natively integrated solutions for enterprise PaaS that run on over 8 million servers. ***Image courtesy of Rescale.***

Simulate Anytime, Anywhere

SaaS, IaaS and PaaS providers let engineers use CAE in the cloud.

BY JESSICA LULKA

The way that engineers simulate is changing. In the last few years, more technology and software has emerged making the necessary toolsets more accessible for non-analysts. But it's not just the software that's changed; the way companies are structuring their offerings has also shifted. Engineers and firms can now purchase software and hardware as a service — offering flexibility, accessibility and mobility.

Because they can use software, servers and the cloud as little or as often as they need, engineers can get a taste of programs or toolsets that might not be accessible otherwise. This has become more common for CAE platforms as engineers are bringing simulation further forward in the design process.

Acronyms, Acronyms, and More Acronyms

Within the industry, "as-a-service" offerings are typically categorized into three main silos: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). While all of these options leverage the cloud, the type of technology that they provide differs. According to Computenext, these acronyms can be described with the following definitions:

- **SaaS:** A Web browser providing a point of contact for software on servers.
- **PaaS:** An environment to run an operating system and servers for deploying specific software programs.
- **IaaS:** Server and cloud computing access via a dashboard and/or API (application programming interface).

These different configurations offer organizations a way to access more power for simulation not only when it's needed, but also at a more affordable price point.

"What we often see is that simulation is a very volatile thing that [engineers deal with]. Compared to CAD, which they use very regularly, simulation might spike during a development project and [be followed] by several weeks where it is not used at all," says David Heiny, managing director of SimScale. Despite the varying levels of support and technology, these "as-a-service offerings" all provide engineers and teams with an option to quickly implement extra simulation resources when needed.

"The main priority of simulation software should always be getting the job done reliably, robust and fast. There's also a huge

When To Use PaaS/Cloud CAE

Before looking at as-a-service providers, it's important to recognize when an outside service beyond software needs to be implemented. According to Rescale, determining if PaaS/IaaS is beneficial can be narrowed down into three stages:

1. When a company identifies that they have more demand than HPC resources and they are losing potential product advantages.
2. When a company has a consistent need for cloud resources and uses a combination of reserved cloud HPC, on-demand HPC, as well as still using their internal cluster.
3. When a company realizes and acts on the full benefits of a PaaS and migrates their entire workload to a cloud environment.

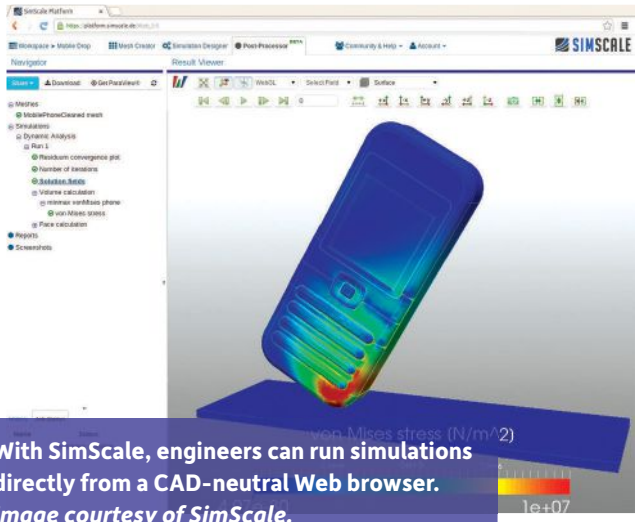

With SimScale, engineers can run simulations directly from a CAD-neutral Web browser. ***Image courtesy of SimScale.***

technical overhead, so when we started SimScale, we saw that the Web browser will evolve to something that is capable of providing access to a simulation engineering workflow," Heiny says. "It [also] really comes down to providing access to simulation methods and capacities in a more convenient manner."

In using the Web browser as the basis for its software, Heiny says it not only offers mobility, but is also a more intuitive format for engineers that enables them to start simulating quickly. To access the resources SimScale provides, users simply create an account, upload a model and start running simulations.

Beyond the Software

However, running simulations requires more than software. Engineers still need access to computing power that can physically calculate results. With providers such as UberCloud and Rescale, teams can use servers, the cloud and software all at once — without needing a massive investment.

Simulation: There's an App for That

Online applications are also helping engineers access simulation quickly on the Web. Here are a few examples of companies offering programs for writing and/or deploying simulation applications:

- **Altair:** CFDCalc offers flow solutions for a specific problem class. Users can run simulations for heat sinks, centrifugal pumps and SMX mixers without needing a CAD model. Altair.com
- **COMSOL:** Application Builder enables users to turn specific simulations into sharable, Web-based models. COMSOL.com
- **Comet Solutions:** The company's SimApp authoring software lets engineers create Web-deployable simulations. Cometsolutions.com
- **EASA:** With SmartApps, users can automate processes and workflows for different levels of model abstraction. EASAsoftware.com

"A primary differentiator of PaaS for simulations is the immediate ability to scale out high-performance hardware," says Sarah Deitz, business development manager at Rescale. "With PaaS, users have the elasticity instantly available through a unified platform."

Furthermore, these companies also bring the necessary expertise to help teams effectively use the technology – easing some of the technical overload. At UberCloud, the community pairs engineering teams with cloud providers and software vendors that best fit the team's simulation requirements. In addition to providing IaaS and SaaS, the goal of UberCloud is to help educate engineers about service-based offerings as well as the benefits of cloud-based simulation.

Considerations

SaaS, IaaS and PaaS can offer extra simulation power, there are several considerations that engineers should take into account before deciding if a cloud-based or as-a-service solution is suited for their needs.

1) **Time and Budget.** Because simulation can fluctuate throughout the year, it's helpful to figure out what budget is available for peak job loads and how quickly simulations will need to be run. By identifying these factors, using a pay-as-you-go, short-term model for software could be beneficial.

2) **Needed Computing Power.** Make sure the service providers can offer the bandwidth necessary for specific simulations. "CAE codes [often] require high performance computing resources, powerful CPUs, GPUs, Infiniband interconnects and more," says Wolfgang Gentzsch, president and CEO of UberCloud. "Exact hardware requirements can be requested from the ISV (independent software vendor), and information about the hardware can then be requested from the cloud provider."

3) **Options for Flexibility and Scalability.** CAE cloud computing resources should be able to adapt with the changing simulation workflow. "A component in choosing their cloud provider is that engineers are not being locked into a single infrastructure network and that their cloud provider can offer various HPC configurations and options to meet the dynamic nature of their simulation workload," says Ilea Gradel, senior account executive at Rescale.

In addition, Gentzsch notes that there are a few other industry and user concerns. These include security, licensing, maintaining control of resources, and data transfer. DE

Jessica Lulka *is associate editor of* DE. *Send e-mail about this article to DE-Editors@deskeng.com.*

INFO ➜ **Computenext:** Computenext.com
➜ **Rescale:** Rescale.com
➜ **SimScale:** SimScale.com
➜ **UberCloud:** TheUberCloud.com

Beyond the Desk

The workstation takes on a new life via virtual desktop infrastructure.

BY KENNETH WONG

Among professional engineers and designers, the network computer is not an attractive term. It's associated with a loss of control over the machine. It's a computer that sits far away from the user, under the tight scrutiny of the IT department. Even if the server room that houses the machine is only a few feet away from the user's desk, bureaucratic "red tape" keeps the user at bay. Need a RAM upgrade to make your assemblies load faster? Need to install a plug-in to make your CAD conversion easier? Submit a requisition slip and wait. By contrast, the desktop workstation, one that physically sits on the user's desk, is viewed as unfettered access to the hardware, and perhaps even a status symbol for the user.

Ben Fathi, CTO at VMWare (left) and Jen-Hsun Huang, NVIDIA CEO (right) at NVIDIA GPU Technology Conference 2014, announcing NVIDIA Grid would run on the VMware Horizon DaaS Platform. ***Image courtesy of NVIDIA.***

So why is the virtual machine, which is essentially a network computer, enjoying a resurgence among some engineering and design firms, even among some universities? (Read "PSA Peugeot Citroën Drives Virtual Desktop Infrastructure Forward," deskeng.com/de/?p=21117). The answer rests with a new generation of virtual desktop infrastructure (VDI) solutions and better middleware that seamlessly connect the user to the remote hardware. The stigma of the network computer fades away when users realize the virtual machine delivers performance and power that rivals a desktop workstation. The firms deploying virtual machines now include not only design and engineering firms but design software makers. Autodesk, a household name in CAD software, not only offers software certified for use in VDI environments; it's using VDI to test and run simulation of its own software internally.

Test-Driving Design Software on VDI

Autodesk's software lineup includes some that are known for heavy system and graphics demands. VDI opponents and skeptics often cite Autodesk Inventor and Revit as the reasons designers and engineers would insist on physical workstations. With the option to display detailed mechanical assembly assemblies and full building projects in photorealistic mode, these software titles exemplify the engineers' needs to rely on professional workstations with GPUs (graphics processing units) — sometimes multiple GPUs in a single machine — to do their design work.

In late 2013, Autodesk began offering potential customers the ability to test drive select titles remotely. The company has long offered people the option to download and test its software in a time-limited trial mode, but the remote test drive offer is different. No need to download gigabytes of installation files and no need to install the software, for that matter. The user runs the trial software from a browser or a thin client. And, Autodesk manages instances of the software in a VDI environment. Autodesk Inventor and Revit are among the titles made available in that fashion.

The AutoCAD giant's VDI is setup on NVIDIA GRID, a GPU-powered virtualization solution from NVIDIA. Previously, the graphics demands of 3D CAD proved a roadblock in running such software titles on a remote machine or a virtual machine, because only GPU-equipped workstations could deliver an acceptable degree of interactive 3D use. But that changed when NVIDIA made it possible to virtualize the GPU through its Kepler architecture.

In 2012, when unveiling the details of Kepler at the NVIDIA GPU Technology Conference, NVIDIA CEO Jen-Hsun Huang said: "I want to announce a GPU that we can all simultaneously share. Today, we're going to take the GPU into the cloud. For the

first time, we've virtualized the GPU." Now, on NVIDIA GRID hardware, virtual machines come with their own virtual GPUs.

"We can have a high degree of confidence in the recommendations we make to customers about virtualization in part because we are doing so much work in virtualized environments ourselves. We have quite a bit of in-house experience," says Anthony Hauck, Autodesk's director of product strategy, AEC generative design ("At Autodesk, VDI Inside and Out," August 2015, WindowsIT Pro, http://m.windowsitpro.com).

The Revit team "drives the application through approximately 16,000 nightly tests simulating user interaction with the product. Often these tests use building models contributed by customers to enhance the realism and accuracy of the interactions. The tests ensure enhancements to the application do not cause any regressions in the associated functionality," says Hauck.

Autodesk's Revit VDI is powered by 12 blade servers, each with dual CPUs and 128GB RAM. Autodesk generally runs 10 to 12 virtual machines per server; therefore, 12 servers have the capacity to support approximately 120 to 144 virtual machines in operation. "We have full utilization of all machines a few times a day, depending on the quantity of additional jobs submitted by development as they check in new work," Hauck says.

Zuken, a printed circuit board (PCB) design software maker, is also relying on VDI powered by Amazon Web Services (AWS) to make its CR-8000 software available for test-driving for interested parties. "CR-8000 demands a fair amount of computing power and 3D graphics. We went with AWS because we found out that some smaller providers couldn't support 3D graphics. We selected the higher-end option when we picked our setup in AWS to facilitate 3D requirements," says Craig Armenti, application engineer at Zuken.

Some of Zuken's competitors have software that works in 2D, which doesn't require intense graphics. But CR-8000 runs in 3D, allowing PCB designers to easily communicate with their mechanical counterparts. Therefore, the VMs supporting Zuken's test-drive setup has to offer the equivalent of a dedicated GPU.

"For now, we're limiting the test sessions to five at a time — five active licenses of CR-8000 at the most running on AWS," says Armenti. "We don't know when these VMs might come online. Some users test-drive the software after business hours when they're at home. So if we are hosting these VMs ourselves, we have to worry about keeping the server up 99% of the time, ensuring Internet connection. On the other hand, Amazon has servers all over the U.S., and if we ever want to expand beyond the country, AWS gives us more flexibility than our own hardware."

Splitting the GPU

Many software makers have refined their code to take advantage of the GPU's presence in most workstations; therefore office productivity software, scientific calculation software, CAD modeling programs, rendering programs and simulation software can derive benefits from what's known as GPU acceleration.

But the GPU is not always necessary for routine computing works, such as Web browsing, word processing, data entry and simple CAD modeling. When the GPU is sitting idle inside a workstation, it represents wasted computing potentials. Workload imbalance and changing peak demands often lead to some engineers craving for more GPU horsepower while others are barely touching theirs. The emergence of the virtual GPU addresses this scenario. Whereas the physical GPU attached to the workstation cannot easily be shared or reassigned, the virtual GPU can be easily shared and reassigned with a few clicks from the VDI console.

For engineers and designers working with heavy CAD models and advanced simulation programs, a designated GPU is recommended. It ensures each user gets the performance equivalent of a physical GPU even when working on a virtual machine. Workers who do occasional CAD work and do not need graphics acceleration all the time may be the ideal candidates for GPU sharing.

In the whitepaper titled "NVIDIA GRID: Graphics Accelerated VDI with the Visual Performance of a Workstation"

Virtual Machines vs. Remote Workstations

All virtual machines (VMs) are accessible remotely, but not all remote workstations are virtualized. VMs are, in a manner of speaking, theoretical computers; they do not exist as physical entities in real life. They live in the cloud (private or public) and mimic the look and behavior of physical computers, but in reality they're the computing resources from a server distributed among multiple users. A virtual desktop infrastructure (VDI) is usually a one-to-many setup: One cluster as the source for multiple VMs. In addition to the hardware, VDI also requires additional server-side and client-side software components. VMware, known for its vSphere software, and Citrix, known for its XenDesktop software, are the two biggest names in the VDI market. Commercial VDI solution vendors also face competition from open-source counterparts, like KVM (kernel-based virtual machine). In enterprise use of VDI, licensing fees for software could add up over time. The open-source alternatives could offer significant savings.

Remote workstations are physical machines made available to the user remotely. They're usually set up in a one-to-one ratio: One physical machine per one remote user. Such is the case with Dell and Teradici's PC over IP (PCoIP) solution. (For more, read "Virtualization: Access Your Workstation from Anywhere" deskeng.com/de/?p=24946). VDI offers IT managers more flexibility in reallocating the computing resources among users based on peak demands and workload changes. Remote machines may be easily deactivated or reassigned to different users; however, they cannot easily be split and shared. Remote machine setups are designed to offer users a fixed amount of computation horsepower no matter the workload.

(December 2013), Alex Herrera, a senior analyst from Jon Peddie Research, wrote: "GPU Sharing is a reasonable solution for many, but not an ideal solution for all. It can perform effectively with simple applications and visuals and support concurrent users (CCUs), but the extensive compute cycles spent abstracting complex 3D rendering will add latency and reduce performance. Furthermore, the reliance on API (application programming interface) translation means 100% application compatibility is impossible to guarantee."

At VMWorld 2015 in August, NVIDIA's rival AMD announced that it will also begin offering a GPU-sharing solution. "An AMD graphics card equipped with our Multiuser GPU technology offers consistent, predictable performance. IT managers can easily configure these solutions to allow for up to 15 users on a single GPU … Each user now has an equal share of the GPU to allow them to design, create, and execute their workflows," the company states.

In its press announcement, AMD called its Multiuser GPU technology the "first hardware-based virtualized GPU solution." This could be somewhat confusing as "virtualization" is usually interpreted as a software-driven mode. "Our solution is hardware-based. The GPU sharing component is in the GPU itself, so it's better than the software-driven approach," says Antoine Reymond, senior strategic alliance manager, AMD.

AMD further explains in its data sheet: "Software virtualization has traditionally been a limiting factor for those who want to fully utilize GPU hardware acceleration for compute tasks under Open CL. With AMD's implementation of the new Multiuser GPU, users are no longer as limited to what they can or can't do in a virtualized environment. Users will have access to native AMD display drivers for OpenGL, DirectX and OpenCL acceleration, enabling them to work with few if any restrictions."

AMD's solution is much closer to the metal (literally) than its competitor's version. "The AMD host driver adds very little burden to the hypervisor," says Tonny Wong, product manager at AMD Professional Graphics. "It is used only to set up the parameters for Multiuser GPU functionality in hardware … AMD solves virtualization in silicon using a PCI-SIG specification called SR-IOV. All that the AMD host driver is doing is telling the hypervisor that we want to split one GPU to many and set up the parameters for that split. All the GPU sharing happens under the hood and is transparent to the hypervisor. The simplicity of the solution is that, when Multiuser GPU is not configured, the product is supported as a pass-through device. Once Multiuser GPU is enabled, the GPU will appear as multiple pass-through devices,"

For knowledge workers who juggle Microsoft Office-type applications, a single AMD GPU may be shared by up to 15 users. But for typical CAD users, AMD recommends sharing the GPU among 6-10 users. To retain the desktop workstation-level performance, AMD recommends sharing among two to six users.

At the same event, NVIDIA also began promoting GRID 2.0, the second generation of its GPU-accelerated VDI solution. The company says GRID 2.0 "doubles user density over the previous version," indicating the same number of GPU can now support a larger pool of users without performance loss.

The virtual machine is distinctly different from its predecessor network computer in one aspect: Professional-grade graphics with little or no noticeable latency. This development removes one more barrier that once prevented CAD and CAM users from venturing beyond their cubicles and the office. **DE**

Kenneth Wong *is* Desktop Engineering's *resident blogger and senior editor. Email him at kennethwong@deskeng.com or share your thoughts on this article at deskeng.com/facebook.*

INFO
- **AMD:** AMD.com
- **Autodesk:** Autodesk.com
- **Citrix:** Citrix.com
- **Dell:** Dell.com
- **NVIDIA:** NVIDIA.com
- **VMware:** VMWare.com
- **Zuken:** Zuken.com

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

Still Sleek

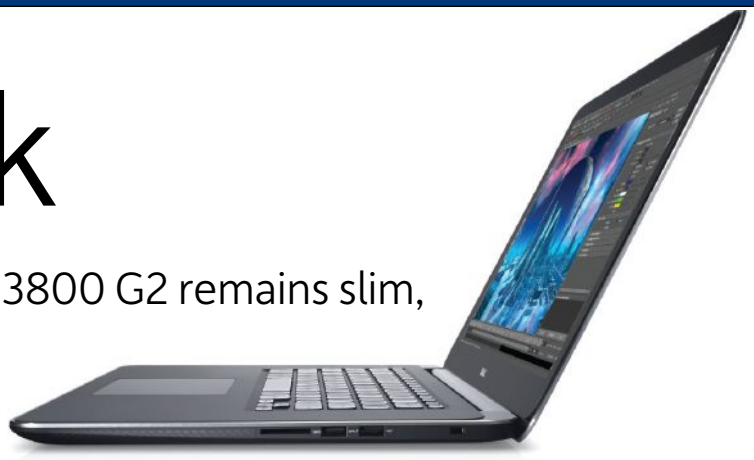

The updated Dell Precision M3800 G2 remains slim, fast and affordable.

BY DAVID COHN

Last year, we raved about the Dell Precision M3800, a remarkably thin, lightweight system that proved a slim laptop could also be very powerful (see deskeng.com/de/?p=15891). At the time, we stated that it ranked "as one of the most dramatic shifts we've ever seen." Needless to say, when Dell asked if we would like to review the updated version of the M3800, we were anxious to.

Outwardly, the new Dell Precision M3800 G2 (second generation) looks identical to its predecessor, which means that it bears a striking resemblance to Apple's 15-in. MacBook Pro. The Dell M3800 is sleek, weighing less than 6 lbs. including its external power supply and cables. Even the minimalist packaging resembles what we've come to expect from Apple.

Housed in a beautiful aluminum and carbon fiber chassis, the Dell Precision M3800 measures 14.56 x 10.0 x 0.71 in. and weighs 4.57 lbs. as tested. The 130 watt AC adapter adds 0.9 lbs., bringing the total weight to 5.65 lbs. That's light compared to mobile workstations that often tip the scales at more than 11 lbs. But, in the past year the rest of the world has caught up. Nearly every system vendor now offers a thin, lightweight mobile system — including the HP ZBook 14, the BOXX GoBOXX G1980 and the MSI WS60, all of which *DE* has reviewed.

INFO → Dell: Dell.com

Dell Precision M3800 G2

- **Price:** $2,109 as tested ($1,649 base price)
- **Size:** 14.56 x 10.0 x 0.71 in. (W x D x H) notebook
- **Weight:** 4.57 lbs. plus 0.9 lb. power supply
- **CPU:** 2.3GHz Intel Core i7-4712HQ with 6MB SmartCache
- **Memory:** 16GB (16GB max)
- **Graphics:** NVIDIA Quadro K1100M w/2 GB GDDR5 memory
- **LCD:** 15.6-in. UHD Ultrasharp 3840 x 2160 (4K x 2K) wide view backlit touchscreen LED
- **Hard Disk:** 256GB SSD
- **Optical:** None
- **Audio:** Built-in speakers, headphone/microphone jack, noise-canceling digital array microphones
- **Network:** USB 3.0 to Ethernet adapter, Intel Dual Band Wireless-AC 7260 Plus Bluetooth 4.0
- **Other:** Two USB 3.0, one USB 2.0, HDMI, mini DisplayPort/Thunderbolt 2.0 combo, 3-in-1 media card reader, integrated light-sensitive HD video webcam
- **Keyboard:** Integrated 80-key full-size backlit keyboard
- **Pointing device:** Gesture-enabled multi-touch touchpad with two buttons

Minor CPU Update

The original Dell Precision M3800 came with an Intel Core i7-4702HQ, a Haswell quad-core CPU with a base speed of 2.20GHz and a maximum turbo speed of 3.2GHz. For the second generation M3800, Dell replaced that processor with a slightly faster Intel Core i7-4712HQ, which has a base speed of 2.30GHz while accelerating to a maximum turbo speed of 3.3HGz. It too is a Haswell processor, and includes the same 6MB SmartCache and a 37-watt thermal design power (TDP) rating to help extend battery life.

Although the CPU includes integrated Intel HD Graphics 4600, Dell also equips the Precision M3800 with an NVIDIA Quadro K1100M discrete graphics card with 2GB of GDDR5 dedicated memory. This powerful GPU (graphics processing unit), with 384 CUDA (compute unified device architecture) parallel processing cores, has a 128-bit interface and a bandwidth of 44.8GB per second. With a maximum power consumption of 45 watts, the K1100M helps the M3800 deliver very good graphics performance while conserving battery power.

Opening the lid reveals a 15.6-in. display and the same full-size 80-key backlit keyboard as the original Dell Precision M3800, which means that it is still lacking a separate numeric keypad. A gesture-enabled multi-touch touchpad with two buttons is centered below the keyboard and a round power button is located to the upper-left. The caps lock button includes a small LED that indicates when caps lock is enabled.

The wide view LED backlit display extends practically to the edges of the lid. Above it is the same noise-canceling microphone array, ambient-light sensor, HD video webcam, and camera-status light as the original M3800.

So, other than the CPU, what makes the second generation M3800 different from its predecessor? As was true for the original

play. But instead of the QHD+ 3200x1800 touchscreen panel in the original M3800, Dell equipped our new evaluation unit with a 4K 3840x2160 IGZO UHD (ultra high-definition) touchscreen panel (a $99 option). Both panels come with 10-point multi-touch and are covered with Corning Gorilla Glass.

More Options

Dell doesn't skimp on memory, providing 8GB in the M3800 base configuration. Our unit came with 16GB of RAM, installed as two 8GB 1.6GHz SODIMMs (small outline dual in-line memory modules), an option that added $112 to the price.

Dell has responded to customer feedback regarding storage options. The base configuration still comes with a 500GB 7200 rpm hard drive, but now you can opt for solid-state drives (SSDs) ranging from 256GB ($59) up to a 1TB mSATA SSD ($999) as well as an Intel 360GB 2.5-in. Serial-ATA solid-state drive ($349). The second 2.5-in. hard drive bay can also be outfitted with a 1TB SSD, for a total of 2TB of solid-state storage (which would add nearly $2,000 to the $1,649 base price). While many customers may shy away from spending that much for internal storage, considering the performance advantage and additional power savings of an SSD drive, most will likely opt for some sort of solid-state drive.

Our system came with a 256GB SSD, adding $249 to the price. That's a rather small drive for a workstation. We think most customers will go for at least a 512GB SSD – a $499 add-on.

Like the original, however, the M3800 G2 does not include an optical drive. While that is becoming less and less of an issue for many users, we had to first copy programs onto an external USB hard drive before we could install them on the M3800 G2. Dell sells an external USB DVD+/-RW drive for $60 and similar drives are available online for as little as $25.

In the new M3800, the mini DisplayPort doubles as a Thunderbolt 2.0 port. The new system also comes with a USB 3.0 Ethernet dongle, whereas the original came with a USB 2.0 version.

The Dell Precision M3800 system we received came with the six-cell 61 watt-hour 12.2 volt base system battery, which kept us running for 5 hours and 34 minutes before saving all data and shutting down. A 6-cell 81WHr battery is a $42 option.

Performance Largely Unchanged

What a difference a year makes: When we reviewed the original Dell Precision M3800, it out-performed many of the larger, more expensive mobile workstations we had reviewed less than a year before. While the new M3800 matched or exceeded its predecessor on the SPEC Viewperf and SOLIDWORKS benchmarks, it actually took a few seconds longer than the older system to complete our AutoCAD rendering test. (See page 30.)

We also ran the SPECwpc workstation performance benchmark. While the slightly faster CPU enabled the M3800 G2 to beat the original M3800 on most aspects of this test, it lagged behind most of the newer systems, performing just slightly better than the HP ZBook 14 we recently reviewed.

Mobile Value

Throughout our tests, the M3800 G2 ran cool, but not quite as silent as the original. Last year, we wondered whether the M3800 even had a fan. This year, under the extreme compute loads imposed by several of our benchmarks, the M3800 G2 fan was audible, although still quiet compared to many other systems.

Our system came with Windows 8.1 Professional 64-bit. Windows 7 and Ubuntu Linux 14 are also available. Dell backs the system with a one-year ProSupport plan plus one year of next business day limited on-site service after remote diagnosis.

At its base price of $1,649, the Dell Precision M3800 G2 is an excellent system. As configured, our unit priced out at $2,109. Even when we configured a new system to match the one we reviewed last year, the new Dell Precision M3800 G2 would cost $2,401, still $486 less than last year. The Dell Precision M3800 G2 remains an excellent choice for any engineer on the go, offering light weight, style and performance at an affordable price. 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* DE *and the author of more than a dozen books. You can contact him via e-mail at david@dscohn.com or visit his Website at www.dscohn.com.*

Mobile Workstations Compared

		Dell Precision M3800 G2 2.3GHz Intel Core i7-4712HQ quad-core CPU, NVIDIA Quadro K1100M, 16GB RAM	HP ZBook 14 G2 2.6GHz Intel Core i7-5600U dual-core CPU, AMD FirePro M4150 and Intel HD Graphics 5500, 16GB RAM	BOXX GoBOXX G1980 2.8GHz Intel Core i7-4980HQ quad-core CPU, NVIDIA Quadro K1100M, 16GB RAM	Eurocom P5 Pro 4GHz Intel Core i7-4790K quad-core CPU, NVIDIA Quadro K510, 32GB RAM	MSI WS60 2.5GHz Intel Core i7-4710HQ quad-core CPU, NVIDIA Quadro K2100M, 16GB RAM	Dell Precision M3800 2.2GHz Intel Core i7-4702HQ quad-core CPU, NVIDIA Quadro K1100M, 16GB RAM
Price as tested		$2,109	$2,115	$3,365	$4,489	$2,600	$2,887
Date tested		5/25/15	2/20/15	2/4/15	2/4/15	1/17/15	3/13/14
Operating System		Windows 8.1	Windows 8.1	Windows 8.1	Windows 7	Windows 7	Windows 7
SPECviewperf 12	higher						
catia-04		15.16	15.09	21.26	64.64	21.26	14.74
creo-01		15.36	16.57	20.28	48.70	19.98	13.37
energy-01		0.34	0.06	0.32	2.61	0.32	0.28
maya-04		13.85	9.09	18.20	48.84	17.90	12.79
medical-01		4.30	2.70	5.71	23.93	5.71	3.72
showcase-01		8.55	7.58	10.35	27.86	10.63	8.50
snx-02		15.30	20.06	22.10	58.41	22.05	14.74
sw-03		25.41	29.21	34.53	97.38	32.32	19.43
SPECviewperf 11	higher						
catia-03		32.18	32.53	47.85	80.24	45.66	33.56
ensight-04		17.38	17.51	24.55	86.39	24.09	17.50
lightwave-01		60.00	65.87	77.62	94.51	64.37	58.84
maya-03		62.83	61.30	81.14	178.55	77.78	61.83
proe-05		13.92	9.70	21.57	22.67	18.26	15.37
sw-02		39.91	43.45	52.31	81.17	47.80	39.48
tcvis-02		28.59	13.38	37.24	70.60	36.95	28.69
snx-01		23.77	23.55	31.76	89.35	31.85	23.76
SPECapc Solid-Works 2013	higher						
Graphics Composite		2.51	2.98	5.59	9.00	3.08	2.41
RealView Graphics Composite		2.54	3.26	5.86	10.61	3.23	2.71
Shadows Composite		2.48	3.20	5.92	10.65	3.23	2.34
Ambient Occlusion Composite		2.36	4.63	7.21	21.36	3.51	2.20
Shaded Mode Composite		2.42	2.90	5.36	8.88	2.96	2.31
Shaded with Edges Mode Composite		2.61	3.07	5.82	9.12	3.21	2.51
RealView Disabled Composite		2.41	2.08	4.61	4.66	2.55	2.40
CPU Composite		2.20	2.65	3.88	4.25	3.06	2.41
Autodesk Render	lower						
Time	seconds	79.38	124.28	55.39	56.88	63.60	71.42
Battery Test (H:MM)	higher	5:34	7:28	2:15	2:10	3:13	6:12

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

Very Fast Indeed

Computer Direct Business delivers the speedy Volta Pro VP1 engineering workstation.

BY DAVID COHN

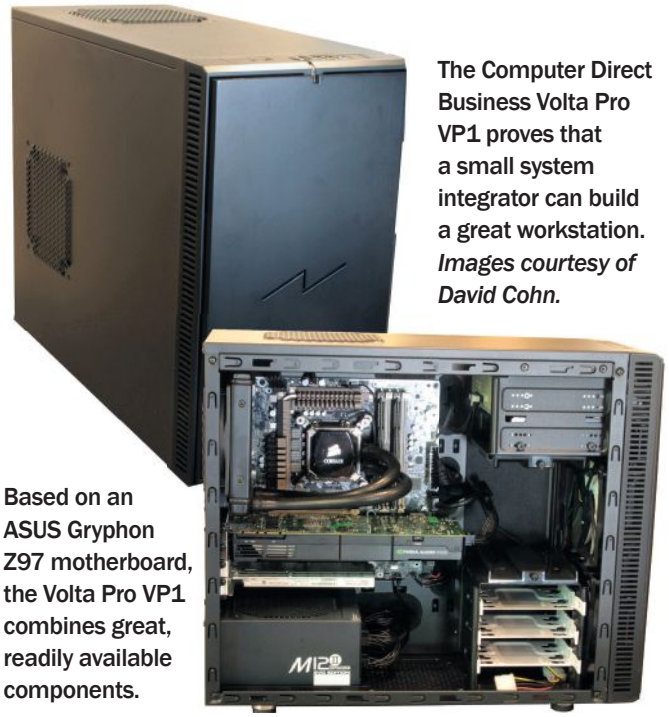

The Computer Direct Business Volta Pro VP1 proves that a small system integrator can build a great workstation. *Images courtesy of David Cohn.*

Based on an ASUS Gryphon Z97 motherboard, the Volta Pro VP1 combines great, readily available components.

Computer Direct Business recently asked if we would like to review their newest engineering workstation. If the name doesn't sound familiar, that's probably because the company is a system integrator, not an OEM (original equipment manufacturer). Computer Direct builds its systems by combining readily available components rather than manufacturing at least some of the parts or subsystems (like Dell, HP and Lenovo). The lack of name recognition may also result from the fact that Computer Direct is based in Greenville, SC — not exactly the heart of Silicon Valley. But we were intrigued, particularly because the company is a dedicated sales and service division of Computer Direct Outlet, a discount retailer and service center operating since 1998. *DE* has reviewed other custom systems in the past and has often been pleasantly surprised by what we receive.

Computer Direct Business sent us its Volta Pro VP1, a well-designed workstation based on a fourth-generation Intel Haswell CPU. The system is housed in a solid black ATX tower case measuring 8.27x19.29x15.55 in. (WxDxH) made in Sweden by Fractal Design. The complete system weighed 31.5 lbs. The case features side and top vents with ample sound dampening and noise absorbing material on the side panels. A hinged door conceals two front panel 5.25-in. drive bays and a pair of fans. The top bay in our evaluation unit housed a Liteon 24X DVD+/-R/RW burner.

A round power button, a pair of USB 2.0 ports, a USB 3.0 port, microphone and headphone jacks are on the front top edge of the case. The rear panel provides four more USB 3.0 ports, an RJ-45 network jack, an optical S/PDIF output port and five audio jacks. The side panels are held in place with non-captive thumbscrews. In addition to the two external bays, the Fractal case provides room for up to six internal bays. Our evaluation unit included three internal drive bays, however.

Classy Components

The Volta Pro VP1 is built around the ASUS Gryphon Z97 motherboard, an excellent board based on the Intel Z97 Express chipset. The motherboard provides four DIMM (dual in-line memory module) sockets supporting a maximum of 32GB of non-ECC un-buffered memory modules of up to 1866MHz. The base Volta Pro VP1 configuration comes with 8GB of RAM using a pair of 4GB DIMMs. Our evaluation unit came with 16GB of memory, installed as a pair of 8GB Crucial Ballistix 1600MHz modules, adding $56 to the system price. The LGA1150 socket on the ASUS motherboard supports a wide range of Intel processors, and Computer Direct Business offers a choice of five different CPUs ranging from the 3GHz Core i5-4460 quad-core processor in the base Volta Pro VP1 configuration to the Core i7-4790K installed in our evaluation unit. While that Intel CPU has a standard speed of 4GHz and a maximum turbo speed of 4.4GHz, Intel designed this processor to be overclocked, and Computer Direct Business did indeed boost the CPU in our system to 4.5GHz. That faster CPU plus overclocking added another $322 to the cost of the base system.

Expansion options on the ASUS motherboard are limited to a pair of PCIe 3.0 x16 slots that support a single x16 board or two x8 boards, plus a PCIe 2.0 x16 (x4 electrically) and one PCIe x1 slot. The base Volta Pro VP1 configuration comes with an entry-level NVIDIA Quadro K620 GPU, but one x16 slot in our evaluation unit housed a high-end NVIDIA Quadro K5200 graphics card with 8GB of GDDR5 memory and 2304 CUDA cores. The high-end NVIDIA board added $2,117 to the system price. The Quadro K5200 provides a pair of DVI ports as well as two DisplayPorts. The rear panel also has DVI, HDMI, and DisplayPorts connected to the Intel HD Graphics 4600 on the CPU.

While the base Volta Pro VP1 configuration comes with a 250GB Samsung 850 Pro SSD (solid-state drive), our unit came with a 400GB Intel 750 PCIe SSD, installed in one of the re-

Single-Socket Workstations Compared

		Computer Direct Volta Pro One 4.0GHz Intel Core i7-4790K 4-core CPU, over-clocked to 4.5GHz NVIDIA Quadro K5200, 16GB RAM	BOXX APEXX 2 2401 One 4.0GHz Intel Core i7-4790K 4-core CPU over-clocked to 4.5GHz, NVIDIA Quadro K5200, 16GB RAM	Xi Mtower PCIe One 3.7GHz Intel Core i7-5930K 6-core CPU over-clocked to 4.32GHz, NVIDIA Quadro K5200, 16GB RAM	Lenovo P300 One 3.6GHz Intel Xeon E3-1276 v3 4-core CPU, NVIDIA Quadro K4000, 8GB RAM	Digital Storm Slade PRO One 3.4GHz Intel Xeon E3-2687W v2 8-core CPU, NVIDIA Quadro K4000, 32GB RAM	HP Z1 G2 One 3.6GHz Intel Xeon E3-1280 v3 4-core CPU, NVIDIA Quadro K4100M, 16GB RAM
Price as tested		$4,441	$5,111	$4,985	$2,072	$5,804	$5,918
Date tested		7/12/15	2/4/15	12/13/14	11/9/14	5/10/14	5/3/14
Operating System		Windows 7	Windows 7	Windows 7	Windows 7	Windows 7	Windows 8.1
SPECviewperf 12	higher						
catia-04		103.66	100.40	98.53	38.19	34.81	42.23
creo-01		91.62	77.69	86.66	34.31	33.15	30.82
energy-01		3.73	3.61	3.49	0.65	0.60	1.74
maya-04		75.92	74.68	72.18	32.31	31.28	33.79
medical-01		31.33	30.01	28.84	12.38	10.75	10.34
showcase-01		49.76	49.76	48.98	22.64	20.65	21.12
snx-02		152.32	83.03	150.42	36.79	34.12	40.37
sw-03		134.67	130.28	126.08	69.37	50.78	38.66
SPECviewperf 11	higher						
catia-03		134.82	131.40	99.71	67.84	69.41	63.80
ensight-04		145.75	152.22	148.83	48.80	47.76	61.56
lightwave-01		109.59	107.01	100.99	88.54	76.90	82.76
maya-03		131.43	245.35	99.44	132.59	101.12	128.09
proe-05		28.51	27.19	18.19	21.34	16.29	17.18
sw-02		97.48	96.35	88.89	72.05	63.66	67.75
tcvis-02		108.24	106.96	78.64	55.66	54.26	58.99
snx-01		135.41	137.53	134.51	53.24	52.98	65.58
SPECapc SolidWorks 2013	higher						
Graphics Composite		11.24	10.27	8.82	6.29	5.37	5.67
RealView Graphics Composite		13.32	12.08	10.03	6.88	5.90	6.16
Shadows Composite		13.37	12.12	10.05	6.89	5.85	6.13
Ambient Occlusion Composite		28.08	24.55	17.58	9.65	9.46	8.48
Shaded Mode Composite		11.25	10.25	8.95	6.17	5.30	5.55
Shaded with Edges Mode Composite		11.22	10.30	8.69	6.41	5.45	5.79
RealView Disabled Composite		5.69	5.37	5.28	4.39	3.70	4.08
CPU Composite		4.87	4.87	4.50	4.18	3.70	3.12
Autodesk Render Test	lower						
Time	seconds	50.83	41.88	42.33	64.08	38.25	45.00

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

maining PCIe x16 slots, adding $344 to the base price. It meant there was only one expansion slot remaining while the drive bays were empty. The company offers drives ranging from a 500GB 7200 rpm drive up to a 1.2TB SSD. Computer Direct Business included a Corsair Hydro Series H60 liquid cooling system. Power was provided by a Seasonic 850watt 80 Plus Bronze power supply. The system was quiet even under heavy compute loads.

Great Performance

By the time we added up the options, our Volta Pro VP1 priced out at $4,441. But how would it perform?

On the SPEC Viewperf tests, the Computer Direct Business Volta Pro VP1 beat every other single-socket workstation we have tested to date, owing largely to its overclocked CPU and high-end graphics card. On the SPECapc SolidWorks benchmark, it again beat other systems equipped with a single CPU. But on the AutoCAD rendering test, the Volta Pro VP1's quad-core CPU could not match the rendering performance of the six- and eight-core systems we've recently reviewed.

We also ran the SPECwpc workstation benchmark. The Volta Pro VP1 turned in great numbers on tests that relied primarily on graphics performance but dropped to the middle or bottom of the pack on elements of the test that were multi-threaded.

Computer Direct Business pre-loaded Windows 7 Professional 64-bit. Windows 8.1 is also available. The company does not offer them as an option when configuring a system online. Strangely, the system comes with only a 1-year warranty on parts, yet most of the components used to build the system carry their own multi-year warranties. Computer Direct Business does offer lifetime labor and tech support, however, and will cover parts for a second or third year for an additional charge.

While it lacks ISV certification and you could essentially build the same system yourself for less, if the Volta Pro VP1 is any indication, the folks at Computer Direct Business know how to build great systems. The Volta Pro VP1 combines excellent components in a well-apportioned package offering great performance at a very attractive price. 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* DE *and the author of more than a dozen books. You can contact him via e-mail at david@dscohn.com or visit his Website at dscohn.com.*

INFO ➔ Computer Direct Business: ComputerDirectBiz.com

Computer Direct Business Volta Pro VP1

- **Price:** $4,441 as tested ($1,602 base price)
- **Size:** 8.27 x 19.29 x 15.55 in. (W x H x D) tower
- **Weight:** 31.5 lbs.
- **CPU:** One 4.GHz Intel Core i7-4790K quad-core (over-clocked to 4.5GHz)
- **Memory:** 16GB DDR3 non-ECC at 1600MHz (2X8GB), 32GB max
- **Graphics:** NVIDIA Quadro K5200
- **Hard Disk:** 400GB Intel SSD
- **Optical:** 24X DVD+/-R/RW
- **Audio:** Integrated Realtek ALC892 8-channel HD audio (top panel: headphone and microphone; rear-panel: microphone, line-in, line-out/front, rear, bass and SPDIF)
- **Network:** Integrated Intel I218V gigabit LAN controller with RJ45 port
- **Other:** Five USB 3.0, two USB 2.0; DVI, HDMI and DisplayPort; two DVI ports and two DisplayPorts on NVIDIA board
- **Power supply:** 850 watts 80Plus Bronze
- **Warranty:** One-year for parts, lifetime labor

CHAPTER 2 EXCERPT

Building a Balanced Workstation

Choose the best hardware components for the design engineering work you do.

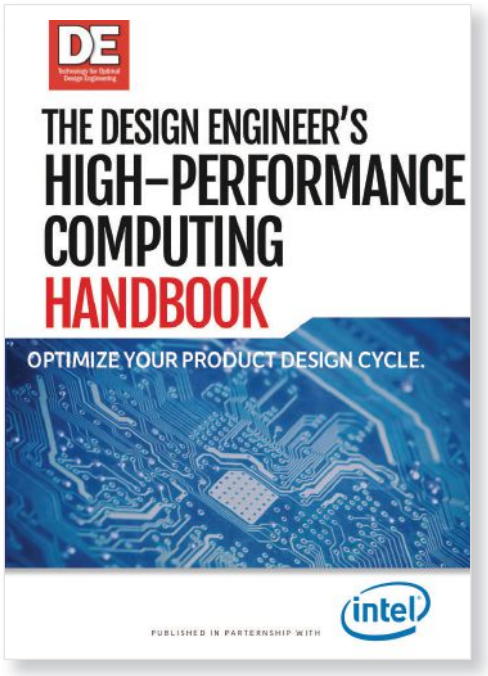

Nobody likes a sluggish computer, but for designers and engineers, a workstation that gets bogged down means valuable design optimization time is ticking away. To keep engineers as productive as possible, companies need to invest in professional-grade workstations, but not every unit is created equal.

"The annualized cost of a high-end workstation on a two- to three-year refresh cycle is normally much lower than the annualized cost of a user's CAD software, yet the correct hardware can have a profound impact on CAD user productivity, and the wrong hardware can bring productivity to a halt," says Sean Young, worldwide segment manager, product development and AEC at HP.

Intel's Andy Parma, Workstation Segment Manager, says engineers should consider balance when considering a new workstation. "When a workstation is balanced, it provides excellent performance across four key elements: processor, memory, storage and graphics," he says. "When these elements are in balance, workstation users can explore more ideas in less time."

A modern high-end CAD workstation will have a multi core CPU, Hyper-Threading and a powerful GPU, along with a solid-state drive (SSD) instead of a traditional hard drive to speed data access. But the minimum hardware requirements recommended by software vendors don't always apply to every machine. Not every workstation necessarily has to have the best graphics. And just throwing a lot of money at the problem in the form of multiple processors with lots of cores doesn't guarantee performance improvements.

"Don't spend a lot on one component and sacrifice in other areas," says Scott Hamilton, Precision specialist at Dell. "In most cases, you get the most benefit from a system where all the components work well together. As long as you look at the workflow of your users and the applications they are using, you'll wind up with something good."

The types of applications you need to run will guide selection, as will determining what apps will be running simultaneously. Photo rendering and simulation have a big

How the HPC Handbook Works

The Design Engineer's High-Performance Computing Handbook is a multimedia resource consisting of magazine article excerpts; the hpc.deskeng.com website that is full of videos, case studies and research; e-newsletters; and regularly released chapters that are free to download. Once you download a chapter, you'll be alerted when new chapters are ready. Each chapter takes a detailed look at a computing topic important to design enigneers. The chapters include pictures, charts, definitions and links to additional information.

The HPC Handbook site is the hub of information for high-performance computing in design engineering. The HPC Handbook is available for download on the site, and the site is constantly updated to include the latest HPC educational resources that are pertinent to design engineering teams. It is divided into sections on Workstations, Clusters/Servers, Cloud Computing and Software, each of which contain the best information available on the Web from *DE* and beyond.

Check it out at hpc.deskeng.com.

impact on hardware configurations, for example, along with the size and complexity of models and assemblies. How you plan to use the workstation will determine whether you'll put your money into more cores, faster processors or more expensive GPUs.

"Put your dollars toward things that will directly impact performance," says David Waters, head of industry alliances, professional graphics, at AMD. "That's the biggest thing right from the beginning."

Solid-State Drives

Every expert *DE* spoke to noted that if you are putting together a professional workstation, then you'll have an SSD instead of a spinning hard disk — preferably for both the boot drive and storage. For the price (only a few hundred dollars more than a spinning hard drive) you get a tremendous boost in performance.

CPU Cycles and Cores

For CPUs, speed is more important than core count if you're running a CAD program. Look at single-thread performance instead of getting caught up in maximizing the number of cores. The number of cores is important as you add other applications, and particularly if you plan to do any simulation, analysis or photo realistic rendering work. Those are the applications that will take advantage of more cores because they are multi-threaded.

Buy a Lot of RAM

The good news when it comes to workstation configuration is that one of the most powerful performance boosters is also one of the cheapest investments: RAM. You need enough memory to avoid swapping hard drive space to store and retrieve active data. Not having enough RAM will slow you down, even with a fast processor.

The GPU

The size and complexity of the models are the biggest considerations when it comes to GPU selection. The system should provide a 15 frames per second (fps) to 20 fps refresh rate; anything below that would look slow or clunky, although what level of performance is "good enough" is a subjective measure. If you can't spend a lot of money on the GPU, you can also consider making adjustments to the design software. You can lower the level of detail in a model and get smooth rotation, it just depends on the comfort level the engineer has with that trade-off and how much detail is required. There are also mid-range cards available that can provide a lot of horsepower, even for larger models.

Bigger, Better Displays

Another element that can drive GPU selection is the display. Most CAD and simulation software users have more than one monitor, a practice that improves productivity. If you have multiple screens, then you'll need professional graphics card like the AMD FirePro or NVIDIA Quadro line that can easily drive multiple monitors. The Intel Xeon processor E3-1200 v3 Product Family supports up to three displays. Each display port is capable of driving resolutions up to 3840x2160 at 60 Hz through DisplayPort and 4096x2304 at 24 Hz/2560x1600 at 60 Hz using HDMI.

Workstation configuration is always limited by the budget, and that means companies will need to make trade offs. Those trade offs should be based on application usage to get the most out of your hardwre, software and manpower investements.

If Not Now, When

With all of the evidence in favor of workstation upgrades, why are many design engineers still limping along on three-year-old (or older) computers?

Some of it is budgetary. Those who hold the purse strings may not have bothered to calculate the quick ROI of an upgrade, or may not fully understand that engineers have different computing needs than their co-workers in accounting or sales departments.

Other reasons for keeping engineers from the computing power they need are cultural. People are comfortable with the status quo because they don't know what they are missing. Some are reluctant to make a change during a large project — and there's always a large project.

To get beyond the budgetary and cultural shortsightedness, look to the future. Models are going to continue to grow in size, and modern design software capabilities will continue to take advantage of the latest advances in computing hardware.

Over the last three years, those hardware advances include more productive computational and graphics processing via new architectures and more cores; faster memory speeds with error-correcting code memory that helps prevent data corruption; and the doubling of disk drive speeds thanks to SSDs. Virtualization — where one workstation can run multiple operating systems or be divided to handle different tasks — is becoming increasingly common, as is the use of multiple monitors.

"Processors, bus speeds, graphics cards and hard drives are all faster," says Lenovo's Tom Salomone. "Engineers may think that hardware is just a little faster or the same, but the truth is it all adds up to let them work much, much faster."

CHAPTER 3 EXCERPT

Parallelization Primer

Parallel processing has revolutionized the way design software operates.

Few technology advances have had quite the same impact on design software as the advent of parallelization or parallel processing. By taking advantage of the power of multiple cores and multiple CPUs and GPUs (locally or on multiple servers or cloud resources), engineers and designers have been able to quickly and cost-effectively wrangle very large simulation and rendering tasks that previously would have required outsourcing and many hours (or days) of delays.

Parallelism leverages concurrency to gain better performance. Parallel computing defines these very large simulation problems in smaller pieces and simultaneously solves each piece using multiple cores, processors or computers. As a result, you can get your answer faster (lower latency) or you can find more answers in the same amount of time by increasing throughput.

"If you are handling and processing a lot of data, you are ripe for using parallelism," says James Reinders, parallel programming evangelist at Intel."

Parallelism provides an opportunity to direct the multiple functions of a computer system in a fashion similar to conducting an orchestra. "How can the system use multiple resources in concert, rather than one a time? 'You have to think about how you are using the total computer system," Reinders says.

This is different than multi-threading, in which a single CPU or core executes multiple processes or threads concurrently by splitting up the data and tasks into sub-tasks on shared memory. On a CPU with Hyper-Threading, for example, you can perform multiple sequential tasks by running several dozen threads in parallel with multiple cores. "Multi-threading is for in-process computation, while parallel computing is for out-of-process computations," says Silvina Grad-Freilich, MathWorks' senior manager for parallel computing and deployment products.

Parallel processing can be enabled via modern CPUs or GPUs. The CPU is good at loading and holding data so that the same data can be used over and over again. GPUs can perform parallel computations that are the same or similar on different data.

"GPU accelerators complement CPUs to provide the best app performance for end users," says Will Ramey, senior product manager of accelerated computing at NVIDIA.

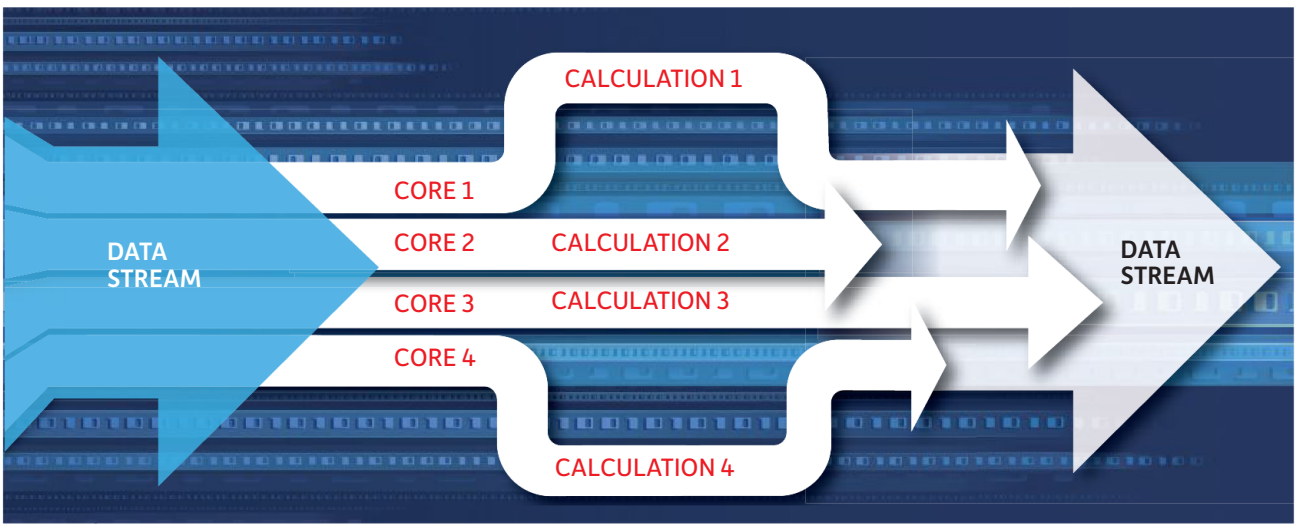

In parallel computing, multiple cores can be used to cooperatively solve large computations by dividing them into smaller computations that can be solved simulataneously, thus speeding the overall solution time.

Faster, Better Results

With software applications that have been designed to take advantage of parallelism, engineers can complete complex tasks much faster.

"With parallelized software, our engineering customers can really achieve enhanced engineering productivity by accelerating simulation throughput," says Wim Slagter, lead product manager for high performance computing at ANSYS. "It also helps them to make more efficient product development decisions."

Engineers can also achieve higher fidelity insight into product performance that couldn't be gained any other way. "Parallelized software capabilities allow engineers to simulate larger models and more complex models, so that more accurate design decisions can be made throughout the design or development cycle," Slagter adds.

"If the simulation runs faster, then you can do more iterations of exploring parameters of the design space much faster," Ramey says. "You can refine those designs faster, which results in better products."

Where Core Count Matters

Where this approach to computation matters most in design is typically in simulation and rendering involving large amounts of data. Computational fluid dynamics (CFD) simulations, for example, benefit from higher core counts. For a large aerospace company, huge simulations rely on throwing large numbers of cores at the computations.

"The applications that benefit the most are those that have computations heavy enough that they can benefit from breaking the data into chunks," says Grad-Freilich. "You need to make the right decision for the technology that you use based on your actual application needs."

However, not every task within an application can be parallelized. Companies may not see the speed-up they expect if there are lengthy sequential processes involved in the solution. "For example, an application may have a sequential part that runs before and after the other computational work, so the overall speedup of the application is less than linear because of the serial nature of the beginning and end processes," Grad-Freilich says.

That limit on acceleration is explained by Amdahl's Law, which limits the theoretical speed-up of applications using parallelism. Because some processes are inherently serial, there is a point at which the application can't be made any more parallel. That point varies depending on the type of work involved.

There are software applications that don't benefit from parallelism, or that don't benefit from an exponential increase in core count. "Mechanical simulations, for example, do not scale up to tens of thousands of cores," Slagter says. "We've shown in our latest release that it scales up to 128 cores for a whole suite of different benchmarks."

"If you have data that must be operated on sequentially, and there's only a small amount of computation for each piece of data, that application would not benefit from parallelization," Ramey adds. "The good news is that when smart programmers get engaged, the algorithms in applications that are fundamentally serial can often be redesigned so they can run in parallel."

Modern design workstations typically have multiple processors, which provides HPC-like capabilities on the desktop. Software of all types has been designed to take better advantage of those capabilities.

"We keep making our data problems bigger and bigger," Reinders says. "Because of that, computers can get more use from parallelism as time goes on. As long as we use more and more data, then parallelism will be required."

The majority of engineering and design software has been architected to take advantage of parallelism. Some software providers have struggled with building tools to take advantage of this approach because of inexperience or because the software was originally designed when single-core was the standard.

"The question becomes how much does this perturb the way the program was written originally," Reinders says. "There are a lot of applications that were written when there was just one core in a machine. This can affect the architecture of the application in a profound way."

Developing software that takes advantage of these capabilities also requires sustained and continued HPC software development to effectively leverage the hardware. "The software should be able to support simulation resources where they are located," Slagter says. "Certifying remote software solutions is important, along with expertise and support from HPC partners like Intel, NVIDIA, Hewlett Packard, etc., so you can make sure that you have an optimized reference architecture for the software and good support."

Learn More

To learn more about building a balanced workstation and parallelization, download *The Design Engineer's High-Performance Computing Handbook,* which has just been updated with two new chapters covering those topics. The full version of chapters 2 and 3 include a discussion of the benefits of error-correcting code RAM, software optimization tips, ROI calculations, a discussion on the blurred boundaries between CPUs and GPUs, definitions of terms and more.

This free resource, brought to you be *DE* in partnership with Intel, is available at hpc.deskeng.com/download.

PLE FYI

An example of what product line engineering means to the design engineer.

BY FRANK J. OHLHORST

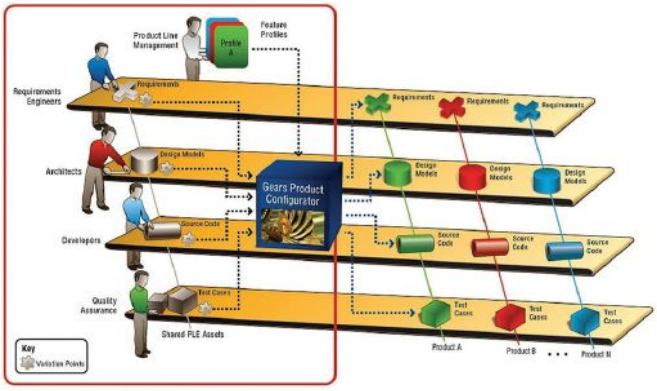

The Gears solution is designed to enable you to create an automated production line comprised of configurable assets, feature profiles and products. *Image courtesy of BigLever Software.*

Product line engineering (PLE) could be the difference between success and failure when building most anything. The basic concept behind PLE is creating, automating and managing the processes that allow a related set of products to share engineering assets and effort in order to achieve an efficient means of production. In other words, PLE makes creating things (physical goods, services, software, etc.) easier and more cost effective.

The magic behind PLE comes from identifying and assigning assets to the production process and then managing those assets efficiently, while also identifying where improvements can be found and redundancies eliminated. It's a concept that sounds like common sense, yet has been difficult to implement beyond the traditional assembly line process of building hard goods.

However, new applications and platforms are changing that and are enabling firms to leverage the concepts of PLE to create almost anything, ranging from products to services to software and infrastructure. "The sophistication of products and systems being deployed throughout the industry requires a commensurate sophistication in the alignment of tools and methods for the traditionally siloed hardware and software engineering," says Charles Krueger, CEO and founder of BigLever Software. It is that inherent sophistication that drives the need for PLE, and in turn PLE reduces the complexity of building, delivering, maintaining and evolving an entire product line portfolio.

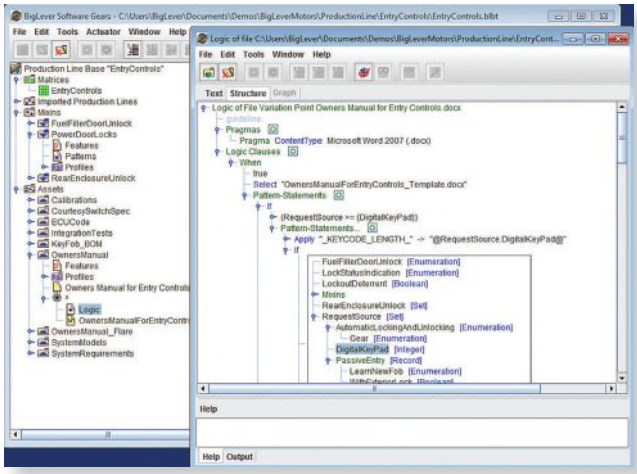

The Gears PLE Lifecycle Framework provides a common set of PLE concepts and constructs that are intended to enable the integration of tools, assets and processes across each stage of the engineering lifecycle – from requirements and design to implementation, testing, maintenance and evolution.

Army Simplifies the Complex

To help implement PLE, BigLever Software is currently partnered with General Dynamics to provide Advanced Product Line Engineering Capabilities for the U.S. Army's Consolidated Product-line Management (CPM) Next Program. The CPM Next contract, which was awarded earlier this year, allows the General Dynamics-led team to lead the development and management of the Army's Live Training Transformation (LT2) product line for combat training.

The Army's adoption of PLE demonstrates how the creation of a complex product line can be transformed to achieve significant improvements in engineering productivity, time to deployment, quality and cost. For the Army, the LT2 product line is one of the most important implementations of PLE, since the end result is the effective combat training of soldiers, while also supporting the rapid deployment of new tactic variations and enhancements to addresses changing missions and threats.

Simply put, by adopting PLE, the Army can better train combat soldiers, regardless of mission and situational awareness changes and achieve goals more efficiently and more reliably. While the Army's adoption of PLE may be considered an extreme case, it validates that PLE can be applied to critical situations and that production line methodologies are applicable to a wide variety of use cases.

The benefits for the Army were significant. "The Army had a product line concept, but different contractors were implementing their products, creating a divergence from the baseline," says Roger McNicholas, director of Training Systems and Technology for General Dynamics Mission Systems. "By implementing PLE, we were able converge 22 different products into a single manageable baseline, which brought order to the chaos."

The goal of the CPM contract was to streamline the process for rapidly adapting training systems and equipment to changing missions and threats. LT2's PLE approach has delivered significant cost savings to the U.S. Army to date, with more than 200 training ranges worldwide that can scale from individual soldiers to full brigades participating in live force-on-force and force-on-target training.

Naturally, combat training proves to be a critical product for the Army to deliver to its soldiers and using PLE to improve the product comes with other benefits. "Some of the immediate benefits included reducing some 26 million lines of software code to just 5 million lines of code. Also the implementation of PLE expedited product maturity and offered a 300% increase in PDSS (post-deployment software support) efficiency," says McNicholas.

That said, there was an even more impressive financial benefit the Army realized from deploying PLE. The reuse cost avoidance is $660 million and sustainment cost avoidance is $86 million, so the LT2 total lifecycle cost avoidance is $746 million over 10 years, according to BigLever. That's a fancy way of saying that the Army potentially saved almost three-quarters of a billion dollars by adopting a production line ideology, backed by the software tools to manage and control the process.

Implementation Strategies

Although PLE can offer significant benefits and savings, it is not a plug-and-play technology. Those looking to adopt PLE will face some hurdles on the path to success. "Implementing PLE is going to take the appropriate software tools, governance and most importantly, a team willing to change," says McNicholas.

Team culture is a big part of the equation; McNicholas noted that "It is surprising how much time it takes to change the culture." Some of McNicholas' other observations include "discovering how much redundancy there is when you peer under the hood of an existing process" and "how collapsing product teams into a factory concept delivers on the promise of reusability of code and processes."

PLE and Design Engineering

While those realizations further demonstrate the benefits of PLE, the question remains, "how can engineering firms benefit from PLE?" From Krueger's point of view, the benefits are obvious. "PLE brings with it a single source of feature truth, which provides a holistic view into the feature variations for a product family. That in turn eliminates the need for multiple feature management mechanisms for multiple tools, which facilitates better interoperability across software and mechanical teams, and ultimately improves cross-functional alignment and communication between engineering and business operations," Krueger says.

Business speak aside, Krueger is espousing the truth about how PLE transforms production of systems and software (or most anything) from a bill of materials (BOM) based process to a bill-of-features process, which is translated into materials. Those materials can be thought of as anything: mechanical, electrical, wiring, lines of code, calibrations, requirements, designs, test cases, documentation and much more. This can make the bill-of-features ideology appropriate for any production line process.

Ultimately, adopting PLE can help engineering firms to:

- Increase a product line's scale and scope
- Make more efficient use of highly paid specialists
- Get new products to market faster
- Increase product quality by reducing mistakes

While PLE software is a tool to make the above happen, it is the ideology of the production line-centric approach that makes the process possible. Engineering firms developing large-scale, complex products or delivering engineering services have much to gain from PLE. **DE**

Frank Ohlhorst *is chief analyst and freelance writer at Ohlhorst.net. Send e-mail about this article to DE-Editors@deskeng.com.*

INFO ➜ **BigLever Software:** BigLever.com

➜ **General Dynamics:** GeneralDynamics.com

➜ **LT2 Product Line Portal:** lt2portal.mil

➜ **U.S. Army:** Army.mil

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

Optimizing and Compromising Lead to Winning Design

The DECKED truck bed storage system is designed to maximize the use of space with divided drawers and a deck that can support a 2000-lb. payload. *Image courtesy of DECKED.*

Altair ProductDesign helps DECKED design, simulate and test an innovative truck bed storage system.

BY JAMIE J. GOOCH

Unless your last name is da Vinci, it's unusual for an engineer to sign his work. That's especially true when the engineers are from a service provider hired to help develop a product. "Customers typically like to keep us a secret," says Tim Smith, vice president of Technical Operations at Altair ProductDesign. That wasn't the case when the company was hired by Jake Peters, the general manager of DECKED, to help develop a truck bed storage system for the startup company.

"If you look at the bottom of the drawer in the storage system, they actually had the signatures of the people from Altair who worked on the product and the Altair logo," says Cliff Boggs, vice president of Business Development for Altair ProductDesign.

"It was very refreshing to have a customer acknowledge us in such a permanent way by having that molded into the bottom of the drawer," Smith says.

The story of how the DECKED truck bed storage system went to market began when Peters approached Altair ProductDesign, a division of Altair that provides product development, industrial design, workflow automation and staffing services.

The storage system was intended to "productize plywood," Peters says, referring to the plywood storage boxes he and his partners saw in pickup truck beds around the ski resorts near Ketchum, ID, where the company is based.

Meeting Many Requirements

By the time they connected with Altair ProductDesign, the DECKED team had done its homework. They knew what the market would demand in terms of size, functionality and price. They also wanted their product to have a high-quality, custom feel that would appeal to customers who loved the outdoors and be rugged enough for customers in the construction trades.

The initial CAD design DECKED brought in failed validation. It was too heavy and would cost too much. Plus, DECKED didn't know whether it would meet their payload requirements.

Altair ProductDesign engineers started anew on the design so they could begin the process using the HyperWorks suite of optimization and simulation software from its parent company. The price point and weight requirements dictated a molded plastic product. Peters also wanted to be able to sell the product online and ship it according to certain dimensions and weight limitations. The DECKED truck bed system also had to fit the beds of many sizes of pickup trucks from different manufacturers.

Peters initially wanted to be able to accommodate both full- and mid-sized trucks with 5.5-, 6- and 8-ft. bed lengths. To accommodate beds of different widths, the initial plan was to divide the DECKED storage system not down the middle, but in a 60/40 split for full-sized beds, which would allow two of the 40% split sections to be used in mid-sized truck beds.

DESIGN REQUIREMENTS AND SOLUTIONS

INITIAL REQUIREMENT	CHALLENGE	SOLUTION
2,000-lb. deck load	Using molded plastic alone was not strong enough.	Simulation revealed initial design would crack when driven over potholes under load; steel supports were added.
Manufacturability	Fully optimized design could not be manufactured.	Reworking design of steel support structure so molders could accommodate it.
250-lb. drawer load	Prototyping showed sliders not adequate easy drawer pulling.	Sealed bearing wheels were used. Requirement dropped to 200 lbs.
Shipped by less-than-truckload (LTL) freight	Fitting everything in a compact, dense carton.	Addressing shipping requirements as part of the up-front design process.
Price point	Ensuring DECKED could be sold for under $1,000.	Keeping part count low and using the same parts for many versions.
Easy attachment to truck bed	A no-drill design that would not require consumers to drill holes in their truck beds.	J-hooks and brackets designed to use the different tie down systems built into various pickup truck beds.
Durability, custom look and feel	Prototyping showed standard, commercial drawer latches did not provide desired look and feel.	Altair ProductDesign created a custom latch to provide the desired consumer experience.
Fit 5.5-, 6- and 8-ft. variants of mid- and full-sized pickup trucks from major OEMs manufactured since late '90s	Getting specific CAD data of truck beds. Balancing costs to manufacture different variations.	OEMs provided CAD data, or pickup truck beds were 3D scanned; abandoned mid-sized and 8-ft. bed length variants.

"It became clear we were going to be suboptimal in performance because we were trying to accommodate too many truck bed permutations — 26 in all — with one design," Smith says.

It was the 60/40 split, to accommodate the mid-size trucks, that was causing issues, Boggs said. Additionally, the 8-ft. bed length variant would increase tooling complexity. Given the fact that 8-ft. bed pickup trucks are only 3-5% of the full size truck market, DECKED decided to focus on 5.5- and 6.5-ft. beds, and abandon the mid-size truck bed permutations. That decision allowed Altair ProductDesign to split the deck in half in a patented design that uses just four tooling setups for 19 design variants.

"To keep the part count as low as possible, we came up with the idea to make the deck on the left the same as the one on the right, just turned 180°," says Smith. Just two different decking designs were needed to accommodate all of the different truck bed sizes. "The center line is interleaved in such a way that the same tooling can be used without modifications, plus it made it possible to meet DECKED's shipping requirements."

Optimization Meets Manufacturing

Altair's simulations showed a molded plastic deck alone would not be able to meet DECKED's 2,000-lb. payload requirements. It would be prone to cracking when driven over bumpy roads. The answer was to reinforce it with steel. Altair ProductDesign used HyperWorks OptiStruct simulation and optimization software to help place the steel members optimally.

"The steel members were going off on many different angles because the topology optimization said that was the best location for these pieces given the load you want," Smith says. "However, we engaged the molder very early in the process and they said flat out: 'We're not going to mold that part.'"

The molder knew the difference in thermal expansion of the steel and high-density polyethylene (HDPE) would lead to an uneven cooling rate longitudinally vs. transversely, so any steel member on an angle would cause the HPDE to crack. The molder was experienced with orthogonal design, and was reluctant to take on the spider web-like optimized design.

"They essentially had zero confidence they could mold it, so they wouldn't sign up for it," Smith said. "We had to rethink it."

In the end, the engineers went with an orthogonal arrangement of the steel supports, but still used Altair OptiStruct to help guide their decisions.

"There were hiccups, but we figured them out," Peters says. "They used software to figure out the optimal design and then used their own judgment to determine what's optimal and [can be] manufactured. That's how we got the design we did."

That design helped launch DECKED and it won the prestigious SEMA Award for Best New Van/Pickup/SUV Product for 2014 during the SEMA Show in Las Vegas. The automotive show features thousands of new automotive parts, tools and components each year — though the DECKED truck bed storage system was probably the only one signed by the engineers who were instrumental in bringing it to market. DE

Jamie J. Gooch *is editorial director of* DE. *Send e-mail about this article to DE-Editors@deskeng.com.*

INFO ➜ Altair ProductDesign: AltairProductDesign.com

➜ **DECKED:** DECKED.com

(Don't) Bring the Noise

Lightweighting and fuel economy standards have created new noise and vibration challenges for automotive engineers.

BY BRIAN ALBRIGHT

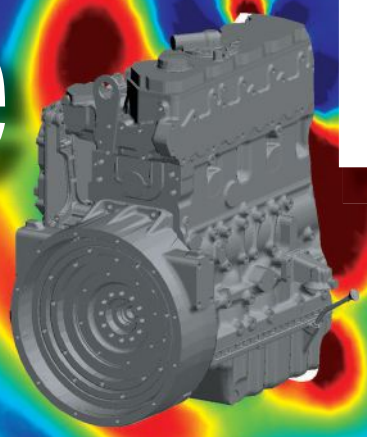

Powertrain acoustic radiation pressure maps help evaluate vehicle noise and vibration. ***Image courtesy of Free Field Technologies/MSC Software.***

The pressure is on automakers to make vehicles lighter and more fuel-efficient. The results have created new engineering challenges around controlling noise, vibration and harshness (NVH).

Increased power to weight ratios, engine downsizing and the switch between power sources in hybrid vehicles are creating more audible noise for drivers. Smaller engines generally create more noise, and automakers have compounded this issue by adding more turbochargers and superchargers, which further increase NVH problems. New materials have also altered the way noise and vibrations pass through the vehicle structure.

"In hybrid vehicles, when you remove the combustion noise, drivers start to hear all sorts of other things, and perceive sounds and vibrations differently. You can even hear fuel sloshing around," says Paul Weal, technical director at Siemens PLM Software. "There's another angle to that as well. Pedestrians can't hear these vehicles, so there are car companies looking for ways to pump noise from the vehicle to warn pedestrians."

Other changes in vehicle design have also produced new noise and vibration problems. "For example, you could change the material or the structure produced," says Diego Copiello, senior application engineer and product marketing manager with Free Field Technologies, an MSC company. "This changes the system dynamics, and leads to a possible increase in noise."

In the past, automakers responded to NVH issues later in the design process by adding noise-canceling materials or insulation. With an emphasis on lightweighting, however, that approach is impractical.

"Weight is normally good for NVH," says Simon Tate, NVH manager at Ricardo. "There is a lot of effort involved in maintaining or improving NVH attributes while at the same reducing vehicle mass and cylinder count on the powertrain."

Further complicating matters is the subjectivity of NVH issues: different drivers perceive NVH differently. "There is a whole science behind measuring NVH quality, and some even go to the extent of presenting audio recordings to a jury of drivers," says Ravi Shankar, director of Simulation Product Marketing for Siemens PLM Software. "There's not a single metric involved."

More Simulations, Earlier Simulations

An increase in NVH issues has led to an increase in demand for more simulations, much earlier in the design process. NVH analysis used to be a process that came toward the end of design, at a point when the cost to change the structure of the vehicle was high — now that process has been shifted to the concept stage.

The process, however, is still the same: "You have to predict or measure the sources of noise, and predict and measure their paths and how those forces and sources are propagating in the vehicle," Weal says.

"You're not only predicting what the noise and vibration levels will be, but also understanding why they are increasing, what the paths are, and how to guide designers and analysts in changing the design to improve NVH," says Mark Donley, NX Nastran product manager at Siemens PLM Software.

Simulations aren't just occurring earlier, they are also becoming more complex. Engineers need a view of how changes will affect the entire design. Multiple simulations have to be linked and conducted simultaneously to measure how noise and vibration

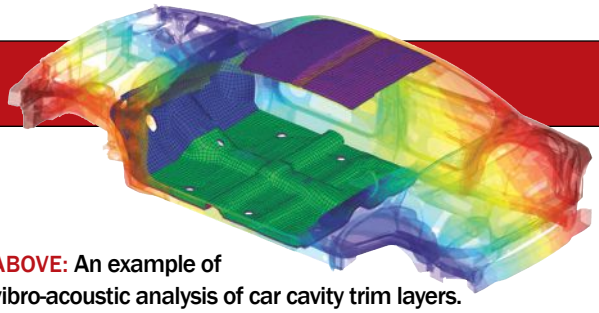

ABOVE: An example of vibro-acoustic analysis of car cavity trim layers. *Image courtesy of Free Field Technologies/MSC Software.*

RIGHT: Ricardo's Vehicle Anechoic Test Facility (VATF) includes a "rolling road" integrated chassis dynamometer. *Image courtesy of Ricardo.*

will travel through the vehicle. This type of dynamic multi-body system analysis allows engineers to better evaluate variables like engine mounting positions much earlier in the process.

Co-Simulation Helps Ford Reduce Drivetrain Noise

Drivetrains are often the source of NVH issues. Lugging is a type of vibration that occurs when a driver accelerates while the vehicle operates in a high gear at a low engine speed. Slipping the transmission can smooth it out, but at a cost in fuel economy.

Adjusting slip has traditionally been a process involving more art than science, with a lot of trial and error. Ford Motor Co. used a co-simulation method to make optimal adjustments without eating up a lot of engineering time using advanced simulation and modeling during the prototyping phase. Ford is using an open standard called Functional Mock-Up Interface (FMI) to combine simulations of the drivetrain, slip controller and full vehicle so engineers can see how slip settings affect noise and fuel efficiency.

Ford used MSC Software's Adams multi-body dynamics solution to create models of the drivetrain and vehicle (that served as the co-simulation master), and the AMESim simulation software to create a 1D model of the torque converter (the co-simulation slave). FMI enables the models to work together in a single simulation. The simulation helped identify an optimal slip of 40 rpm.

"Recent advances in co-simulation have allowed us to use different software to look at those interactions," Tate says. Ricardo has its own software, VALDYN, for multi-body dynamic and kinematic simulation for valvetrain and drive system analysis. "We use those techniques to develop valve lift profiles for low noise while maintaining the breathing characteristics of the powertrain. That wasn't possible in the past," he says.

"You need different sets of tools that work together to evaluate those system-level effects," Weal says. "We're definitely seeing more use of co-simulation and FMI."

Quiet Evolution of NVH Fixes

There is only going to be more demand for better NVH analysis and new approaches to reducing noise. Tate says NVH tools will evolve to provide simpler interfaces so that engineers and executives who are not acoustics experts can understand the results of the analysis through auralization. "We're seeing more tools with the ability to communicate how changes in design affect the creation of sound using actual simulated sounds," he says. "There are simulation tools that allow you to make changes at the powertrain or chassis level and see and hear how that affects the drive."

There is also more automation involved. "In these managed environments, the system can identify which parts have changed and automatically reassemble them into a full engine body model much more quickly than before," Donley says. "Before, it could take weeks to get those models and then find out that the change had negatively impacted NVH performance."

Auto manufacturers are also employing discontinuous Galerkin method (DGM) analysis for solving acoustics. "Previously those were available for solving specific aerospace acoustic issues," Copiello says. "Now these methods are a factor in automotive. Combined with GPU (graphics processing unit) accelerations, DGM can solve pass-by noise problems with simulation."

Many NVH issues can be addressed through physical alterations, but some noise-canceling solutions take a novel approach. A few automakers have employed active noise control that uses the car audio system to generate sound waves with the same amplitude (but inverted phase) to cancel out low-frequency noise. Siemens Automotive, for example, created an active noise control system that attaches to plastic air intake manifolds and can "listen to" and cancel out unwanted vehicle noises. While the engine is running, the system monitors noise through a microphone. Analog signals pass through an A/D converter to a digital signal processor (DSP), which determines their frequency. The DSP then produces an opposing output signal and sends it to a speaker.

The need for more advanced NVH analysis earlier in the design process will continue to accelerate. Six automakers have introduced vehicles with 3-cyinder engines augmented with turbochargers. Consumer acceptance of these vehicles hinges, in part, on how drivers perceive the quality of the ride, and the level of noise and vibration they generate. DE

Brian Albright *is a freelance journalist based in Columbus, OH. He is the former managing editor of* Frontline Solutions *magazine, and has been writing about technology topics since the mid-1990s. Send e-mail about this article to DE-Editors@deskeng.com.*

INFO ➜ **Free Field Technologies:** FFT.be

➜ **Ricardo:** Ricardo.comk

➜ **MSC Software:** MSCSoftware.com

➜ **Siemens PLM Software:** Siemens.com/plm

3D Printing Thinks Really Big

Additive manufacturing systems offer amazing build volumes.

BY PAMELA J. WATERMAN

There was a time when 3D printing a 6-in. long car model was a mind-blowing process. Now companies are building full-size car bodies, functional bridges, live-in houses and operational jet engines. Additive manufacturing (AM) has become not just big business but a business of big proportions.

Thinking about scaling up your own designs? *DE* scoped out industrial-strength AM systems that can take on the big jobs, using build-volumes with a minimum of just about 1 meter (39 in.) dimension in one or more of the X/Y/Z build directions. Interestingly, most systems met that criterion for X or Y, or both X and Y values; however, a few systems qualify here based strictly on their extended Z-axis capacity. Also, with most equipment optimized for a single material family, the listings are grouped under plastics-plus, metals, and sand/ceramics applications. We're leaving the bridges and buildings for another day.

The large-scale XLine 2000R laserCUSING metal AM system from Concept Laser is just shy of a 1 meter build dimension. ***Image courtesy of Concept Laser.***

Super-Sizing that Plastic, Rubber or Wood-like Part

3D Systems – ProX 950 (1500x750x550 mm; 59x30x22 in.) Perhaps it's no surprise that, given its decades of experience as the original AM company, 3D Systems markets a large-volume system called the ProX950 production unit. This is a stereolithography (SL) system that can run 13 different UV-curable liquid resins from the company's Accura line, in black or white ABS (acrylonitrile butadiene styrene) as well as in materials that are flexible, clear, tough or castable.

3DP Unlimited – 3DP1000 (100x100x500mm; 39.3x39.3x19.7 in.) This open-bed printer comes from a company experienced with controlling accurate industrial-strength linear motion. Printer type is fused filament fabrication (FFF) with layer resolution down to 70 microns. Three-millimeter filaments handled include, but are not limited to, polylactic acid (PLA), ABS, polycarbonate (PC), nylon, Ninja Flex and high impact polystyrene (HIPS).

BigRep – BigRep ONE (1100x1050x980 mm; 43.3x41.3x38.6 in.) Offering a 1-cubic-meter build volume, this FFF system features ruggedness and accuracy aimed at industrial customers. The new, modular, dual-nozzle extruder can print support material and multiple colors; the universal spool-holder fits all filament spool sizes and holds up to four 10 kg (22 lb.) rolls in an enclosed cabinet. An improved heat-bed comes with auto-leveling.

Cincinnati Inc. – Big Area Additive Manufacturing (BAAM) 100 Alpha Size 2 (6096x2362x1829 mm; 240x93x 72 in.) and BAAM 100 Alpha Size 1 (4000x2000x864 mm; 157.5x78.75x34 in.) Entering the world of AM backed up by more than 100 years' experience with metal fabrication processes, Cincinnati Inc. offers two extremely large plastics-based 3D printers that can extrude 38 to 100 lbs./hour of pellet-feedstock thermoplastics.

Cosine Additive – Additive Machine 1 (AM1) (1100x850x900 mm; 43.3x33.5x35.4 in.) This Houston-based company is big on composites. Its AM1 FFF system supports any third-party filament, but the buzz surrounds its use of polycarbonates. The system can process filaments containing a wide variety of additives including chopped and continuous carbon fiber, glass fiber, glass spheres, mica and stainless steel and bronze powders.

German RepRap – X1000 (1000x800x600 mm; 39.4x 31.5x23.6 in.) Industrial users will be pleased with this larger version of the company's previous FFF systems. Based on the open-source RepRep designs, the X1000 features an enclosed unit with a touch display, browser-based user interface and integrated nozzle cleanup function.

Massivit3D – Massivit3D (1500x1200x1800 mm; 59x47.2x70.9 in.) This Israeli company, founded in 2013 by experienced 2D and 3D printing executives, managers, scientists and engineers, has developed an AM process called gel dispensed

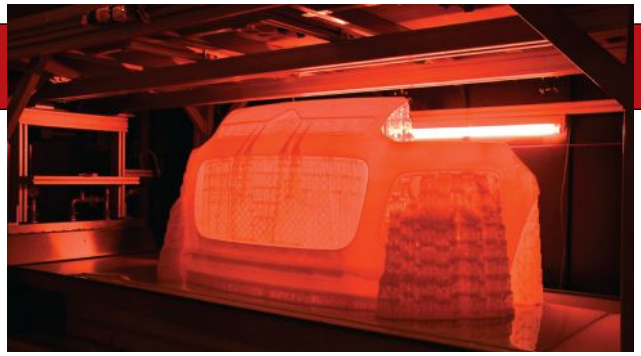

Full-size resin vehicle bumper printed on a Mammoth SLA AM system at Materialise. ***Image courtesy of Materialise.***

printing (GDP). The debut system can already print a UV-cured gel, using a shell/support approach, at rates of up to 1000 mm/sec (39 in./sec) on both X and Y axes.

Materialise – Mammoth (2100x700x800 mm; 82.7x27.6x31.5 in.) The Mammoth system from Materialise in Belgium is a stereolithography apparatus (SLA) resin system; parts are available through the company's service bureau only, but the company operates 14 of these enormous systems with typical lead times of four to eight working days. Currently offered materials are functional white, grey and transparent ABS-like plastics.

Sicnova – JCR 1000 (1000x600x600 mm; 39.4x23.6x23.6 in.) Sicnova of Spain began in 2007 as a 3D printing consulting and reseller business, but decided that the market needed a larger FFF system. In March 2015, Sicnova introduced the JCR 1000, an enclosed, heated environment system with two print heads and interchangeable nozzles; the system can print with a large range of thermoplastics, nylon, rubber and carbon fiber materials.

Stratasys – Objet1000 Plus (1000x800x500 mm; 39.4x31.5x19.7 in.) In 2015, Stratasys added yet another PolyJet technology system to its line-up, this time extending its Objet line of multi-material printers with the Objet1000 Plus. Dual jetting allows combining base resins that are UV-cured into composite digital materials with properties ranging from ABS-like to rubber-like, with a wide range of colors and transparencies.

Voxeljet – VX800 (850x500x1500/2000 mm; 33.5x19.7x59/78.7 in.), VX1000 (1060x600x500 mm; 41.7x23.6x19.7 in.), VX2000 (2060x1060x1000 mm; 81.1x41.7x39.4 in.) and VX4000 (4000x2000x1000 mm; 157.5x78.7x39.4 in.) Each model of these large-scale AM systems from Voxeljet is available for purchase or for parts/mold production at the company's service bureau facilities. Parts made by 3D printing a binder onto powdered plastic can be infiltrated with epoxy or wax; applications include prototypes or masters for lost-model casting.

Mega Metal AM Systems

BeAM – MAGIC (1500x800x800 mm; 59.1x31.5x31.5 in.) and CLAD (1000x700x700 mm; 39.4x27.6x27.6 in.) The configuration of three- to five-continuous axes lets these directed-energy systems manufacture or repair metal parts using laser-melted metal powders. The MAGIC machine features a controlled atmosphere as standard; this is an option on the CLAD unit.

Fabrisonic – 7200 (1829x1829x914 mm; 72x72x36 in.) Fabrisonic's ultrasonic additive manufacturing (UAM) process builds up parts from thin metal sheets, while permitting separate parts

Sciaky's EBAM 300 series metal 3D printer can produce parts that are 19 feet long, 4 feet wide and 4 feet high. ***Image courtesy of Sciaky.***

A massive plastic-extrusion additive manufacturing system is in development by Thermwood, a traditional CNC equipment manufacturing company, is pictured. ***Image courtesy of Thermwood.***

Know of any big additive manufacturing systems that didn't make our feature this month? Please submit them to de-editors@deskeng.com.

and materials to be inserted in-process, from wires and electronic sensors to composite fibers and shape-memory alloys. The system can join fine layers of multiple metals without forming brittle intermetallic junctions.

Norsk Titanium – (1200x1200x1800 mm; 47.2x47.2x70.9 in.) Norsk Titanium has developed a plasma-arc-based direct metal deposition technology to melt titanium wire for near-net shape manufacturing, as well as to repair titanium parts. The company recently announced plans to open a 200,000 sq.-ft. facility in the U.S., where it will operate several dozen of the units in a public/private partnership to produce aircraft parts.

Optomec – 850-R (900x1500x900 mm; 35.4x59.1x35.4 in.)

MAGIC LF 6000 additive manufacturing system from BeAM. The company's CLAD technology creates or repairs parts with laser-melted metal powders. ***Image courtesy of BeAM.***

LENS AM technology, based on laser melting a focused stream of powdered metals, supports repair, rework and modification of high-value industrial components. The 850-R features five-axis motion and deposits up to 0.5kg/hour of metals such as titanium, nickel, stainless steel, cobalt, copper and aluminum.

Plus Mfg – +1000K (1200x900x600 mm; 47.2x35.4x23.6 in.) Wire feedstock is used in this fully enclosed, arc metal deposition (AMD) system for creating near-net-shape parts. The company's goal is to help users create parts quickly and with low energy draw; typical deposition speed is more than 2.5 kg/hour (upgradeable), currently processing steels, aluminum and titanium.

Sciaky – EBAM 300 (5791x1219x1219 mm; 228x48x48 in. (typical)) Six decades of experience with electron beam welding led to Sciaky's development of Electron Beam Additive Manufacturing (EBAM) systems as a fast, economical alternative to forging and casting of large parts. Each system is custom made. Materials processed include wire feedstock of titanium, titanium alloys, stainless steels, Inconel, tantalum, tungsten and niobium.

Sand and Ceramics for Really Big Parts

3Geometry – DSM1200 (1200x1200x420 mm; 47.2x47.2x16.5 in.) Based in India, 3Geometry markets a family of AM systems developed for tool-less manufacturing. The DSM1200, available in regular and high speed build versions (3600 cm3/hour or 10,000 cm3/hour) uses laser sintering technology to create sand-based molds for metal casting. Different sands are available for different end-metal applications.

ExOne – Exerial (Two build sections, each 2200x1200x700 mm; 86.6x47.2x27.6 in.) and SMax (1800x1000x700 mm; 70.9x39.4x27.6 in.) – ExOne systems use binder-jet technology to print sand cores and molds for casting parts. The SMax has a single job box; the Exerial model includes two build sections within the same system for continuous part production and processing.

Prodways – ProMaker V4000 (1060x2250x2040 mm; 42x89x81 in.) This system is the largest of the ProMaker V series from Prodways, a company of Groupe Gorgé based in France. UV laser technology supports curing of highly viscous ceramic materials, creating parts to be used as prototypes or in short production runs. When fired, the finished material is high density and highly homogeneous.

3D Systems' ProX 950 large format SLA 3D printer, with a Quickcast 3D-printed casting pattern. ***Image courtesy of 3D Systems.***

Voxeljet – These large systems from Voxeljet, as mentioned above, can 3D-print a binder onto plastic powders and sand materials, the latter for creating cores and molds for casting. The company's phenolic resin binders allow use of untreated silica sands with different pH values and ceramic materials. DE

Editor's Note: *Read more on this topic at deskeng.com/de/big.*

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

INFO
- **3Geometry:** 3Geometry.com
- **3D Systems:** 3DSystems.com
- **3DP Unlimited:** 3DPUnlimited.com
- **BeAM:** Beam-Machines.com
- **BigRep:** BigRep.com
- **Cincinnati Inc.:** E-CI.com
- **Concept Laser:** ConceptLaserInc.com
- **Cosine Additive:** CosineAdditive.com
- **ExOne:** ExOne.com
- **Fabrisonic:** Fabrisonic.com
- **German RepRap:** GermanRepRap.com
- **Materialise:** Manufacturing.Materialise.com
- **Norsk Titanium:** NorskTitanium.no
- **Optomec:** Optomec.com
- **Plus Mfg:** Plus-Mfg.com
- **Prodways:** Prodways.com
- **Sciaky:** Sciaky.com
- **Sicnova:** Sicnova3D.com
- **Stratasys:** Stratasys.com
- **Voxeljet:** Voxeljet.com

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

CD-adapco Enhances Product Design with Faster Simulations

The company selects Panasas to update its IT for Multidisciplinary Design eXploration.

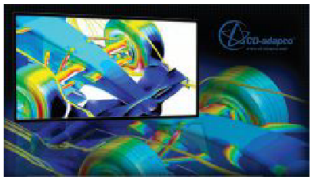

Simulation allows engineers to test new ideas using software models that mimic the behavior of real-world objects, which greatly simplifies the process of trying out alternatives.

Models are only approximations of the real world. As designers need more accurate simulations, they push the limits of technology. That's where CD-adapco comes in. The company is developing an approach known as Multidisciplinary Design eXploration (MDX). It involves multiple physics domain simulations, all within a single software tool, each addressing a specific part of the problem that allow an engineer combine each domain to simulate the entire system as a whole.

Breaking the Bottleneck

Not surprisingly, MDX demands much higher performance from technical infrastructure. "When we first started doing just CFD analysis by itself a few years ago, a single simulation would stop everything else running on the system," says Steven Feldman, senior vice president of IT at CD-adapco. His staff traced the problem to legacy network-attached storage (NFS), which could not keep up with the input/output (I/O) requirements of the highly data-intensive CFD (computational fluid dynamics) calculations. Systems that couldn't keep up with CFD-only analyses had no chance of performing MDX calculations well.

Feldman's team began its search for a storage system that would eliminate the I/O bottleneck. To avoid overtaxing the company's thin IT staff, the new system had to be easy to use, highly reliable and backed by additional support. A search of commercial storage vendors led Feldman's team to develop a short list that included Panasas. Now it was time to put the vendors through their paces.

MORE ➜ deskeng.com/de/?p=26616

ANSYS and Fujitsu Accelerate Simulation in the Cloud

Targeted cloud technology and parallel computing alleviate Chiyoda's taxed IT infrastructure for simulation.

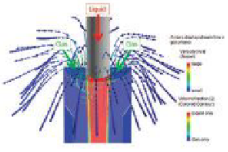

Engineering firms are challenged to provide optimal customer service, while keeping staffing, IT resources and other investments low. As simulation grows as a core engineering competency, it can be difficult for consulting firms to accommodate the shifts in IT demand. Renting additional server space can take up to a week to acquire the technology; a pace that is too slow for a fast-based business.

Business Challenges

Chiyoda Corporation, a leading Japanese engineering company, relies on ANSYS Fluent to attack engineering challenges for clients in the global energy business. Chiyoda's internal IT resources were too overtaxed to process large, complex simulations quickly. To create a more flexible IT infrastructure — and help make the most of its ANSYS HPC Pack licenses — ANSYS introduced Chiyoda to Fujitsu Ltd., a partner offering cloud computing resources and expertise. Chiyoda implemented ANSYS Fluent, ANSYS CFD-PrepPost, ANSYS HPC Pack, Fujitsu Technical Computing Cloud and the UberCloud Experiment.

MORE ➜ deskeng.com/de/?p=26611

Advertising Index

★ SPOTLIGHT ★

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

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.

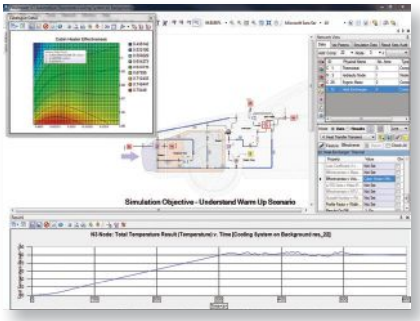

Mentor Graphics Releases Flowmaster 7.9.4

Software now includes a Multi-Arm Tank Component for heat transfer models.

System-wide simulation takes on new meaning beginning with Flowmaster 7.9.4. For the first time, according to Mentor Graphics, Flowmaster offers integration with the third-party pipe stress analysis tools CAEPIPE and CAESAR II from SST Systems and Intergraph Corp, respectively.

Version 7.9.4 has two other notable enhancements. One is a Multi-Arm Tank Component that lets users choose heat transfer models: polytropic or full heat transfer. Two, an enhanced Experiments capability lets users modify key system parameters and make immediate changes.

MORE ➜ deskeng.com/de/?p=25889

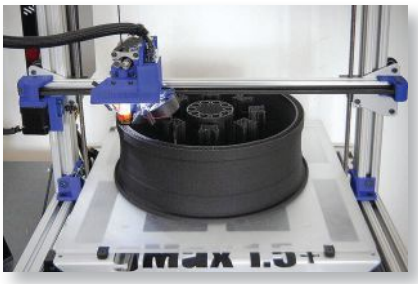

gCreate 3D Printers Offer Big Build Areas

The new gMax 1.5 series runs on a variety of model slicing software.

gCreate, an outfit in the emerging hotbed of additive manufacturing technology known as Brooklyn, NY, recently launched its gMax 1.5+ and gMax 1.5 XT+ series of desktop-sized 3D printers.

The gMax 15+ series 3D printers use FFF (fused filament fabrication) technology.

Print volumes are 16 x 16 x 12 in. for the gMax 1.5+ and 16 x 16 x 21 in. for the gMax 1.5 XT+, respectively. That's big enough to 3D print a 14-in. tire rim.

The systems are compatible with approximately 50 materials, such as ABS, PLA, Nylon and NinjaFlex.

MORE ➜ deskeng.com/de/?p=26034

BOXX Technologies Supports Virtualization

The company's XDI series is suited for heavy design engineering workloads.

The XDI (XTREME Desktop Infrastructure) virtual workstation series from BOXX Technologies provides any outfit with the horsepower to maintain peak application performance while hosting up to eight concurrent engineers on a single, centralized server. They offer all the benefits of virtual desktop infrastructure — remote application access, security, high-speed shared storage and so on — optimized to handle the compute-intensive applications and large data sets that solutions like CATIA, Solid Edge and 3ds Max generate routinely.

XDI also comes equipped with multiple Xeon processors and PCoIP capability.

MORE ➜ deskeng.com/de/?p=26088

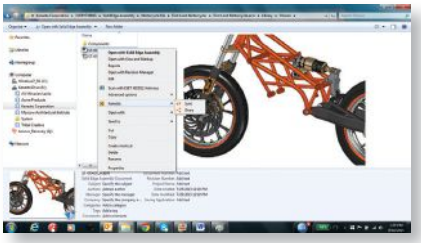

Kenesto Updates its Cloud-Based LAN Drive

Functionality goes beyond a traditional network drive or general-purpose cloud.

The new round of Kenesto Drive updates includes functionality traditionally only available through physical LAN-mounted drives and PDM systems. The big news here is an auto-lock feature. What this does is block other users from editing a document you have opened. They still have read-only access to in-use documents, and Kenesto automatically unlocks the document when a user is finished. This all means that users can modify documents, including CAD files, without worrying about accidental overwrites or wondering how you're going to buy more drives.

MORE ➜ deskeng.com/de/?p=26185

Make the Connection

Find the simple way through complex EM systems with CST STUDIO SUITE

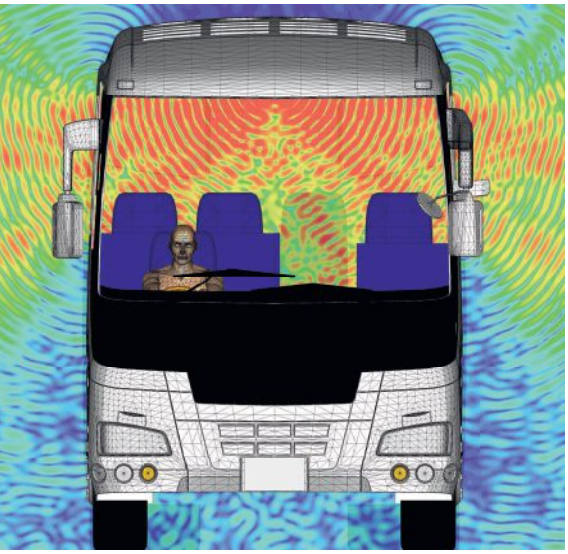

Components don't exist in electromagnetic isolation. They influence their neighbors' performance. They are affected by the enclosure or structure around them. They are susceptible to outside influences. With System Assembly and Modeling, CST STUDIO SUITE helps optimize component and system performance.

Involved in antenna development? You can read about how CST technology is used to simulate antenna performance at **www.cst.com/antenna**.

If you're more interested in filters, couplers, planar and multilayer structures, we've a wide variety of worked application examples live on our website at **www.cst.com/apps**.

Get the big picture of what's really going on. Ensure your product and components perform in the toughest of environments.

Choose CST STUDIO SUITE – Complete Technology for 3D EM.

CST

FROM MODEL

TO APP

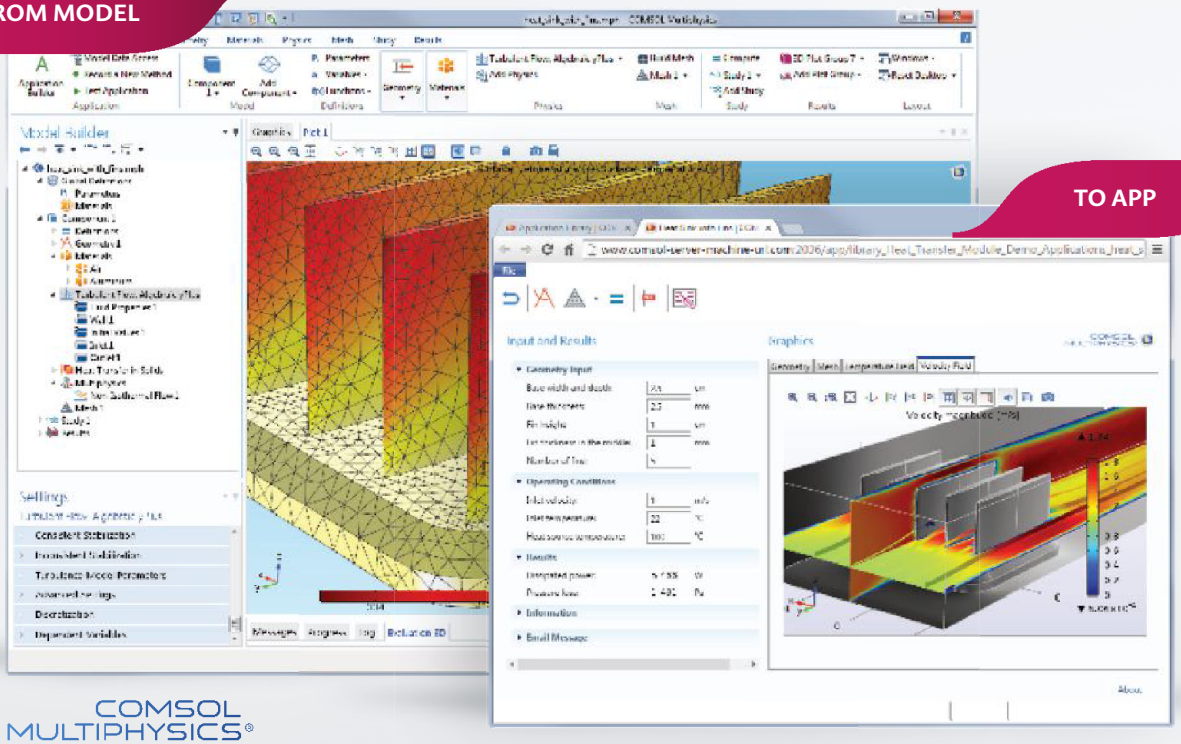

COMSOL MULTIPHYSICS®

COMSOL SERVER™

Verify and Optimize your Designs with COMSOL Multiphysics®

NOW FEATURING APPLICATION BUILDER & COMSOL SERVER™

The Application Builder provides you with tools to easily design a custom interface for your multiphysics models. Use COMSOL Server™ to distribute your apps to colleagues and customers worldwide.

Visit comsol.com

Product Suite

› COMSOL Multiphysics®
› COMSOL Server™

ELECTRICAL
› AC/DC Module
› RF Module
› Wave Optics Module
› Ray Optics Module
› MEMS Module
› Plasma Module
› Semiconductor Module

MECHANICAL
› Heat Transfer Module
› Structural Mechanics Module
› Nonlinear Structural Materials Module
› Geomechanics Module
› Fatigue Module
› Multibody Dynamics Module
› Acoustics Module

FLUID
› CFD Module
› Mixer Module
› Microfluidics Module
› Subsurface Flow Module
› Pipe Flow Module
› Molecular Flow Module

CHEMICAL
› Chemical Reaction Engineering Module
› Batteries & Fuel Cells Module
› Electrodeposition Module
› Corrosion Module
› Electrochemistry Module

MULTIPURPOSE
› Optimization Module
› Material Library
› Particle Tracing Module

INTERFACING
› LiveLink™ *for* MATLAB®
› LiveLink™ *for* Excel®
› CAD Import Module
› Design Module
› ECAD Import Module
› LiveLink™ *for* SOLIDWORKS®
› LiveLink™ *for* Inventor®
› LiveLink™ *for* AutoCAD®
› LiveLink™ *for* Revit®
› LiveLink™ *for* PTC® Creo® Parametric™
› LiveLink™ *for* PTC® Pro/ENGINEER®
› LiveLink™ *for* Solid Edge®
› File Import *for* CATIA® V5

COMSOL