DE
Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

November 2014 / deskeng.com

MAKING THE CASE FOR PROFESSIONAL ENGINEERING WORKSTATIONS P. 8

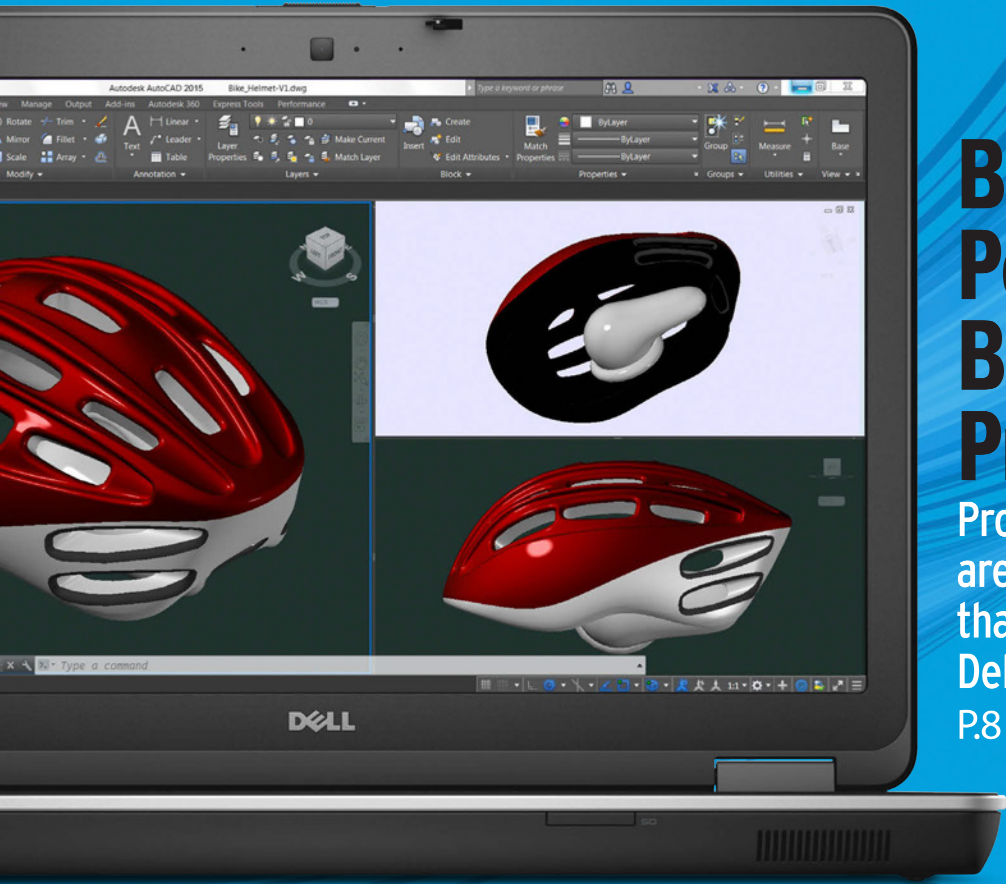

Breakthrough Performance, Breakthrough Price

Professional workstations are more cost effective than ever thanks to the Dell Precision M2800.

P.8


Infiniti Red Bull Racing Wins
for the 4th Consecutive Year

DE

Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

November 2014 / deskeng.com

The Future of Mobile & HPC

Infiniti Red Bull Racing Wins for the 4th Consecutive Year

2010, 2011, 2012, and 2013 Formula One World Constructors' and Drivers' Champions

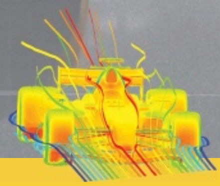

ANSYS®

Realize Your Product Promise®

Winning the Formula One World Constructors' Championship is a monumental accomplishment. Winning it four times in a row is practically unheard of. But for Infiniti Red Bull Racing, it was just another day at the office.

Using ANSYS simulation technology, Infiniti Red Bull Racing is creating virtual prototypes of its race cars, so engineers can quickly and inexpensively optimize everything from aerodynamics to brakes to exhaust systems. Infiniti Red Bull Racing is delivering on its product promise by remaining dominant in one of the most competitive environments imaginable.

Not a bad day at the office.

INFINITI Red Bull RACING | ANSYS INNOVATION PARTNER

Visit ANSYS.COM/redbull to learn how simulation software can help you realize your product promise.

Degrees of Freedom

by Jamie J. Gooch

The Future of Mobile & HPC

The future of mobile could be described as simply "the future," because just about everything is becoming a mobile data capture and communication device — from wristwatches to cars to eyeglasses. Even objects that are decidedly fixed in one location, such as your front door lock or the factory assembly line, are borrowing design elements from mobile devices, including sensors, radios, antennas and batteries.

At its most basic level, the ubiquity of mobile technologies has two ramifications for design engineers: 1. They need to employ specific design, simulation and testing techniques to create mobile products; and 2. They can do so by tapping into an increasing number of high-performance computing (HPC) resources — from HPC servers they can now access remotely, to HPC mobile workstations. The computing resources engineers need to usher in smaller, connected devices are always at their fingertips.

The future of mobile and HPC are inextricably linked.

Designing for Mobile

Our connected lifestyles have already had a huge impact on design technologies and design cycles. More products are downsized for mobility and packed with technologies that require specialized engineering skills, multiphysics simulations and more powerful computing hardware to design. Electronic and software engineers need to collaborate with mechanical engineers to ensure connected products perform as specified, leading many design software developers to make collaboration a cornerstone of their applications. And data being collected by connected devices, some of which already comprise the early stages of the Internet of Things (IoT), needs to be stored, filtered and turned into actionable information that can be used to improve product designs.

According to the Cisco Visual Networking Index, traffic from wireless and mobile devices will exceed traffic from wired devices by 2018. The company expects there to be an 11-fold increase in global mobile data traffic between 2013 and 2018. That represents a lot of people already using a lot of mobile products — 7 billion connected devices by some counts. And even though the number of people connected to the internet has grown by more than 600% since 2000, more than half the population isn't online yet. Add the number of future connected products to the number of people soon to be connected, and it's easy to see we're just getting started.

To respond to the need for increasingly complicated mobile products, design engineers must have access to increasingly powerful computers. They also need to develop or specify the specialized, embedded systems that will enable future smaller, low-power mobile devices. The future of mobility and HPC are inextricably linked.

Mobile Engineering

Fortunately, HPC choices abound. From workstations, both virtualized and local, to private and public cloud resources to clusters and data centers — innovations in engineering computing are enabling mobility, just as more advanced embedded systems are meeting engineering design needs. The adoption of advanced computing solutions by design engineers has led to an exponential increase in the rate of innovation.

Breakthroughs are happening at breakneck speed, which means design cycles are further compressed. Here again, mobility and computing are merging to allow engineers to work more productively, no matter where they are.

Engineers aren't just the designers, developers, testers and analysts of mobile technologies, they're also its users. Mobile engineering is quickly evolving from its roots in hauling workstations to the boardroom or reviewing designs via a tablet with offsite clients. It's now possible to use mobile devices to manipulate large CAD files and run complex simulation via virtual machines and the cloud, not to mention the previously unheard of power and speed now commonplace in mobile workstations.

As the muscle of supercomputers that fill entire buildings is being democratized and brought to bear on creating smaller, lighter, always-connected products, a final challenge remains the user interface. Many engineering tasks — from collecting test data to 3D scanning to collaborating on designs are already mobile. Actual design and simulation work, however, is still best accomplished with a keyboard and mouse, often on large screens or multiple monitors. Early attempts to port the design experience to tablets and phones have not been widely accepted. But as mobile devices become more ubiquitous, expect to see more specialized apps that allow more design engineering tasks to be completed — not just on phones and tablets, but via a number of mobile devices that work together to take advantage of each one's user interface strengths, whether that is touch, voice, motion tracking or a heads-up display. **DE**

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

Smarter Embedded Designs,
Faster Deployment

The combination of NI LabVIEW system design software and reconfigurable I/O (RIO) hardware helps small design teams with varied expertise develop demanding embedded applications in less time. Using this graphical system design approach, you can take advantage of the same integrated platform to program embedded processors and FPGAs for faster application development in industries ranging from energy to transportation, manufacturing, and life sciences.

LabVIEW system design software offers ultimate flexibility through FPGA programming, simplifies code reuse, and helps you program the way you think–graphically.

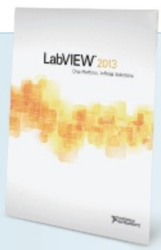

>> Accelerate your productivity at ni.com/embedded-platform

800 453 6202

NATIONAL INSTRUMENTS

DE November 2014 VOLUME 20/ISSUE 3

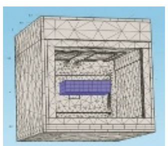

ON THE COVER: *Images courtesy of iStock.com, Microsoft, Apple, HP and Lenovo.*

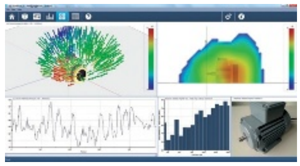

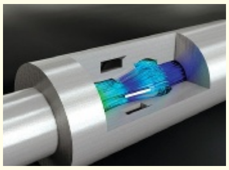

5.0 NOW AVAILABLE
comsol.com/release/5.0

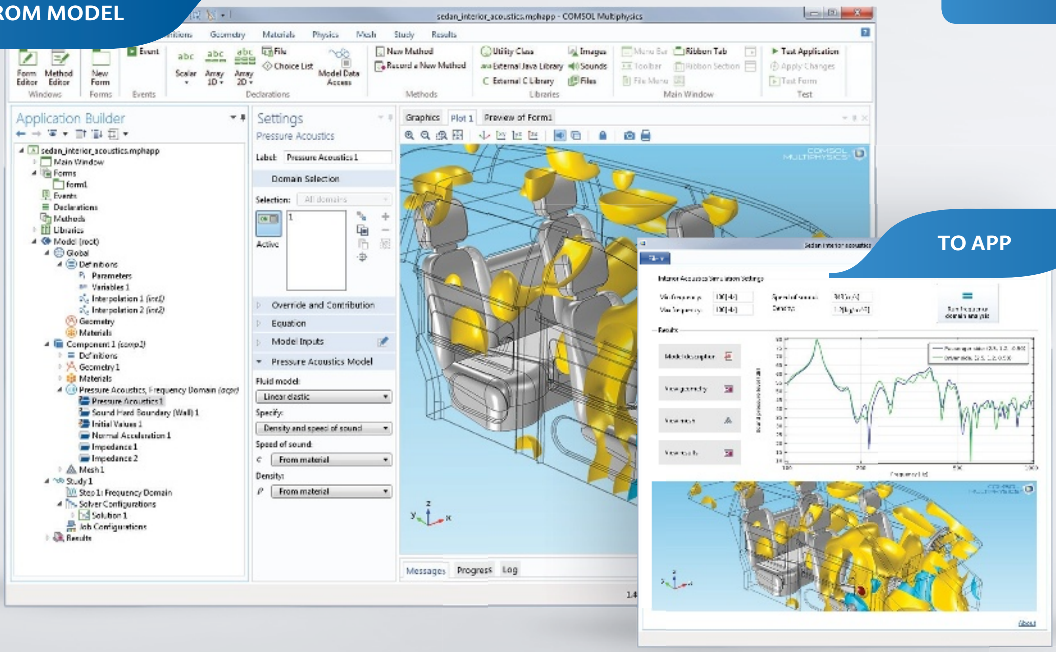

Verify and Optimize your Designs with COMSOL Multiphysics®

NOW FEATURING THE APPLICATION BUILDER

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/release/5.0

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

DEPARTMENTS

Visit Deskeng.com

- **Navigation**
  Quickly find the content you're looking for in dedicated Design, Simulate, Test, PLM, Prototype/Manufacture, Engineering Services and Engineering Computing sections. Drill down deeper via individual topic areas, technology vendors and authors within each section.

- **Integration**
  All of *Desktop Engineering*'s content — from videos to Engineering Services Directory listings, to our blogs and social media posts — appear in our news feeds throughout the site.

- **Presentation**
  Larger images, text and video players with a neutral color scheme make it easier to interact with the site's content.

Check out the new deskeng.com and let us know what you think.

Desktop Engineering®

EXECUTIVE EDITOR & PUBLISHER
Tom Conlon

EDITORIAL
Jamie J. Gooch | Editorial Director
Kenneth Wong | Senior Editor
Anthony J. Lockwood | Editor at Large
Jess Lulka | Assistant 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
603-563-1631 • Fax 603-563-8192
Erich Herbert | Sales Executive (x263)

ART & PRODUCTION
Darlene Sweeney | Director (x257)

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

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

www.peerlessmedia.com

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

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

Desktop Engineering® (ISSN 1085-0422) is published monthly by Level 5 Communications, Inc., 1283D Main Street, P.O. Box 1039, Dublin, NH 03444, 603-563-1631. Periodicals postage paid at Dublin, NH, and at additional mailing offices. Desktop Engineering® is distributed free to qualified U.S. subscribers.

SUBSCRIPTION RATES: for non-qualified; U.S. $108 one year; Canada and Mexico $126 one year; all other countries $195 one year.

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

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

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

THE NEW MULTIFUNCTION SURECOLOR® T-SERIES

- Produce presentation-quality 24" x 36" color copy in under 40 seconds*
- High productivity with color scanning speeds up to 6" per second*
- Advanced image processing hardware for superb image quality
- Scan to file, email and remote SureColor T-Series multifunction printers
- Single-roll or Dual-roll models available in 36" and 44" print widths

Epson® SureColor T-Series

MFP Configuration Starting at $8,495

EPSON
EXCEED YOUR VISION

*Speeds are based upon print engine speed only. Total throughput times depend upon factors such as computer, file size, printer resolution, ink coverage, and networking. Prices are MSRP, before rebates. Please check with an Epson Professional Imaging Authorized Reseller for actual price as dealer prices may vary. EPSON and SureColor are registered trademarks and EPSON Exceed Your Vision is a registered logomark of Seiko Epson Corporation. Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Copyright 2014 Epson America, Inc.

Adobe® PostScript® 3™

EPSON.COM/**PLOTTERINFO**

The Case for Professional Engineering Workstations

Professional workstations are crucial to design engineering, and they're affordable.

There is a tremendous opportunity for you to create more innovative product designs more efficiently, providing incredible value to your company, all by simply upgrading to a professional engineering workstation. Professional workstations are more powerful and more affordable than ever, but not everyone realizes the impact they can have. Make the case for professional engineering workstations by focusing on what they would mean to you, the design engineer, your colleagues in the information technology department, and your managers.

Pro Tools for Professionals

Every job performed by design engineers and their co-workers in other departments is important, but they don't all require the same tools. Andy Rhodes, executive director of Dell Precision workstations, puts it this way:

"A surgeon doesn't wake up in the morning and pick up a kitchen knife and go to work to perform brain surgery," he said when introducing the new Dell Precision M2800 mobile workstation. "They use a scalpel. The two things are both knives, but one is a professional tool for a professional job. That's the same thing happening on the workstation side. They are professional engineers, but they don't have the professional tools to do their jobs."

Without the right tools, design engineers are frustrated by small delays that pile up as they wait for a model to rotate or miss key details while a visualization stutters. What's worse, today's larger, more complex models can overwhelm older hardware and cause inadequate systems to crash.

The M2800 and its desktop counterpart, the Dell Precision T1700 promise workstation performance at a breakthrough price. For design engineers, that means they are built with professional-grade Intel processors, a choice of GPUs, higher-quality memory, and redundant storage to improve reliability and performance. In fact, the M2800 was shown to perform better than a Dell Latitude business class laptop equipped with the same processor and RAM in all the verticals tracked by the SPECwpc benchmarks, including manufacturing and engineering.

Professional workstations are also independent software vendor (ISV) certified to run the latest engineering software, such as AutoCAD, SolidWorks and Creo. In fact, Dell's ISV certifications can cover 90% of the market on some models. The Dell Precision Optimizer, which isn't offered on consumer-level PCs, automatically ensures those ISV applications are running at their maximum potential by fine-tuning Dell Precision workstations beyond standard factory settings. For example, SPECapc benchmarks show the Dell Precision Optimizer can boost the performance of PTC's Creo by 121%.

IT Can Empower Power Users

The Information Technology department is a key ally in securing professional workstations for design engineers. They know engineers are power users with different computing needs than other departments. What they might not know is how professional workstations can save IT time, resources, and potentially money.

Workstations are designed for easy deployment, maintenance and expansion. Details like the handles, easy-access front hard drives and tool-less features on some desktop models save deployment and expansion frustrations. Likewise, professional workstations are also designed to require less maintenance. For example, the M4800 and M6800 mobile workstations are equipped with dual fans and an advanced cooling design for better reliability. With better quality components, workstations can work 24x7, improving efficiencies by speeding up design work during the day and running complex analyses at night.

Professional workstations also offer options to further decrease downtime. Dell's Reliable Memory Technology (RMT), for one, will quarantine memory errors discovered by Error Code Correcting memory. After a reboot, RMT prevents the workstation from writing to areas of bad memory. That reduces system crashes, lost data, IT service calls and improves productivity and extends memory life. Plus, the Intel® Xeon® E3-1200 v3 product family includes Intel® vProTM technology, which makes it as secure and manageable as any PC in your organization's fleet.

Intel caching accelerating software for workstations

Make Your Case

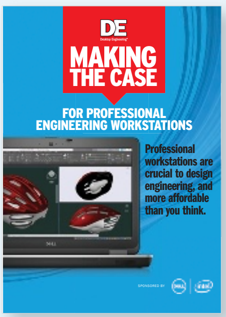

To make the full argument for going pro, download "Making the Case for Professional Engineering Workstations." This free paper was produced by *Desktop Engineering* and sponsored by Dell and Intel. It details the engineering productivity enhancements inherent in professional workstations, breaks down the return on investment and provides real-world examples of how companies have become more competitive and innovative thanks to professional workstations. **Download it here: deskeng.com/de/proworkstations.**

(CAS-W) is another feature exclusive to professional Dell Precision workstations. CAS-W provides increased hard drive storage performance — near solid-state drive speeds — without the cost of SSDs and without the need for IT to spend time configuring applications.

In addition to workstation performance tuning, the Dell Precision Optimizer is also a boon to the IT department. It can save IT administrators time by categorizing updates, enabling automatic updates of Dell-certified drivers and allowing administrators to filter updates by type and urgency. It also allows engineers to easily monitor their own systems' utilization and send analysis reports to IT who can then verify whether they have the correct system configuration.

Screen image courtesy of Autodesk.

The Dell Precision M2800 is a professional engineering mobile workstation that, as of Oct. 1, 2014, could be purchased for as little as $1,079.

Executive Decisions

Executives and managers want to boil productivity gains into hard numbers. They want to know how fast an investment will pay off in accelerated time to market, competitive advantage, and better quality products that help eliminate recalls and returns.

Companies can't profit from compute-intensive engineering workflows like simulation-led design if engineers are forced to take a coffee break every time they try to manipulate a large CAD file or run a simulation.

"The world is still in this place where sometimes a hardware purchase is not being allowed for the right set of technology because of cost," said Dell's Rhodes. "If you just look at a CAD engineer with an average salary of $70,000, look at the ROI of providing them with a first-class workstation. The time it saves them while doing their job — it pays for itself in about three to four weeks."

And that's a conservative estimate. Dell pitted two Dell Precision T1700 configurations against a standard Dell OptiPlex 9020 desktop on a SPECviewperf 11 benchmark using, PTC Pro|ENGINEER, Dassault Systèmes SolidWorks and Autodesk Maya datasets. The T1700 entry-level workstation with an NVIDIA Quadro K600 performed up to 885% (9.9x) faster and the T1700 with the NVIDIA Quadro K2000 performed up to 1210% (13x) faster than the Dell OptiPlex 9020 standard desktop. That equates to an ROI of less than four work days.

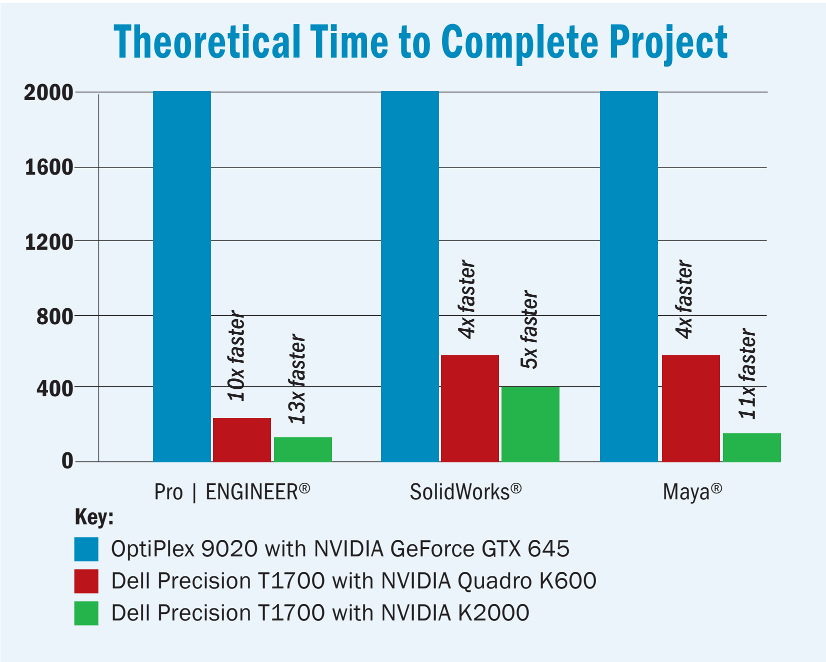

"If you're an engineer and your system is rendering, you're not an engineer," says Rhodes. An engineer who isn't doing any engineering is an expensive resource.

If design engineers' computers are not professional workstations, they're not hitting peak efficiency in simulation-led design workflows. A simple upgrade to a professional workstation with the latest design software is the right tool for the engineer's job. It will improve engineering and IT productivity, quickly pay for itself and allow your company to compete.

Download "Making the Case for Professional Engineering Workstations" at deskeng.com/de/proworkstations.

What Does It Mean To Be Faster?

As humans, we always strive to do things better, to push ourselves just a little bit further. The drive to move forward is even more pervasive in business. Bigger, better, faster is the rally cry for innovation, and engineers are tasked with doing all that they can to push product designs to the next level.

Simulation has become a key enabler for engineering groups to head down this path to more innovative design. The ability to test drive ideas using simulation tools in a virtual world allows teams to explore a greater number of design concepts and identify flaws earlier on the cycle when is far less costly to make changes.

Yet a successful simulation strategy hinges on having a hardware platform optimized to deliver results. If engineers find themselves limited by the size or the number of models they can simulate or lack the horsepower to test out designs in a realistic virtual setting, they are compromised by what today's simulation software can offer. What is the point of running a robust simulation problem if it ties up a workstation for hours or days on end, taking an engineer offline, unable to work. If simulation becomes the bottleneck for productivity, engineers will be less inclined to conduct extensive investigative studies simply because they can't afford to miss a project deadline or disrupt aggressive time-to-market cycles.

Intel® Xeon® Processor E5-2600v3 Product Family in a Snapshot

Here's how the Intel® Xeon® Processor E5-2600v3 Product Family stacks up in terms of features and functions:

2 Intel QuickPath Interconnect

8 Max Memory Channels (requires 2 processors)

Over 1TB Max Memory Footprint (requires 2 processors)

20 MB Max L3 Cache per processor

80 Max PC13 Gen 3 I/O Lanes

32 threads with Max Intel Hyper-Threading Technology Threads

Both Intel Integrated I/O (IIO) & Intel Data Direct I/O (IDIO)

Accelerate Innovation

But what if the workstation is no longer the drag on the design cycle, but rather the enabler, boosting simulation performance by a factor of nearly 3x and improving overall engineering productivity? The Intel® Xeon® Processor E5-2600v3 product family, a new line of workstation processors optimized for the highest performance and maximum memory support, does just that, empowering engineers to accelerate discovery and improve innovation throughput.

With support for two processors, up to eight memory channels, over 1TB of memory, and two Intel QuickPath Interconnect buses, workstations based on the Intel® Xeon® Processor E5-2600v3 product family move data faster for processor-intensive simulation, rendering, and ray-tracing tasks. The ability to move data faster means simulations run faster—in some cases, up to a 3.7x performance boost compared to comparable analyses running on a four-year old machine powered by an Intel® Xeon® Processor X5690.

The added horsepower of an Intel® Xeon® Processor E5-2600v3 workstation has other benefits over its four-year-old predecessor. Engineers can move beyond using simulations for single-point predictions to stochastic modeling, a practice of automating systematic, rapid changes to help determine system performance throughout a product's lifecycle. With this approach, which has traditionally required high performance computer (HPC) processing, an engineer rapidly iterates designs to determine how changes like curved corners or a reduction in material will aid in achieving key design goals, like minimizing weight, for example, or withstanding certain temperature ranges.

The Intel® Xeon® Processor E5-2600v3 product family ups the ante significantly, especially for simulation experts and power users. While they may have been able to get by with older workstation technology, it hasn't been without constraints on their design capabilities. Upgrading to a workstation optimized to speed up simulation practices removes any limitations and lets them kick design innovation up a notch. For more information on the Intel® Xeon® Processor E5-2600 product family, go to www.intel.com/workstation.

Don’t Skimp on ECC Memory

Professionals configuring a workstation optimized for robust simulation should not underestimate the importance of Error Correcting Code (ECC) RAM.

In fact, one in three computers will experience some sort of memory error. The impact may lie undetected—something like an error in unused memory or an odd pixel color. But the circumstances could also be far more catastrophic like an error resulting in irretrievable data and the subsequent hours of recreating lost work.

With data integrity key to a design professional’s livelihood, why take that risk?

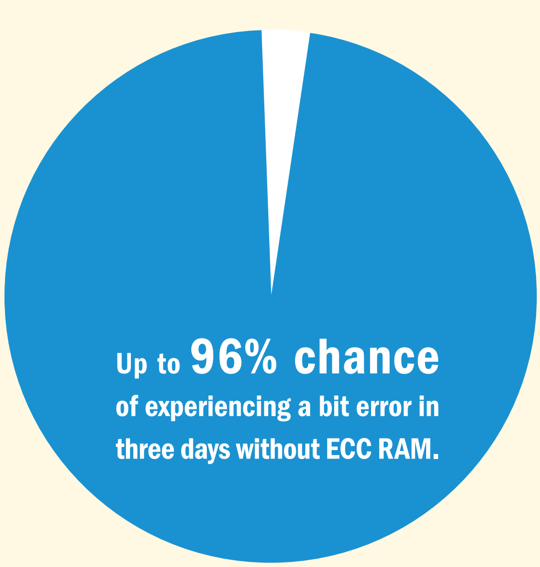

Source: http://lambda-diode.com/opinion/ecc-memory

Speeds up to 4x faster than your four-year-old workstation

Performance Summary: Xeon® E5-1600 v3 vs 4-Year-Old Machine

Four-Year-Old Workstation Refresh Yields Up To 6.2x Performance Boost

Performance Summary: Xeon® E5-2699 v3 vs X5690

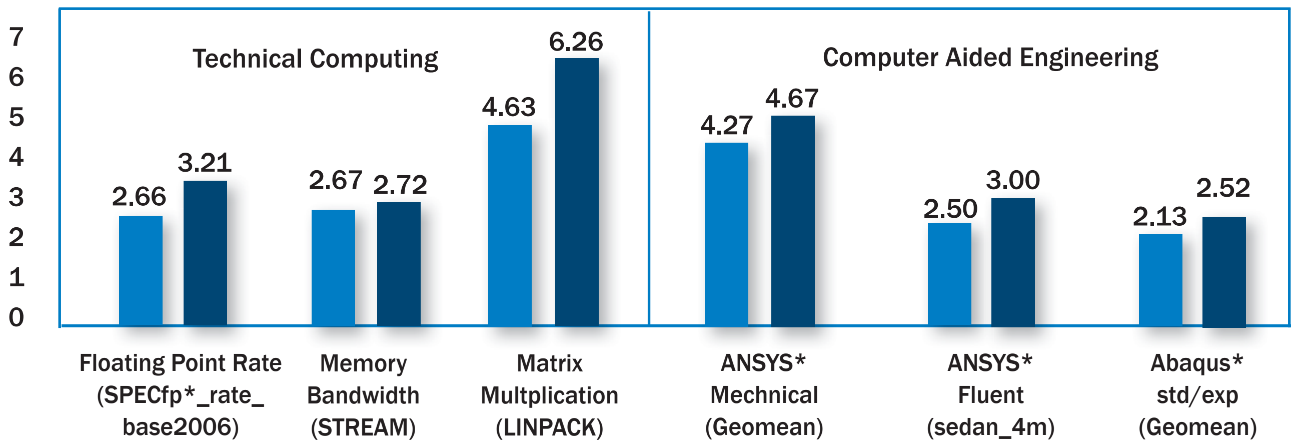

by Kenneth Wong & Beth Stackpole

Altair Rolls Out HyperWorks 13.0

Some software businesses might treat release 13 with a sense of anxiety; others might sidestep it altogether by calling it something else or skipping a release number. But not Altair. When it's time to roll out HyperWorks 13.0, Simone Bonino, Altair's marketing director, sees no reason to flinch. He points out, "As engineers, we appreciate the fact that 13 is a Wilson prime number, a Fibonacci number and the smallest emirp (prime spelled backward)."

The HyperWorks suite's emphasis is optimization, the company's core competency. "We currently have approximately 15 Optimization Centers all over the world, primarily in automotive and aerospace," Bonino says. "Because they operate as a vehicle to transfer optimization and lightweight design knowledge from us to our customers, they have a relatively short lifespan, usually no more than three to four years."

Updates to OptiStruct and HyperStudy

In Altair's approach to optimization, the company's OptiStruct and HyperStudy software are the centerpieces. Jeff Brennan, chief marketing officer of Altair, explains in the media-preview webinar that there are "significant advancements in our ongoing efforts to expand our linear analysis solutions into nonlinear analysis. In this release, the big highlight is the ability to do large displacement analysis in nonlinear analysis with hyper-elastic materials." Other new OptiStruct functions in this release include rotor dynamics and nonlinear heat transfer.

As a recognition of the growing need to combine solvers for multiphysics problems, HyperWorks 13.0 accommodates coupling of MotionSolve, Altair's multibody mechanical solver, and AcuSolve, Altair's general-purpose fluid flow solver.

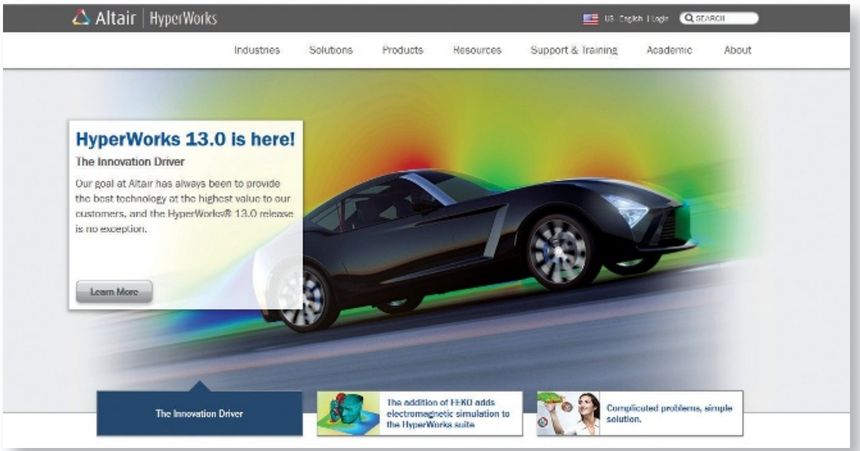

With optimization as its centerpiece, Altair rolled out HyperWorks 13.0 in early October.

The company has also added additional physics in its structural analysis engine RADIOSS to simulate crack propagation and airbag deployment, two phenomena commonly studied by automotive manufacturers.

Multithreading

"The real benefits of simulation come from running multiple simulations, potentially hundreds of thousands of them. That requires robust compute power," says Brennan. By refining its multithreading code, the company says it manages to make many of its solvers run faster. RADIOSS, for instance, is said to be 1.43 times faster than HyperWorks 12.0 on Intel Haswell processors.

"If you have multicore computers with multicore technology, you want to be able to use every processor in your computer to solve your problem. Now, we're looking at a technology where we can have crash models with 15 million elements, and we're paralleling up to 3,000 processors," says Uwe Schramm, Altair's CTO for solver technologies.

Multithreaded performance is critical to Altair's vision with HyperWorks Unlimited, a specialized appliance that could function as a private cloud or onsite high-performance computing (HPC) system. By leasing the hardware, users get unlimited access to bundled Altair software (hence, the name HyperWorks Unlimited). The appliance is aimed at those who want to pursue optimization studies involving hundreds, or even thousands of design iterations.

The Excel Handshake

In HyperWorks 13.0, the company plans to embrace Excel as an input device by allowing the use of Excel spreadsheets as entries for both OptiStruct and HyperStudy. The new release will also support PowerPoint export — not just as static images, but as linked data. Therefore, PowerPoint slides can remain synchronized to the post-processing results referenced in them.

Automotive designers and manufacturers may be interested in HyperWorks 13.0's multi-model optimization capability, which allows users to specify a common design component for a family of vehicles in an optimization study. The company is also bolstering its tools for designing with composite materials, a common approach in lightweighting projects.

Some say they're FAST...

But do they have the SCALE to deliver?

At Proto Labs, we're optimized to deliver quick-turn additive, machined and molded parts in as fast as one business day. Our three distinct services allow us to manufacture prototype and production parts for thousands of customers each day, many with dozens of designs that need to be simultaneously tested. Since 1999, we've shipped millions of parts to product designers and engineers around the world.

Our proprietary software and automated technology let us make parts for you faster than anyone else, but we back it up with large-scale manufacturing facilities on three continents that house hundreds of 3D printers, CNC machines and injection-molding presses.

Whether your project requires a single prototype or thousands of parts from 50 different designs, we have the capacity to meet your manufacturing needs every time.

2015 Cool Parts Calendar
Request your free calendar at
go.protolabs.com/DT4D

proto labs®
Real Parts. Really Fast.™

Visa/Mastercard Accepted | ISO 9001:2008 Certified | ITAR Registered

Both HyperMesh and HyperForm take the lead in this area by giving users the ability to study and understand composite fiber weaving and draping. The interactive post-processor in the new release lets you select and see the stress level, "even down to individual fiber level," says Brennan. "The new material laws added to RADIOSS and OptiStruct put (Altair) at the forefront of advanced materials."

Mesh Improvement

HyperWorks 13.0 features adaptive-wrap meshing, described by Brennan as "a complete rework of the previous shrink-wrap meshing." This will be particularly useful for creating watertight mesh models with recognizable features based on CAD models, he explains. The new release will use "rapid contouring with caching" to cut down on model reload and processing. The virtual wind tunnel simulator based on AcuSolve is also available in the new version.

In June, Altair completed the acquisition of EM Software & Systems, which developed the FEKO software for "antenna design, antenna placement, electromagnetic compatibility analysis, bio-electromagnetics, radio frequency components, 3D electromagnetic circuits, design and analysis of radomes, and radar cross-section analysis," according to a press announcement. The purchase is an indicator that Altair, already an established name in automotive and aerospace, is adding tools for developing Internet of Things (IoT) products into its portfolio in anticipation of growth in this area. The acquisition may bring Altair new opportunities in consumer goods and life science.

Release 13 may be even more meaningful to Schramm, the man behind the company's solver technologies, and a big soccer fan. As Bonino points out, Thomas Müller — one of the heroes of the 2014 World Cup — wears a No. 13 jersey.

— K. Wong

London Calling, Nissan Responds with New Taxi Design

London's iconic black cab, as synonymous with taxis as the yellow cab is for New York City, is about to hit the streets with a full design makeover. The move is in part to ensure it conforms to new regulations governing the Hackney Carriage status, and to serve as a launch pad for an electric vehicle (EV) version by 2015. London Mayor Boris Johnson has called for all city cabs to be emission-free by 2020.

Nissan's new taxi re-imagines London's traditional black cab. *Image courtesy of SAC Engineering.*

Nissan, which is using its multipurpose NV200 platform as the base for the taxi redesign, is planning to release a 1.6-liter petrol engine fleet with an automatic gearbox in December.

One of the key challenges of the project was making low-volume modifications to a high-volume vehicle — the NV200 van, in this case — to ensure it complies with the strict regulations governing Hackney Carriages, which includes a required 25-ft. turning radius, among other criteria.

In addition to the suspension modifications required to accommodate the new turning circle, the London Taxi specification calls for a side door modified for a ramp, and wheelchair access. Front bumper panels have been redone. The makeover includes round headlamps and a remodeled grille that mirrors the traditional "face" of the original London black cabs, while LED lights frame the taxi sign to increase visibility.

"It's a large transformation, but not a lot of panels are transformed," notes Steve Charlton, managing director at SAC Engineering Design, a high-end engineering consultancy that worked on the Nissan London Taxi design with ADV Manufacturing, one of the key design partners on the project. Specifically, the NV200's door, roof, front-end suspension, inner structure for rear seats, roof and front end surfaces were targeted as part of the redesign project, Charlton says.

Because of the nature of work related to the panel design, SAC relied heavily on Dassault Systèmes' CATIA — specifically, its sheet metal functionality. The additional sheet metal licenses enabled SAC to apply company standards to ensure design quality, and to address manufacturability constraints early on in the design phase, which accelerated assembly design.

The fact that CATIA is so widely used throughout the industry, and because SAC has been a user since 1993, positioned the CAD program as a natural fit for this project, Charlton says. "If we need external help with the suspension, for example, we can enlist partners and share models" without a problem, Charlton explains. "It's pretty much the industry standard for this sort of thing."

London's NV200 cab isn't Nissan's first foray into taxi cab design, however. The company has an extensive global taxi program, and has been involved with cab designs for New York, Barcelona and Tokyo.

— B. Stackpole

SolidWorks Supports Paperless Manufacturing

Is the paperless office a reality, or a Utopian concept? Whatever you might think, SolidWorks seems convinced there's sufficient interest in it to warrant a new product. At the press preview of SolidWorks 2015, the company shared details about a new product, dubbed SolidWorks MBD (Model-Based Definition).

In a SolidWorks blog, Jeremy Regnerus, the company's senior user advocate and community manager, writes: "MBD provides an integrated, drawing-less manufacturing solution for SolidWorks 2015. With these tools, you can define, organize, and publish 3D Product Manufacturing Information (PMI) and 3D model data in industry standard file formats … (It) defines 3D PMI such as dimensions, datums, geometric tolerances, surface finishes, welding symbols, bills of material (BOM), callouts, tables, notes, metaproperties, and other annotations within the SOLIDWORKS 3D environment. The process is both intuitive and interactive, and helps multiple people within the supply chain understand the design without the need for 2D drawings."

With MBD, SolidWorks envisions manufacturers publishing annotated 3D views — one set for shop floor, another for quality inspection, another for a specific vendor, and so on — to be distributed in the eDrawings or 3D PDF format. But it's much more than a simple publishing tool, according to SolidWorks global product manager Oboe Wu.

"It's an integrated drawing solution to support the entire MBD workflow," Wu says. "It's not just 3D PMI, but also PMI organization, template customization and 3D data sharing."

Anark, a SolidWorks partner, also hopes to capture the same market with its product Anark Core MBEWorks. Anark's MBE stands for Model-based Enterprise. According to Anark, MBEWorks "converts SolidWorks 2012 parts and assemblies with 3D dimensions, tolerances, GD&T, and other 3D notes into high-resolution 3D PDF and HTML documents containing manufacture-quality BREP geometry, interactive associativity between features (ASME Y14.41 'visual response' characteristics), and 'machine readable' semantic definitions for each component."

Chris Garcia, Anark's executive VP of global commercial operations and business development, previously worked as VP of R&D at SolidWorks.

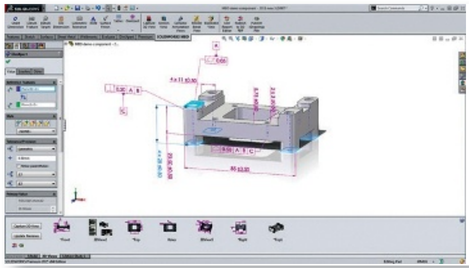

SolidWorks hopes to support the paperless workflow with a new product, SolidWorks MBD. *Image courtesy of SolidWorks.*

Resistance Remains

MBE and MBD seem to be a natural outcome of the manufacturers' shift to digital prototyping, supported by extensive use of 3D design software. The move to replace paper drawings with lightweight digital 3D drawings is gaining momentum among top-tier aerospace and automotive manufacturers, but at lower levels, some resistance still remains, as seen in a two-year-old discussion thread on a SolidWorks forum (forum.solidworks.com/thread/56419).

SolidWorks product manager and blogger Matthew Lorono poses the question: "Do you intend on implementing Model-Based Definition? (Get rid of drawings?)"

User Glenn S. says, "I work at a university-based facility that designs and tests roadside safety and anti-ram structures … Most (probably all) of our fabricators want drawings. Our construction crew will always need them (if we issue them iPads, they would need a new one every two weeks), and [drawings] are also required for the reports that we issue documenting the structure that was tested."

For some, the resistance is not based on practical concerns or technological hurdles, but instead in culture and habit. Kenneth B. notes that "working from paper is more natural. I prefer to read printed books as opposed to digital media … [Digital] annotation/notes will always be a [pain] to have to constantly go to an electronic device to refer back to."

A favorable comment comes from Mark K., who says, "I'd just love to see computer screens at every workstation so they could view eDrawings, PDFs … instead of having paper drawings everywhere. That would be a major step for us." But even he has to admit, "I can't see how you would put together an assembly without some written documentation, drawing, or otherwise." MBD/MBE advocates might point out that such assembly instructions can also be produced as dynamic, animated 3D drawings.

Filling a Need

MBD is a plug-in that requires a license of SolidWorks. Priced at $1,995, with a $495 annual subscription, MBD costs as much as some full-featured CAD programs. Citing customers he has spoken to, SolidWorks' Wu reports that "one-third of engineering dollars is spent on creating and maintaining 2D drawings." That convinces him, he adds, that the headache justifies spending the procurement cost of MBD software.

As more manufacturers and field crews switch to mobile devices, MBD will become easier to implement. Smaller firms, though, may need more convincing, encouragement and economic incentives to consider MBD.

— K. Wong

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.

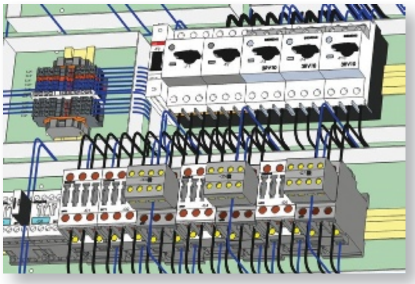

Zuken Introduces E^3.series 2014A

Updated software adds features for enhanced manufacturing.

The E^3.series is a modular set of integrated electrical CAD and documentation solutions. It lets you use multi-view project files where all project design data such as schematics, cable plans, control panels, form boards and fluid detail are kept and dynamically linked. It has built in real-time design rule checks, intelligent libraries of symbols and components and file viewing and other collaboration tools.

The E^3.series 2014A release focuses on enhanced manufacturing with new features such as a new PLC data exchange and the eCl@ss Advanced tool.

MORE ➜ deskeng.com/de/?p=18790

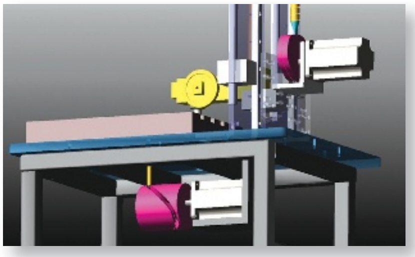

MSC Software Launches Adams 2014 Application

New version supports co-simulation with MSC Marc.

Adams 2014 introduces native nonlinear part modeling and analysis functionality in the form of a new Adams-native modeling object for very large deformation use cases called FE Part. It is meant to provide engineers with a way to more accurately calculate dynamic loads for geometrically nonlinear parts within a multibody dynamic system. FE Part offers a basic 2D option and a 3D Beam formulation option. The 3D Beam formulation is a fully geometrically nonlinear representation that can be used for beam-like structures. Complementing this is a new FE Load feature.

MORE ➜ deskeng.com/de/?p=19102

Roland DGA Releases monoFab series

The desktop fabrication product line includes a 3D printer and milling machine.

Roland DGA has released what it calls the monoFab series of desktop fabrication systems. The series is comprised of two machines: the monoFab SRM-20 subtractive milling machine and the monoFab ARM-10, the company's first ever 3D printer. The idea behind the series is to offer designers and engineers a set of desktop tools that lets them integrate the additive manufacturing capabilities of 3D printing with the precision subtractive fabrication capabilities of milling.

The ARM-10 has a good-sized build area that can accommodate making multiple objects. Build area specs are 5.11 (130 mm) W x 2.75 (70 mm) D x 2.75 in. (70 mm) H.

MORE ➜ deskeng.com/de/?p=19160

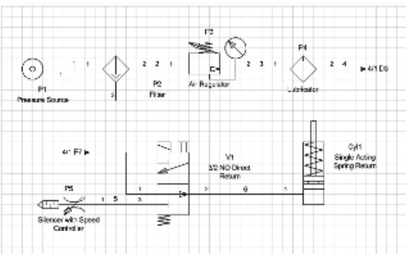

Radica Software Updates Electrical CAD Platform

The program is supported by Microsoft's Visio technical diagramming tool.

Radica Software recently released the E7 version of its Electra electrical CAD platform. It has a bunch of sweet-sounding attributes to it, and chief among them is that it runs on top of Microsoft's Visio technical diagramming and layout tool.

Electra E7 uses a lot of automation to simplify creating and documenting circuit drawings. This includes features like automatically tagging symbols and wires, assigning the right component, generating terminals and building a table of contents. Electra can also find all connection points and float them to a group level.

MORE ➜ deskeng.com/de/?p=19253

2BOT Develops Compact Subtractive Modeling System

A company called 2BOT says it has developed a subtractive process that can be used in a classroom or office setting. 2BOT calls its new system the ModelMaker. It acts like a desktop mill using 3D data and proprietary material to produce models built in SketchUp, AutoCAD, and similar modeling programs.

As far as technical specs go, the ModelMaker is a desktop-sized system, with a 25 x 25 x 13 in. footprint. 2BOT claims the system is accurate to 0.030 in., and has a cutting speed of around 1 in./sec. The standard material cassette that fits into the ModelMaker is 12 x 12 x 2 in., but 3D designs can be sliced into multiple production runs to produce larger models.

MORE ➜ rapidreadytech.com/?p=7606

Arduino Enters the 3D Printing Space

One of the latest companies to enter the 3D printing market is Arduino, which is mainly known for its open-source, single-board microcontroller. Arduino's new 3D printer is named the Materia 101, and officially launched at Maker Faire Rome. Arduino received help in developing the new AM system from Sharebot.

The Materia 101 has been designed to be about as affordable as any desktop

system can hope to be, with prices of €600 ($800) for a kit or €700 ($1,000) for a fully assembled system.

MORE ➜ rapidreadytech.com/?p=7623

PrintAlive Bioprinter Offers Hope for Burn Victims

While any advance in additive manufacturing (AM) technology is good news, advances in an ever-growing medical portion of the market can be cause for celebration. Medical AM, including accurate surgical models, prosthetics, and other medical devices offer new treatment methods and help reduce costs. That's not even including bioprinting, which is possibly the most exciting area of development for the medical field.

A group of students at the University of Toronto have advanced the technology of skin replacement surgery with the PrintAlive Bioprinter. The device was named the Canadian winner of the 2014 James Dyson Award, which qualifies the research team for a shot at the grand prize total of $60,000 in funding.

The PrintAlive Bioprinter could eliminate the skin graft step by printing new skin directly onto the damaged area using the patient's own cells to reduce the chances of rejection.

MORE ➜ rapidreadytech.com/?p=7612

Rapid Ready Review: 'Print the Legend'

The year 2012 is generally seen as the year when AM went mainstream. Instead of being hidden away, it started making the news as the next great manufacturing revolution. "Print the Legend" is a documentary about AM that managed to coincide with that zeitgeist and capture the growing excitement and maturation of the 3D printing market.

Broadly speaking, the documentary has three main characters. These are Bre Pettis, co-founder of MakerBot; Max Lobovsky, co-founder of Formlabs; and Cody Wilson of Defense Distributed fame. The film follows each individual from humble beginnings to lofty heights, and manages to tell some of the story of AM along the way.

MORE ➜ rapidreadytech.com/?p=7596

Dremel 3D Printer Uses Autodesk Spark

Dremel's new AM system runs on proprietary 1.75mm filament, which is available in 10 colors. The 3D Idea Builder has a build envelope of 9 x 5.9 x 5.5 in. (230 x 150 x 140mm), and a variable layer thickness of up to .1mm. No word on build speed, but keep in mind this is meant to be a home AM system.

Dremel has also set up a deal with Autodesk to help provide digital objects that are ready to print, including objects that can be altered to add personal touches. The strategic partnership will also offer a free and simple design toolset for users looking to make their own 3D designs. Additionally, Autodesk will be tasked with offering support to new users to help them through the design process.

MORE ➜ rapidreadytech.com/?p=7577

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

LiDAR Puck Showcased on TALON Robot

At October's RoboBusiness conference in Boston, Velodyne LiDAR and QinetiQ teamed up to demonstrate the VLP-16 LiDAR puck.

The VLP-16 was on the TALON robot, which is known for its resilience in military, law enforcement and first responder applications. Initially launched in 2000, it is able to climb stairs, negotiate rock piles, overcome concertina wire and plow through snow.

The VLP-16's specifications include a range of more than 100 meters, with a target range of 150-200 meters. The unit has a power consumption of less than 10W and weighs about 600. It provides 16 channels, approximately 300,000 points/sec and a 360 degree horizontal field of view.

MORE ➜ deskeng.com/de/?p=20051

IoT Wzzard Sensing Platform Launched

B&B Electronics, a developer of network connectivity products, has introduced the Wzzard Intelligent Sensing Platform. The platform is intended to be a medium to develop scalable, reliable, Internet of Things (IoT) networking in remote and demanding environments.

The Wzzard platform uses iterative control limits and gateway data aggregation to create a network. In order to generate connectivity, Wzzard Intelligent Edge Nodes connect via a conduit fitting cable to any industry-standard sensor. The nodes are already configured with internal temperature and accelerometer options, and can be controlled using Android or iOS smartphone devices.

MORE ➜ deskeng.com/de/?p=18171

Stanford University Engineers Create Ant-Sized Radio

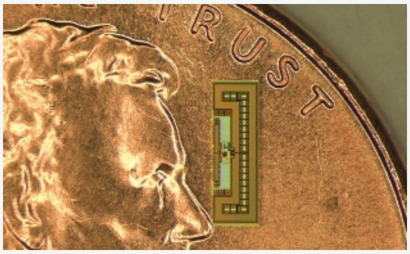

Stanford University, in collaboration with University of California, Berkley, has created a radio the size of an ant. Designed to execute and relay commands, the goal is to have this device be easy to manufacture and be placed in Internet of Things (IoT) devices.

The radio is so energy efficient, a university press release states, that it gathers all power necessary from the same electromagnetic waves that carry signals to its receiving antenna. This means the device does not require batteries to operate.

"The next exponential growth in connectivity will be connecting objects together and giving us remote control through the web," said Amin Arbabian, an assistant professor of electrical engineering at Stanford. Looking forward, Arbabian hopes to create networks of these chips deployed every meter or so throughout the house. He believes this technology can provide the web of connectivity and control between the Internet and smart household devices.

MORE ➜ deskeng.com/de/?p=19305

Brain Study Could Advance A.I.

University of Reading researchers have gathered a new understanding of how the brain processes information and say their findings could advance artificial intelligence.

The study had participants play a 3D game where they needed to adjust the slant of a surface so a moving ball bounced off it and through a target hoop. Throughout the study, the bounce of the ball was altered so the surface behaved differently.

Researchers found that the brain uses internal simulation to change the slant's angle for consistent scoring. It also proved that humans are an adaptable species — a trait engineers are now studying and looking to utilize in effective autonomous robot development. The industry is further studying how humans' sensory systems can achieve what is not feasible for robotic systems.

MORE ➜ deskeng.com/de/?p=18444

Mission to Launch Space-Based Supercomputing

ConnectX, a collaboration of experts in the aeronautics, mathematics, physics, engineering and business fields, aims to replace current cloud computing with a Space-based supercomputing platform. With this new structure, the company will send, receive, store and process massive amounts of Big Data at a much lower costs, it stated in a press release.

The next-generation Big Data computing and communication platform would reportedly use powerful, miniature satellites that leverage the space environment. The system would also replace binary code with a new symbol structure for optimized transmission, storage, processing and delivery of data.

MORE ➜ deskeng.com/de/?p=19933

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

Designing the Internet of Things

Individuality, cloud connectivity and aesthetic appeal are the cornerstones of smart products.

BY KENNETH WONG

For Ken Macken and his colleagues at ollo wearables, the phrase "Keep me in the loop" is more than a figure of speech. The personal telecommunication device they're developing is codenamed Loop. At a recent press event hosted by the California-headquartered semiconductor firm Broadcom, ollo wearables CEO Hugh Geiger appeared with an early model of Loop around his neck. With smooth surfaces and tapered ends, it looked a bit like a crescent-shaped boomerang with a mouthpiece.

"We're trying to reimagine the smartphone experience," Geiger said at the event. "We want to explore new ways you might interact with your phone."

With the microphone positioned below his lips, Geiger issued dialing commands without ever using his hands to access a menu or push a button.

Macken, who had worked with vision-guided robotics and minting machines in his previous jobs, describes his background as "old-school mechanical engineering." Soon after launching ollo wearables, Macken quickly discovered the importance of aesthetic appeal.

"If we make something that works well in its engineering, but is not aesthetically pleasing, nobody will buy it," he observes. "People prefer to buy a device that looks great, but doesn't have a lot of features over the one with all the features they want, but looks horrible. And they're willing to pay more for the former. We found that out from our own research."

In traditional engineering, function take precedence, he notes. "There, you make sure all your mechanical and electrical components work well, then you wrap a nice shell around it," he says. "But here, we design something beautiful first, something the user would want to keep on himself or herself all the times, and then we go about fitting the components in it. We transitioned from an engineering company to a design company."

That mindset shift is critical in developing highly personalized Internet of Things (IoT) devices, Macken says.

ollo wearables' Loop telecommunication device, designed in Autodesk Inventor and rendered in Luxion KeyShot.

Shortcut to Compliance

Macken and his colleagues have developed as many as six different variations of Loop. Macken is also the principal CAD modeler on the Loop project. He uses a combination of Autodesk Inventor and SketchUp to sculpt the design in 3D. "I import hand-drawn sketches into Inventor and then create digital sketches from them," he says. The same CAD model serves as the basis for photorealistic renderings done in Luxion Keyshot, used for presentation and marketing.

Aside from the on-off button, there are hardly any other mechanical components in Loop. Most functions — locating contact info, dialing, setting up alerts, etc. — will be driven by voice commands. In fact, even the on-off function can be voice-activated, Macken reveals. But the engineers chose to include a hard button because they discovered that, when it comes to switching the device on or off, consumers are not yet ready to give up the tactile sensation.

One ongoing, critical decision for ollo wearables is selecting a chipset that can power these functions through a combination of embedded electronics and software. "We can probably get a gray-label chipset from overseas and make it work," says Macken. "But we're looking for support, high reliability and compliance."

Using a chipset that's known to comply with industry standards and local regulations, he points out, goes a long way to remove potential roadblocks in getting their own device approved for sales. Macken and his colleagues are currently testing out Loop with chipsets from Qualcomm, Texas Instruments and Broadcom.

Personal User Experience

In the test versions of Loop, the device runs on custom firmware, but ollo wearables plans to migrate to Android OS. "Loop is more than a device," says Macken. "It's a software and services platform."

The emphasis on software makes the device easy to update without any mechanical intervention. Perhaps more important, the software and services facilitate the personalization that's a common IoT characteristic.

Loop won't be as sassy or artificial intelligence (AI)-like as the fictional OS Sam in the futuristic romantic comedy "Her," but Macken says, "We're trying to anticipate what a user wants to do and execute it before being asked to do it. For example, Loop might read out the weather every morning after the alarm if that's what you typically do. And if you usually call your mom every Sunday around 11, on the Sunday you don't call, Loop can issue an alert: Did you forget to make that call?"

The Cloud is Essential

Many of the predictive functions in Loop are driven by historic user behavior, like the type of calls regularly made on a certain day or the type of data requested from the device. Because it's not feasible to always store and maintain that ever-growing data in a device the size of a USB thumb drive, streaming data to and from the cloud is the best option.

"Our device is connected to the cloud, controlled by the cloud," Macken says. "There are trigger events that occur in the cloud. They activate other triggers. [The cloud] is a huge part of our efforts."

As a small startup, ollo wearables doesn't maintain its own data centers. Instead, it partners with Rackspace, a managed cloud solution provider. The Rackspace database for ollo wearables is duplicated in several geographic locations, so it can quickly stream to users located in different regions.

New Design Considerations

Broadcom recently released its software developer kit (SDK) for Wi-Fi- and Bluetooth-enabled devices. The SDK is dubbed Wireless Internet Connectivity for Embedded Devices, abbreviated as WICED ($19.95, pronounced "Wicked," like the musical). It's a system on a chip (SoC) with its own software. The company expects engineers to use WICED to prototype connectivity functions in cameras, washers, dryers, refrigerators and other IoT products.

"Your simulation will now have to include the 'smart' part of the 'smart products,'" says Brian Bedrosian, Broadcom's senior director of embedded wireless. "Because connectivity is a big part of this type of devices, the CAD designer would have to think about antenna placement, RF (radio frequency) shielding, sensors and others."

Shelley Gretlein, director of platform software at National Instruments (NI), agrees, noting, "The digital design can be treated independently of the sensor functions; however, there are considerations that must be taken into account as these are combined. The functionality of the sensor can be confirmed, but the placement of the sensor relative to heat sources, power outlets and accessibility points is critical for a connected device."

In developing a wearable heart-rate monitor, for example, the designer would need to consider not just allocating sufficient space for the electronics and chipsets, but how, during normal usage, the user's physical activities might interfere with the sensor monitors and RF transmissions.

"Every device is not going to be connected to the cloud all the time. Some of the decision making and processing have to be done locally," Bedrosian points out. "So the [engineers] have to decide when and how the device will connect to the cloud."

Gretlein says, "In Industrial Internet of Things (IIoT), a more connected system impacts all members of the design team. In fact, systems that only account for IoT in one part of the design are often inferior products, or don't integrate as fully as they could or should, depending on the specific device. With that said, there are certainly more demands on the software side of IoT design than in the past."

The IoT Pie

Broadcom, Qualcomm, Texas Instruments and Intel are among the growing number of players jostling for space in IoT. Those already established in the embedded systems market propose their programmable chipsets ought to be the de facto IoT platform. Graphics processing unit (GPU) giant NVIDIA dips its toe into the same waters with its NVIDIA Jetson TK1, a base board with Tegra K1 SoC and quad-core ARM CPUs.

On the other hand, when Oracle refers to its "IoT platform," the company is not talking about a board but its Java-based software architecture to facilitate machine-to-machine communication. Though historically a product lifecycle management (PLM) software company, PTC has in the past few years made a number of strategic acquisitions to bolster its IoT arsenal. Along with application lifecycle management and service lifecycle management software, PTC now owns ThingWorx, a software environment for developing applications for connected devices. The birth of the IoT Consortium in early 2013, led by Logitech and other firms, attests to industry momentum and enthusiasm behind IoT commerce.

In its report titled "Ten IT-enabled business trends for the decade ahead" (May 2013), McKinsey Global Institute notes, "The proliferation of devices and applications is so rapid that this trend, which in 2010 we called the Internet of Things (IoT), we have expanded to call the Internet of All Things." In another report titled "Disruptive Technologies: Advances that will transform life, business, and the global economy" (May 2013), McKinsey quantifies IoT's impact in dollars: Its estimated potential economic impact is $2.7 trillion to $6.2 trillion annually.

No matter how you slice it, that's an irresistible pie.

— KW

GE's Brain Surgery

If you're an appliance maker that manufacturers and trades in large volumes based on established supply chains, how would you respond to the IoT consumer's preference and insistence on a personalized user experience? If you're GE, you'd plant the seed for a community of tinkerers, hackers and programmers who could develop apps and gadgets to augment your ovens, dishwashers and laundry machines.

GE recently released the Green Bean Maker Module ($19.95), a circuit board that the company calls "a direct path into the brain of your GE appliances." Plugging the Green Bean into the appropriate port in a GE appliance, anyone with programming know-how may interact directly with the microprocessor inside the unit.

Taylor Dawson, a GE appliance engineer and GE product enthusiast, characterizes Green Bean as a device "that appeals to tinkerers and makers capable of developing hardware or software for their dishwashers or laundry machines. Our vision is to create a community of people who can share and refine these applications."

The Idea Garden

Even in its nascent stage, the budding Green Bean user community, dubbed FirstBuild, has already produced projects and ideas aplenty. From Joan from Louisville, KY, came the idea for an app that generates recipes based on the ingredients available in a refrigerator. Nathan from San Francisco thinks it's a good idea to equip the trash can with a barcode scanner. That way, it can "add items to your shopping list whenever you throw away an empty container," he writes.

FirstBuild also issues challenges to invite new ideas. Winners stand to earn as much as $2,500 for their ideas, along the chance to earn royalties for subsequent GE products that incorporate the idea. At press time, the challenge urges food enthusiasts to come up with a way to bring indoors what has always been an outdoor activity: grilling. James from Auburn, AK, thinks you can stop the smoke right at the source by incorporating a water container beneath the burner. But Darren, another community member, thinks built-in venting systems and fans are the way to go.

"We want the community members to take pride in what they have done," Dawson says. "We want them to say, I had an impact on a major GE appliance that's now in use in millions of homes."

GE's FirstBuild initiative suggests the company considers the talent of the crowd to be a tremendous source

GE's Green Bean module, described as "a direct path into the brain of your GE appliances."

of IoT innovation. Shelley Gretlein, director of platform software at National Instruments (NI), similarly points to Instructables, a maker community (acquired by Autodesk in 2011), as a fertile ground for homegrown IoT projects.

"More connections and more innovation mean more opportunity for device and system simulation, as well as an increase in the need for more CAD designs," Gretlein says. "This growth, combined with the increased capabilities of 3D printers, creates a productive system among device design in CAD, system simulation and almost-instantaneous device creation/printing."

Self-Aware Components

The integration of sensors and electronic components adds complexity to product testing and simulation.

"Previously 'dumb' nodes in a system now need to be tested for their communication and interactions with the rest of the system, adding a test burden that some designers may not account for in the initial prototype," Gretlein says. "Sensor integration in the IoT world will look much different than the past. Sensors will know about themselves, the product, and the operations. This built-in communication allows for in-field serviceability and upgrades."

For traditional mechanical engineers, IoT components require much more than physical space allocation inside the device's housing. It demands a different way of thinking.

"In a much more extreme IoT design, the physical design could be fused with the information design," Gretlein explains. "Rather than manufacturing being distinct from IT, these functions could be indistinguishable. In this situation, the material would 'know' what it is and what it's used for, and carry that information with it as it's formed into the final device. This interdisciplinary collaboration is a change in process design — and fundamental engineering thinking." **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 → **Autodesk:** Autodesk.com

→ **Broadcom:** Broadcom.com

→ **GE Green Bean:** Firstbuild.com/mylescaley/Green Bean-maker-module/

→ **Intel:** Intel.com

→ **IoT Consortium:** IOfThings.org

→ **Luxion:** Keyshot.com

→ **McKinsey Global Institute:** McKinsey.com

→ **National Instruments:** NI.com

→ **NVIDIA:** NVIDIA.com

→ **Logitech:** Logitech.com

→ **ollo wearables:** ollowearables.com

→ **Oracle Corp.:** Oracle.com

→ **PTC:** PTC.com

→ **Rackspace:** Rackspace.com

→ **SketchUp:** SketchUp.com

→ **Qualcomm:** Qualcomm.com

→ **Texas Instruments:** TI.com

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

Get Started with an HPC Cluster

As HPC cluster technology becomes more accessible, companies of all sizes have the potential to tap unlimited horsepower to create new simulation-driven engineering workflows. But before diving headfirst into what's likely uncharted territory, there are a number of considerations small- and medium-sized enterprises (SMEs) need to hammer out to ensure they make full use of this key technology.

By some industry accounts, SMEs are in the driver's seat for about 80% of product design in their role as suppliers and design partners to larger corporations. Yet this group faces challenges to scaling its existing workstation infrastructure with HPC resources to accommodate simulation-driven design workflows.

Simulation specialists aren't necessarily HPC experts, and smaller shops may lack the expertise and IT resources to effectively deploy, configure, and operate HPC clusters without a lot of handholding. HPC clusters, while far more affordable than traditional big-iron HPC systems, are still a significant investment, especially for smaller shops.

Consider Existing and Future Workloads

HPC clusters offer a great deal of choice with regard to processors, memory configurations, high speed networks, and cluster software. You don't want to risk under powering the new environment or conversely, over paying for expensive HPC compute power that is left to sit idle. Applications and user experience should drive cluster specifications.

The HP-Intel Innovation Initiative sponsors a number of programs to simplify the configuration and deployment of HPC clusters and to steer companies to the optimal HPC environment.

The program supports development of application-specific Solution Reference Architectures (SRAs). The SRAs provide advice on optimal cluster designs, based on performance testing at HP's in-house Benchmark Centers and inputs from software developers. These SRAs are available for leading CAE applications from vendors such as ANSYS® and Dassault Systèmes SIMULIA®. (See solutions at www.hp.com/go/hpc.)

HP Cluster Platforms are flexible, factory integrated solutions available from HP and its resellers. These clusters can be built to support the SRA recommendations, and can be expanded easily over time to support changing workloads.

Minimize Cost and Maximize ROI

HPC clusters are core capabilities for manufacturers, delivering ROI with benefits such as lower time to market, and allowing deeper insights into R&D problems. HPC clusters are not an insignificant investment, but today's technology is providing unprecedented performance for entry HPC clusters.

Turnkey entry clusters provide an affordable path to HPC and are available with efficient, compact designs and on-site support for easy integration into the SME's infrastructure. For example, ANSYS-specific HPC Starter kits are available from HP reseller, Dasher Technologies. They feature a turnkey, optimized hardware/software platform, pre-configured with ANSYS CAE software, delivered and installed on site, with remote management and support (www.dasher.com/solutions/ANSYS). Monthly leasing can minimize budget impacts.

The HP ProLiant SL2500 Gen8 Server provides a compact, yet powerful, HPC cluster solution perfect for an SME customer's first HPC cluster. ***Image courtesy of HP and ANSYS.***

Another option is HP Helion (hp.com/helion/hpc), a self-service HPC solution that provides a portal to a private cloud. For enterprises considering deploying a private cloud, HP Helion enables inclusion of an HPC service into that private cloud.

Manageability and Simplicity

HPC clusters use management tools that may be unfamiliar to IT administrators. Consulting engagements and managed services can help fill the gap, but it's key to have robust cluster management and workload scheduler capabilities via an intuitive user portal, so non-HPC experts can self-service common tasks.

HP® and Intel® have also teamed up with the National Center for Manufacturing Sciences (NCMS) and prominent universities like George Mason to create HPC Innovation Hubs, which provide on-demand access to HPC resources to SMEs that aren't in a position to buy and operate their own HPC clusters.

With so many affordable and accessible resources to choose from, SMEs no longer have to hold back from making HPC clusters a key part of their simulation design efforts. To find out more about the HP-Intel HPC Innovation Initiative, go to www.hp.com/go/compete.

Testing Simulation in the Cloud

A research team does an engineering simulation experiment in the cloud for a severe service control valve.

BY MARK A. LOBO, P.E.; JON DEN HARTOG, P.E.; DERREK COOPER AND WOLFGANG GENTZSCH

Engineers and scientists can use various types of simulation during design and development to reduce costs, reduce project risk, accelerate time to market, and deliver higher-quality products. Using specialized computing hardware for this simulation work can dramatically improve productivity by removing bottlenecks posed by some workstations. According to a study of the U.S. Council of Competitiveness, however, less than 10% of manufacturers are using technical servers for simulations. The vast majority of companies perform virtual prototyping just on their workstations, but 57% report they have application problems that they can't solve because their desktops are too slow or because the geometry or physics are too complex.

There is another option today to get additional technical computing power: the cloud. Computing resources in the cloud allow engineers to continue using their workstations for daily design and development, as well as common office tasks, and to submit larger, more complex, more time-consuming jobs to be solved elsewhere. Benefits of cloud computing are on-demand access to "infinite" resources, pay per use, reduced capital expenditure, higher-quality results, lower risks and product failure rate, and greater business agility by dynamically scaling computing resources up and down as needed.

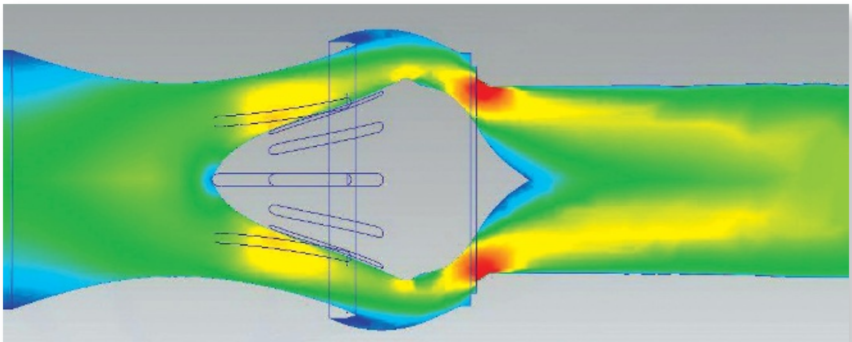

Figure 1: Idealized flow path for choked flow control valve.

The UberCloud Experiment started in July 2012, exploring hands-on cloud computing for engineers and scientists and to learn about cloud hurdles and how to remove them. Since then the UberCloud Experiment has attracted 2,000 organizations from 72 countries, and 155 teams in computational fluid dynamics (CFD), finite element analysis (FEA), biology and other domains have been built, and many case studies published in the 2013 Compendium and the 2014 Compendium of UberCloud Experiment case studies, available here: deskeng.com/de/simulation-cloud. The UberCloud Experiment provides a platform for scientists and engineers to explore, learn and understand the end-to-end process of accessing and using cloud resources, and to identify and resolve the roadblocks. End-users, software providers, resource providers and computing experts collaborate in teams to jointly solve the end-user's application in the cloud.

Taking Cloud on the Cloud from Concept to Reality

Flow control valve specifications include performance ratings for a valve to be properly applied in fluid management systems. Control systems sort out input parameters, disturbances and specifications of each piping system component to react and produce a desired output. System response is chiefly a function of the accuracy of control valves that respond to signals from the control system. Valve performance ratings provide information to the system designer that can be used to optimize control system response and predictable restriction vs. the input

TABLE 1: COMPARISON OF DESKTOP, CLOUD AND HPC SOLVING OPTIONS

SIMULATION SOLVING APPROACH	APPROXIMATE TIME TO COMPLETE	INVESTMENT REQUIRED
Single Desktop Machine + Local Computing	800 hours (1 month)	Engineering Workstation + Simulation Software License
Single Desktop Machine + Cloud Computing	24 hours (1 day)	Engineering Workstation + Simulation Software License + $1,200 Cloud Compute Fee
Single Desktop Machine + Private HPC Cluster + Multiple Solver Licenses	24 hours (1 day)	Engineering Workstation + Simulation Software License + 30 Node Compute Cluster + 30 Simulation Solver Licenses

Source: The UberCloud Experiment

signal range, the inherent characteristic, which requires hundreds of simulations to verify and optimize.

One experiment by Mark A. Lobo, P.E., Principal, Lobo Engineering PLC focused on a control valve where the effects of a complex body cavity and trim design were critical to overall performance. The flow path for the valve is shown in Fig. 1. The restriction components, or "trim," reduce the annular area as the cavity profile on the right moves to the left. The location of highest velocity is indicated in red. The design objective was to direct the fluid away from the containment surfaces that otherwise would be damaged by cavitation.

In Lobo Engineering's UberCloud experiment, the cloud was used to complete several sweeps of simulations to characterize the performance of this valve in various operating conditions. Rather than solving these simulations sequentially, as would be done with a typical license and hardware setup, most of the simulations were solved in parallel in the cloud to maximize throughput and allow studies of various conditions and geometries to be studied relatively quickly.

The Virtual Valve Test

Autodesk Simulation CFD, the simulation tool used in the project, employs solution methods based on well-understood fluid dynamics principles, and provides the option to solve locally or in the cloud. The iterative approach to solving the embedded fluid flow equations is presented graphically to provide an appraisal of the program's progress converging toward a solution. Adjustments can be made to settings

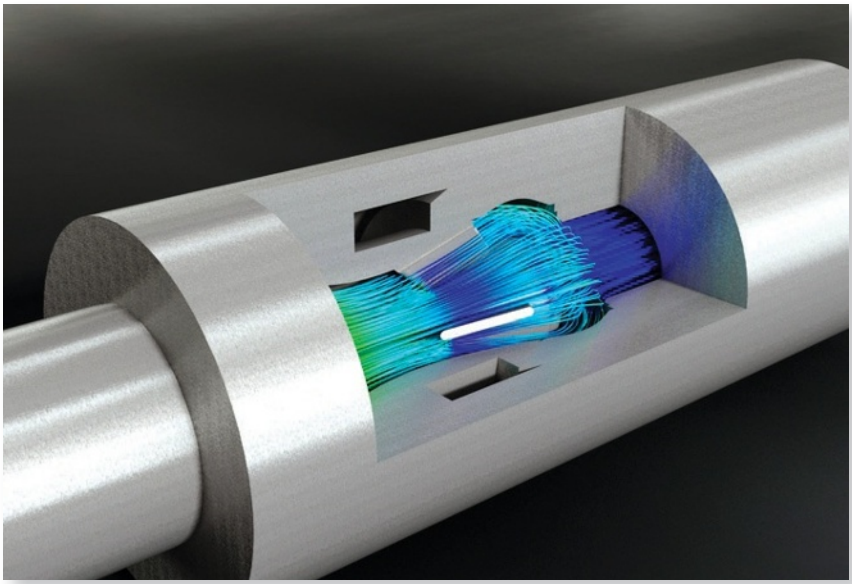

Sending solutions to the cloud could reduce the time to collect a complete data set by an order of magnitude.

Meet the Control Valve Experiment Team

TO discover the potential of cloud computing for compute-intensive engineering tasks, such as simulation, the UberCloud Experiment pairs end users with application experts to help them explore the use of cloud computing. The control valve experiment team consisted of:

- End user Mark A. Lobo, P.E., Principal, Lobo Engineering PLC. Lobo has broad experience in industrial product engineering, specializing in application of modeling and analysis software for control valve design.
- HPC/CAE application expert Autodesk, which supported this project by providing access to simulation software and cloud computing resources. Specific application support was provided jointly by Jon den Hartog, P.E., senior product line manager, Autodesk, and Heath Houghton, simulation CFD product manager, Autodesk.

such as mesh controls, solution methods, advection model and more. After acceptable results are realized locally with confidence in the program settings, multiple scenarios can be sent to the cloud to free up local resources and reduce analysis time.

The valve test simulation boundary conditions are set in the same manner as a physical test setup. The upstream and downstream pressures are applied, and Simulation CFD solves for the flow rate. During the course of this project, more than 200 simulation scenarios were solved in the cloud. The cloud solves were performed on a pool of Amazon EC2 m2.4x large instances, each with eight physical cores and 64.8 GB RAM. The local solves were performed on a HP Z800 Dual 2.4GHz Quad Core E5620 Workstation, 24GB RAM, with AMD's ATI FirePro V5800 1GB Graphics Adapter. Flow data was collected for six control positions of a representative control valve, and 10 combinations of pressure drop for each position, or travel step, from 20% open to 100% open. The data was analyzed in the same manner as if it was generated through physical testing to produce the inherent characteristic for fully developed turbulent flow, and the effect of choked flow on performance.

The premise of this project was not only to explore virtual valve testing, but to evaluate the practical and efficient use of CFD by the non-specialist design engineer. As a benchmark, the end user had no prior experience with the software when the project initiated, no formal training in the software, and depended on the included tutorials, help utility and thorough documentation to produce good results and good data. In the end, the time and effort required to run virtual valve tests compared well to equivalent physical tests. The investment to equip a valve manufacturer's test lab would be one to two orders of magnitude greater than the resource investment required to have the capability to run a "test in the cloud" of a control valve.

Benefits of the Cloud

Because of the computational requirements for each simulation, engineers often cannot set up and run the spectrum of simulations they would like within the time available. Limited hardware and software licenses often act as the bottleneck in a simulation project — leaving engineers to virtually test only a limited number of cases.

It is certainly possible to get around this problem by making a significant investment in computing hardware and software licenses in the form of HPC clusters and simulation solvers. This approach is appropriate for specialized simulation analysts at large organizations, which can typically maintain high enough utilization on the hardware and software to get a good return. The vast majority of engineers, however, have fluctuating demands for simulation that depend upon their project work. Most engineers involved in design work have intermittent simulation needs, but when they do require it, they demand hardware and licenses as any full-time analyst would.

Download the 2014 UberCloud Compendium

THE UberCloud Compendium is valuable resource for engineers, scientists, managers and executives who believe in the strategic importance of High-Performance-Computing-as-a-Service (HPCaaS) in the cloud. It's a collection of selected CAE case studies from the participants in Rounds 3 and 4 of the UberCloud Experiment. Among these case studies, you may find scenarios that resonate with your own engineering computing challenges. Download a free copy of the 80-page compendium that includes 17 different case studies across various industries. Reading them can help you to benefit from the challenges they encountered, problems solved, lessons learned and their expert recommendations.

The 2014 UberCloud Compendium can be downloaded for free here: deskeng.com/de/simulation-cloud.

Cloud computing enables the design engineer to access a large amount of computing power in a cost-effective way. Rather than owning the hardware and software licenses, engineers can pay for what they need when they need it, rather than making a substantial upfront investment.

The cloud reduced solution time by at least a factor of three, with the local workstation dedicated to the CFD analysis. During the project, local solution times increased by another factor of three when the local workstation was in use for intensive modeling work while a simulation was running. If a workstation is required to be actively used for other work during simulations, sending solutions to the cloud reduces the time to collect one data set for 10 scenarios by an order of magnitude.

Accounting for runtimes and data download upon completion of each of the runs, it would be possible to solve the 200 simulations for this project within a day. For an engineer with one simulation license on a single workstation, this would have required 800 hours — approximately 30 days — to complete if the simulations were running nonstop, one after another. Table 1 compares the approximate time and investment that would be required for various solving approaches.

Engineers who are interested in participating in a similar experiment to perform their simulations on an external technical cloud computing resource are invited to register at TheUberCloud.com/hpc-experiment. **DE**

Mark A. Lobo *is principal, Lobo Engineering PLC.* **Jon den Hartog** *is senior product line manager, Autodesk.* **Derrek Cooper** *is director of Autodesk's digital simulation product line.* **Wolfgang Gentzsch** *is an industry executive consultant for high performance, technical, and cloud computing. Contact them via theubercloud.com or send e-mail about this article to DE-Editors@deskeng.com.*

INFO → **AMD:** AMD.com

→ **Autodesk:** Autodesk.com

→ **HP:** HP.com

→ **UberCloud Experiment:** TheUberCloud.com

→ **U.S. Council of Competitiveness:** Compete.org

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

Control the Cloud

Public or private? There are tradeoffs in building your own cloud inside your organization, but it might be the right decision for engineering teams.

BY PETER VARHOL

As engineers seek ever-more-efficient ways of using high-performance computing (HPC) resources in running analyses and simulations, the number of available alternatives is greater than ever. In particular, the cloud is emerging as an intriguing option, especially for those groups not needing dedicated HPC support.

But even with running jobs in the cloud, there are choices to be made. When most think of the cloud, they envision a location unknown, but off-site. In reality, there is another option: a so-called "local cloud." Also known as a private cloud, this is typically a hardware appliance that serves up engineering applications locally to all engineers in the organization.

Why have a private cloud appliance? In many cases, the answer is obvious. The engineering data remains within the organization's firewall, where it can usually be secured more easily and reliably than with a commercial public cloud. While many public clouds offer higher levels of security, with encrypted data transfers, many engineering groups are simply reluctant to let proprietary design information offsite, out of the local network.

Local appliances can also provide HPC resources with a low systems management overhead. Appliances that offer full turnkey solutions, complete with analysis and simulation software applications, can effectively be almost free of system administration requirements.

While these appliances are powerful systems, they aren't clusters — which can have far more horsepower, but can sometimes demand larger and more complex IT management resources. For larger engineering organizations, they won't replace traditional clusters from a sheer performance standpoint. Where they can make a difference is in smaller groups that are using individual workstations, or for groups that cannot fully utilize a cluster. Private clouds can also be useful in larger groups that have occasional needs for more computing horsepower than their cluster can provide.

Large data centers, such as this one from IBM, are used as public cloud environments. They offer easy scalability, but let important engineering data outside the firewall. ***Image courtesy of IBM.***

Plain Appliance or Turnkey Solution?

Most of the large traditional system vendors, such as IBM, Lenovo, HP and Dell, offer appliances that integrate compute, storage, networking and system management into a single box. With the appropriate software licenses, engineering groups or a systems integrator can install their own analysis and simulation software on these appliances for use by the group. For those engineers with non-standard or unique application needs, starting with a commercial appliance and building out a solution is another way to share applications among an entire team.

An emerging solution is Microsoft Azure, which the company expects to be delivering in an appliance in the near future. Microsoft has already announced an Azure appliance for storage, archive and disaster recovery, called StorSimple, and is expected to follow that up with compute appliances. However, it's important to realize that StorSimple is actually a hybrid solution, with an appliance and the ability to offload to the Azure cloud. Should Microsoft also take that route with a compute appliance, it could offer some interesting advantages over a strictly public or private cloud alone. Teams could reap the advantage of sharing a common platform with exactly the applications they need, while also being able to offload archived data offsite.

For those whose software needs are satisfied by a single engineering applications vendor, it can be easier if that vendor offers an appliance, either as a systems integrator or in partnership with an HPC systems provider. One example of that type of solution is the ANSYS-IBM partnership to provide the full ANSYS suite on an IBM server appliance. One open question is whether that partnership will survive the sale of IBM's X86 server business to Lenovo. That proposed transaction was approved by the U.S. Treasury Department in August, and Lenovo expects the purchase to complete by the end of the year. Barbara Hutchings, ANSYS' director of strategic partnerships, still expects the partnership to continue after the acquisition.

Another fully turnkey solution is Altair's HyperWorks Unlimited. This is a private cloud with fully configured hardware and software, offering unlimited use of all Altair software on the appliance. The company says that it can be set up and running jobs in a few hours or less.

There are three different hardware offerings, ranging from six to 20 cores and 64GB to 256GB of memory. Altair uses SGI hardware, but sells and fully supports HyperWorks Unlimited itself, including the hardware. Like the software, the hardware is leased rather than sold, so ongoing support is an integral part of the lease.

It's also possible to turn an existing cluster into a cloud, for better manageability. Penguin Computing's Scyld HCA turns any HPC Linux cluster into a managed HPC hybrid cloud environment. Among the features it adds to the traditional cluster are providing virtual login host lifecycle management, object and data storage, user and group administration, metering and reporting, authentication, a Web-based portal, and a graphical user interface (GUI) for hardware and workflow management. If you're not yet ready to move entirely to a cloud appliance, this approach could be a good initial step.

Not Just for Computation

NVIDIA offers a different twist on the private appliance. While NVIDIA has been building a reputation for high levels of computational performance, its appliance offering is explicitly designed for graphics processing and performance, focusing on rendering complex graphics. The NVIDIA GRID Visual Computing Appliance (VCA) runs engineering design applications and sends their graphics output over the network to be displayed on a client computer.

The NVIDIA GRID VCA is certified and supported by Dassault Systèmes to virtualize and remotely deliver SolidWorks 2014 over the network. The NVIDIA GRID can also be used to accelerate AutoCAD designs and send the results to one or more clients. Right now, these are the only application partners available, but NVIDIA expects to be adding more in the future.

The NVIDIA GRID appliance consists of eight GPUs with a total of more than 23,000 compute unified device architecture (CUDA) cores, and 20 Xeon CPUs with hyperthreading. Each GPU has 12GB of memory, while the system memory is 256GB. It renders either application blazingly fast for up to eight concurrent users.

This solution provides flexibility for

engineers looking for a fast way of doing design rendering while also spreading the cost out among a small engineering team. Conceivably, it can also be used strictly as a compute server, with applications that have appropriate NVIDIA GPU support, but that may not be a cost-effective use.

The Good and the Bad

In many of these alternatives, software licensing policies remain a stumbling block. Some vendors' licensing policies are still complex, and weighted to be increasingly costly for more powerful systems. Groups that want to integrate their preferred application software with an appliance themselves have to understand, budget for and manage application licenses for the appliance.

While local appliances enable a group to feel more in control of its computing resources than with a public cloud, with that control comes a higher level of awareness. For example, if a group decides it needs additional cores or memory for a particular job, those resources are easy and seamless to obtain from a public cloud.

With a private cloud appliance, however, it requires either upgrading the existing hardware or buying an additional appliance. Both are time-consuming and expensive endeavors. The appliance can be shared, so it's still often a better solution than newer and more powerful workstations — but it still requires a significant budget. That's why some teams are looking at them as supplements to existing resources rather than workhorses.

A private cloud is typically a hardware appliance that can be hooked up and forgotten about. It may also include pre-installed engineering software, or engineers can customize it for their specific needs. *Image courtesy of vBC Solutions.*

On the positive side, the use of an appliance can be more flexible and cost-effective than running the same applications on individual engineering workstations. If a small group of engineers can successfully make more efficient use of the computing power in a single device, it can deliver more power to individual jobs at any given time.

Manageability is also one of the key advantages of the private cloud appliance. For groups that have limited IT support, or have to perform their own system administration activities, the appliance can save their time for real engineering work. In most circumstances, you can simply hook it up, turn it on, configure it and forget about it.

And while some still worry about security, many public clouds offer advanced security features, including encrypted transfers and protected virtual machines (VMs) on partitioned servers. While there is no such thing as absolute protection, under many circumstances it can be about as safe as it is within the firewall.

The bottom line is that the cloud for HPC, and even for rendering, is gradually becoming more popular. This is especially true for organizations without dedicated IT resources, and those without ongoing analysis needs. Engineering groups looking to address growing HPC needs with flexibility should start looking to the cloud. Whether that is a public or private cloud depends on the unique requirements of the group. Each has advantages and limitations, so any group should have a good understanding of their needs before determining a direction. DE

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

INFO → **Altair:** Altair.com

→ **ANSYS:** ANSYS.com

→ **Dassault Systèmes:** 3DS.com

→ **Dell:** Dell.com

→ **HP:** HP.com

→ **IBM:** IBM.com

→ **Lenovo:** Lenovo.com

→ **Microsoft Azure:** Azure.Microsoft.com

→ **NVIDIA:** NVIDIA.com

→ **Penguin Computing:** PenguinComputing.com

→ **vBC Solutions:** vbcloud.com

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

Parallel Programming

To take advantage of multiple cores, software must be programmed for parallel computing. It can be a challenge for software developers accustomed to years of serial programming to make the transition, says Zvi Tannenbaum, CEO of Advanced Cluster Systems, Inc. (ACS).

ACS's Supercomputing Engine Technology (SET) is a software development framework and implementation library. Tannenbaum says it is intended to make parallel programming and scaling of serial software easy for those familiar with serial programming, including software vendors and even small- and medium-sized businesses. Additionally, the SET framework can create ad hoc parallel computing machines out of collections of individual Mac and Linux computers, and utilize local computing clusters.

For more information, visit advclustersys.com.

— Jamie Gooch

New Processors Drive Innovation

The quest to deliver higher performance levels, while keeping power consumption and costs under control.

BY FRANK J. OHLHORST

If one thing could be said about the latest crop of processors, it's that Moore's Law is still in full effect. Evidence for that decades-old theory can be found in Intel's and AMD's latest offerings, backed by an onslaught of new video card offerings.

Intel Brings More Cores to the Party

With its Sept. 8 launch of its latest Xeon processors (code named Haswell-EP and Haswell-EN), Intel aimed to maintain its lead in the high-performance and server computing market. The Xeon E5-2600 V3 family of processors, based on Haswell microarchitecture, features as many as 18 cores. It uses the C612 chipset, with improved performance in an R3 socket over its predecessor, the Xeon E5- 2600 V2, which is based on Ivy Bridge architecture.

"The higher core count brings significant performance improvements over the previous generations, as well as improved peripheral expansion," says Intel Workstation Segment Manager Wes Shimanek.

During the launch event at the Intel Developer Forum (IDF) in California, Intel's SVP of its Data Centre group, Diane Bryant, said, "From a performance perspective, we are delivering worldwide performance levels, tripling performance [thanks to] the Xeon V3's 18 cores [that] offer a 50 percent increase over the prior generation."

Beyond an increase in cores and a redesigned microarchitecture, performance gains offered by the latest Xeons come from innovations such as the switch to DDR4 memory of up to 2,133MHz, outpacing Ivy Bridge's 1,866MHz DDR3 memory.

"We are continuing to be the most energy-efficient processor on the planet, and have achieved 27 world records in performance across the product line," Bryant said.

On the heels of the Intel launch, major workstation vendors announced new systems that will leverage the power offered by Intel's latest Xeons.

HP was quick to demonstrate its latest high-performance workstations at a press-only event at its Fort Collins, CO headquarters in early September. At the event, Jeff Wood, worldwide product manager for the HP workstation business, explained that the new HP Z Desktops and ZBook Mobile Workstations are built to address the constraints of compute-intensive industries.

By combining the latest technologies, HP is offering workstations that speed the most intensive processes for design and simulation. For example, the top-of-the-line Z840 supports up to 36 cores, 1TB of error-correcting code (ECC) memory, and two NVIDIA K6000 (12GB) graphics cards in a full-size tower configuration.

Jim Zafarana, VP and general manager of HP's commercial solutions unit, said at the event "the new Haswell processors, when paired with the latest in high-performance graphics cards, allow HP to bring exceptional levels of performance to engineering, design and simulation users."

HP isn't the only one delivering the latest Haswell processors in workstation form. Dell has introduced tower and rack workstations based on the new Intel Xeon E5-2600 v3 and E5-1600 v3 processors. The company claims that the new workstations, which are equipped with 2133MHz DDR4 memory, perform as much as 4.7 times faster than the previous generation, making it possible to design and iterate faster than ever.

"Every key component of our new workstations has been updated, resulting in outstanding performance gains to help our customers design without limitations," says Thaneth Angkasirisan, country manager, Dell Corp.

AMD's Latest FX Processors Focus on Value

AMD also had some CPUs news of its own in early September. The company introduced a trio of eight core CPUs built on AMD's 32nm Piledriver architecture. From a performance standpoint, it would be unfair to compare AMD's 32nm architecture to Intel's latest Haswell 22nm architecture; however, AMD's is focusing on value. AMD's 4GHz FX-8370 and 3.3GHz FX-8370E are priced as $199.99 each, while the 3.2GHz FX 8320E is priced at $146.99.

AMD's new FX processors are geared toward enthusiasts, and the company is touting high clock speeds as an advantage, as indicated by Bernd Lienhard, corporate vice president and general manager, Client Business Unit, AMD.

"Enthusiasts can experience seamless multitasking and unleash the real capabilities of multi-threaded applications with additional enthusiast performance-tuning features unlocked as standard," Lienhard says. "World-record frequencies are just a start."

For engineers looking to build multi-display, low-cost systems for HD content creation, the latest FX processors from AMD may be a viable choice — especially when cost is a key concern.

AMD, NVIDIA Bring Next-generation GPUs to Market

Both AMD and NVIDIA have launched new graphics cards that feature faster and more efficient silicon, pushing graphics

	Quadro K6000	Quadro K5200	Quadro K4200	Quadro K2200
Base core	GK110 (Kepler)	GK110 (Kepler)	GK104 (Kepler)	GM107 (Maxwell)
GPU production process	28 nm	28 nm	28 nm	28 nm
CUDA cores	2880	2304	1344	640
Texture units	240	192	112	40
Raster units	48	32	32	16
GPU frequency	900 MHz	667 MHz	771 MHz	1046 MHz
Graphics memory size	12 GB	8 GB	4 GB	4 GB
Single-precision FP performance	5.2 TFLOPS	3.0 TFLOPS	2.1 TFLOPS	1.3 TFLOPS
Dual-precision FP performance	1.73 TFLOPS	0.13 TFLOPS	0.09 TFLOPS	0.04 TFLOPS
Graphics memory type	GDDR5	GDDR5	GDDR5	GDDR5
Memory bus width	384-bit	256-bit	256-bit	128-bit
Memory frequency	1500 MHz	1500 MHz	1350 MHz	1250 MHz
Memory bandwidth	288 GB/sec	192 GB/sec	173 GB/sec	80 GB/sec
OpenGL	4.4	4.4	4.4	4.4
OpenCL	1.1	1.1	1.1	1.1
DirectX	11	11	11	11
Heat dissipation	225 Вт	150 Вт	105 Вт	68 Вт
Video outs	1 x DVI-I 1 x DVI-D 2 x DP	1 x DVI-I 1 x DVI-D 2 x DP	1 x DVI-I 2 x DP	1 x DVI-I 2 x DP

NVIDIA's updated Quadro products feature more onboard memory (twice the amount of their respective predecessors) and are designed to work better with complex models and at higher, 4K resolutions.

	AMD FirePro W9100	AMD FirePro W9000	NVIDIA Quadro K6000	NVIDIA Quadro K5000
Base core	Hawaii	Tahiti	GK110	GK104
GPU production process	28 nm	28 nm	28 nm	28 nm
Shader processors	2816	2048	2880	1536
Texturing units	176	128	240	128
Raster units	64	32	48	32
Core frequency	930 MHz	975 MHz	900 MHz	700 MHz
Graphics memory size	16 GB	6 GB	12 GB	4 GB
Single-precision floating-point performance	5.2 TFLOPS	4.0 TFLOPS	5.2 TFLOPS	2.2 TFLOPS
Dual-precision floating-point performance	2.6 TFLOPS	1.0 TFLOPS	1.73 TFLOPS	0.09 TFLOPS
Graphics memory type	GDDR5	GDDR5	GDDR5	GDDR5
Memory bus width	512-bit	384-bit	384-bit	256-bit
Memory frequency	1250 MHz	1375 MHz	1500 MHz	1350 MHz
Memory bandwidth	320 GBsec	264 GB/sec	288 GB/sec	173 GB/sec
OpenGL	4.3	4.3	4.3	4.3
DirectX	11.1	11.1	11.1	11.1
Heat dissipation	275 W	274 W	225 W	122 W
Video outs	6 x mini DP	6 x mini DP	1 x DVI-I 1 x DVI-D 2 x DP	1 x DVI-I 1 x DVI-D 2 x DP

With specs evolving and the GPU wars in full swing, one cannot think about AMD's FirePro GPUs without comparing them to market leader NVIDIA, which aims to hold onto the performance GPU market with its Quadro GPUs.

processing unit (GPU) performance upward.

Building upon its recently introduced Maxwell graphics architecture, NVIDIA has launched the GeForce GTX 980 and GTX 970. The new cards add more efficient multi-frame anti-aliasing, as well as Dynamic Super Resolution, which will render content at 4K and then downscale it for an improved 1080p image. NVIDIA has also added improvements for virtual reality processing by including the assignment of one or more GPUs to render the visuals for a specific eye.

Although the GTX 980 and the GTX 970 are geared toward the gaming market, NVIDIA's latest Maxwell-based cards do offer advantages for those looking at simulation and design. NVIDIA's Tom Petersen touts the enhancements as ideal for hardcore gamers, but notes that "Voxel Global Illumination is a technology that can simulate light in real-time, creating scenes that are much more realistic, vastly improving the quality of real-time rendering."

However, NVIDIA's performance-leading GPUs remain with the Quadro product line — with the well-established K6000 at the top of the heap. A new lineup of Quadro GPUs, in the form of the K5200, K4200, K2200 and K620, are squarely aimed at archrival AMD's FirePro series of GPUs.

AMD's latest GPUs include the FirePro W9100, which the company says is designed for the next generation of 4K workstations accelerated by OpenCL, and can deliver as much as 2.62 teraflops.

"The AMD FirePro W9100 is designed to deliver ultimate visual supercomputing performance for the next generation, ultra-high-resolution workflows," says Raja Koduri, corporate VP, AMD Visual Computing. "This professional graphics card is optimized and certified for leading workstation applications, and ensures ultra-high geometry performance and smooth handling of complex models."

More Processor News on the Horizon

Announced in July, Intel's latest Core M is the first consumer-level processor to use the latest 14nm design. Positioned as a low-power usage, system on a chip (SoC) processor, the Broadwell-Y based Core M creates a great deal of opportunity for those designing mobility products or embedded devices.

Broadwell-Y can be used as the basis to manufacture a range of products, from high performance to low power. Improvements include a redesign of each transistor fin, which is now denser, making the 14nm much smaller and more efficient, and allowing for the creation of "fanless" devices.

Broadwell-Y's SoC design delivers lower capacitance. Intel says it reduces power input by 25% and minimum operating voltage by 10%, which saves an estimated additional 20% in power consumption. Also, Broadwell-Y offers 10% to 15% performance per transistor increase through the use of low-voltage, optimized 14nm transistor technology, backed by 2x lower leakage. This further reduces power consumption by 10%, as most of the power is utilized instead of being lost to leakage.

"The goal with Broadwell-Y was to make an 8mm to 10mm thick device, with a 10.1-in. display that was fanless with a 3 to 5 Watt SoC," explains Stephan Jourdan, Intel Fellow and director of SoC Architecture at Intel. "Intel needed to optimize the CPU, GPU, PCH and move onto the 14nm process to create a sub-5 Watt SOC that could fulfill such a role." DE

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

INFO ➜ AMD: AMD.com

➜ **Dell:** Dell.com

➜ **HP:** HP.com

➜ **Intel:** Intel.com

➜ **NVIDIA:** NVIDIA.com

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

The Promise of Instant Remote Access

Teradici releases a software-only solution for remote workstation access.

BY DAVID COHN

Dell and Canadian company Teradici used the SIGGRAPH 2014 conference in Vancouver to announce a joint effort to make it easier than ever before for users to instantly and securely access their workstations from anywhere using an inexpensive software-only virtual machine (VM) solution (see "Wowing Them in Vancouver at SIGGRAPH 2014" on deskeng.com/?p=18757). The news from Vancouver was actually an extension of a partnership going back more than 5 years, in which Teradici PC over IP (PCoIP) technology has become a key ingredient of desktops and workstations in Dell's portfolio of enterprise products.

The new Teradici PCoIP software became available to Dell customers through the Dell Online Store, and through Dell's channel partners for existing and new Dell Precision workstations beginning in August. The software will be rolled out globally to other customers in the first half of 2015.

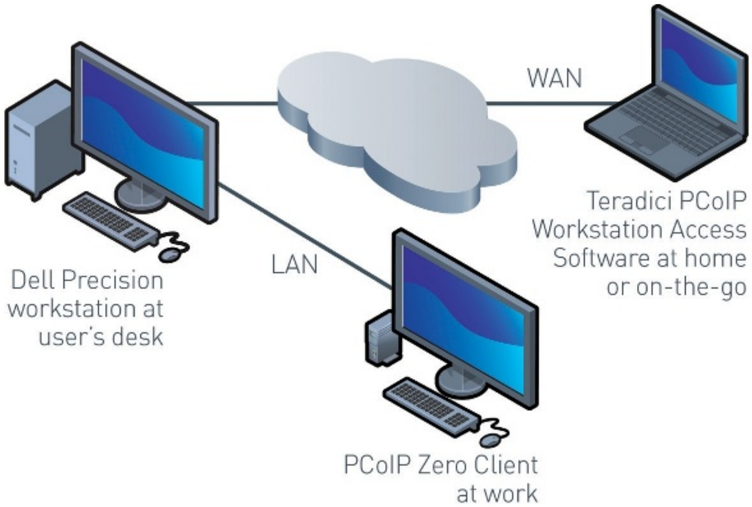

Teradici's PCoIP software promises an easy-to-install and easy-to-manage solution with a seamless remote end-user experience — allowing users to access a workstation at their desk from either a zero client at work or a software client at home or on the go.

The Teradici announcement generated a lot of buzz at SIGGRAPH. What makes Teradici's new product so exciting is the promise of a software-only VM solution that is easy to install, use and manage. Once installed, the Teradici PCoIP software runs in the background as a Windows service. It promises to be compatible with any PCoIP end point, including the company's existing PCoIP zero clients and its new software client. Both iOS and Android versions of that client software will soon become available, making it possible to access and run workstation-based applications such as Autodesk Inventor and Dassault Systémes' SolidWorks — or any other Windows-based application — from an iPad or Android tablet, as well as from a remote PC or Mac.

According to Teradici representatives, if you can install other programs on your computer, you can install the Teradici PCoIP software. The software is ideally suited to small- and mid-sized businesses in architecture, construction, engineering and manufacturing, or companies with fewer than 100 workstations in a department. The company says the software can be installed in less than five minutes, giving users instant access to CAD/CAM applications on Windows devices for a seamless remote connection anywhere.

The initial release of Teradici's software provides a convenient, low-cost, simple way for users to get up and running, with no impact to their current deployment and no need to have to get up to speed on new IT management tools. Regardless of the application being used, the Teradici PCoIP protocol compresses, encrypts and transmits pixels rather than data, thus protecting against possible loss or theft of intellectual property. It also supports dual monitors and resolutions up to 2560x1600. It maintains full 3D application programming interface (API) compatibility (OpenGL and DirectX), so the remote experience is the same as if you were sitting at your workstation. The technology is compatible with any GPU hardware, but is optimized for NVIDIA K2000 or higher graphic processing units (GPUs).

Teradici's PCoIP software is an expansion of its existing portfolio that supports encoding up to 50 megapixels per second on any Windows-compatible device. For higher-end workstation performance, the company continues to sell remote workstation host cards that deliver even faster performance.

According to Teradici, the software is easy to deploy and does not require set up of a third-party broker to get started. The local workstation user can share his or her screen with one remote user leveraging the same Windows account. The host software runs on 64-bit versions of Windows 7 or 8, while the client software is currently available for both 32- and 64-bit versions of Windows XP and Windows 7, Windows 8.1 64-bit, and Mac OS X 10.7 or later, with support for iOS and Android devices coming soon. The software is not intended to compete with products such as Hewlett-Packard's Remote Graphics Software (HP RGS), which is a client-server remote desktop software solution.

The new Teradici PCoIP software is priced at $199, plus $40 for a one-year mandatory support and maintenance contract, which includes free software upgrades (major and minor releases) and 24/7 support services. The software is sold on a per-workstation basis, with one license per machine and no restrictions on the number of users who can remotely access it at different times.

According to a recent survey published by Jon Peddie Research, end users report a strong interest in remote graphics and virtualization. But significant numbers felt that the technology was too expensive, too technical or too complicated for adoption at their firms. Teradici's PCoIP software may be the solution to those concerns. *Desktop Engineering* plans to publish a complete hands-on evaluation of this promising product in an upcoming issue, so stay tuned. **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, graphics cards and monitors since 1984. He is a contributing editor to Desktop Engineering and the author of more than a dozen books. You can contact him via email at david@dscohn.com or visit his website at DSCohn.com.*

INFO → Teradici Corp.: Teradici.com

Teradici PCoIP Workstation Access Software

- **Price:** $199 plus $40/year support and maintenance
- **Host software requirements:** Windows 7 64-bit, Windows 8.1 64-bit
- **Host GPU compatibility:** Any discrete GPU, optimized for NVIDIA K2000 or higher
- **Remote client software requirements:** Windows XP 32-bit or 64-bit, Windows 7 32-bit or 64-bit, Windows 8.1 64-bit, Mac OS X 10.7 or later

Image courtesy of iStock

IT consultants can help you deploy and maintain HPC resources.

BY BRIAN ALBRIGHT

As more engineering firms and departments turn to high-performance computing (HPC) clusters to run their applications and simulations, the need for outside help to manage and maintain these systems has also increased. Because many companies don't have in-house HPC expertise in the IT or engineering departments, IT consultants can step in to help guide the HPC strategy; in some cases they provide managed services that allow their clients to outsource HPC support.

"A lot of these companies have engineers or researchers that are running commercial applications like Autodesk on desktops, and then realize at some point that they have to go beyond the desktop," says Deepak Khosla, president and CEO of eXcellence in IS Solutions (X-ISS), a Houston-based HPC consulting and services firm. "The workloads coming in are too large to finish the job in a reasonable amount of time."

"They are open to a centralized computing model as long as it improves efficiency," says Rod Mach, president of TotalCAE, an Ann Arbor, MI-based firm that provides cluster installation, support and cloud-based solutions for CAE and CFD applications. "There's an adjustment period, because it's a different way of working and there are new things to learn."

While smaller and mid-size firms are able to deploy HPC clusters on-site, often with the help of their commercial software vendor, they may not have the internal resources to properly maintain and configure the cluster. "Most IT staff are not knowledgeable on HPC clusters because they are Linux-based, and these are more than just a collection of servers," Khosla adds.

Mach agrees. "Engineering is a specific niche, and corporate IT often doesn't have the time or resources to invest in learning those applications," he says. "They want to be able to quickly support new apps, and have them run efficiently."

In some cases, engineers or researchers wind up helping maintain the cluster. "That usually doesn't end well," Khosla says. "First, the value of that person is not in managing the cluster, it's in engineering. And second, they may be able to make things work for themselves, but they often mess it up for other users in the process."

HPC consultants can provide guidance that will help direct the firm toward either an on-site or cloud-based solution, comparing the relative return on investment (ROI) for both. "The software we represent has a high computer power demand, and the engineers approach us because their own IT department is focused on a certain way to solve the problem," says Carl Smith, professional services manager at IMAGINiT, an engineering software and services provider. "When engineers encounter these efficiency problems, often the internal IT staff will suggest buying a new server or adding some more RAM. Often management will task the IT staff with looking for an alternative to spending money on new infrastructure."

The Outsourced IT Department

When engineering firms turn to consultants for HPC help, typical pain points include:

- they need applications to run more smoothly;
- simulations need to be faster; or
- they are struggling with storage of simulation results.

For cloud-based implementations, users want to know what alternatives are available, and what the security model looks like.

In addition to providing guidance on HPC system selection, IT consulting firms can also provided outsourced IT services that can encompass complete management of the entire HPC infrastructure and applications, including installation, upgrades and all maintenance and troubleshooting.

For companies with on-site HPC resources, these outsourced support services can largely be rendered remotely.

"Once things are set up, the only time you need someone on-site is when the hardware fails or there's an access problem," Khosla says. "We can manage everything remotely, even when the hardware dies. As long as it's in the rack and connected, we can bring it up from the bare bones and redeploy."

Smaller, and even some mid-sized firms often don't have a dedicated IT staff, so using outside resources is a familiar model.

"A lot these companies outsource to a varying degree," says Smith, offering as an example how a contractor might come in "every Tuesday and Thursday and do a health check."

Mach notes that TotalCAE provides a full outsourced IT department for engineering groups, including all applications, licensing, server maintenance and other services.

"We like that model because no matter what's going wrong, they come to us," he says. "There's only one person to call. You don't have to try and decide whether the problem lies with corporate IT or us. We're an island on the corporate network."

HPC Assessment Process

Companies that are doing an initial HPC evaluation will typically go through an assessment process to determine what type of solution they need, and whether those resources will reside on-site or in the cloud.

As part of that process, the firm should determine whether any quality compliance or government security requirements are in place. "Some of these compliance requirements don't allow the clusters to be cloud-based," Smith says.

What kind of workload will the HPC cluster need to handle? "In some cases, the workload requirements are going to knock cloud options out of the running," Khosla says. "If that's not the case, then you have to do an ROI calculation. Assuming you have information on how much data you'll be pulling and storing, you can take the best price and compare that to the cost of handling it internally."

If you don't have good information about the workload, proof-of-concept testing can help determine whether a cloud or on-site solution would be better for your business.

Mach says that his company will typically meet with clients to go through the types of applications they are running, determine what their goals are, the turnaround time involved, and how to optimize licenses costs and runtime.

"We take that into account when designing a system or workflow for them," Mach says. "We work with engineers primarily as the customer. Often there is a meeting with IT, but engineering is the direct interface."

A cost evaluation is also important, particularly when deciding between on-premise or cloud-based solutions. It may make sense to have cloud computing available so that simulations can run in the background without tying up computers, but in many cases a local solution will be required. Factors to consider include internal networking capabilities, security requirements, compliance issues and cost.

"The main interest in cloud computing is saving money, and the cloud doesn't always necessarily save you any money," Mach says. "If you run simulations all day, it saves no money. Licensing is solver-specific, and not every CFD (computational fluid dynamics) solver licensing scheme is pay-as-you-go."

Cloud deployments don't require ongoing management, but as Khosla points out, getting set up for a cloud solution is not trivial.

"There are all of these things you have to consider around security, data transfers and workflow, and companies will often need help with that," he says.

Find the Right Partner

When evaluating potential IT consulting partners for an HPC project, look for companies with experience in your environment and with your enterprise resource planning (ERP) and business systems.

"You want to make sure the résumé behind that team has experience with a business the same size as yours," Smith says.

The company should also be familiar with the applications you are using. "If they don't know the applications, they can't benchmark the models and can't provide the performance you want to get, or they just don't speak your language," Mach says. "That's especially important for engineering. It doesn't matter what all the benchmarking says; if it doesn't run on your model, then it is irrelevant."

Khosla emphasizes that application knowledge is more important than specific vertical industry knowledge. "You have to know the application, and it helps to know the industry a bit, but that's not required," Khosla says. "We aren't serving as the application expert. We're making sure whatever application is running is performing at its best in the HPC environment."

Test and pilot the solution or services in advance, and don't surprise your internal IT department with the project, Mach suggests. "Don't go too far down the path without involving the corporate IT department, or you could get pushback," he says. "No one likes surprises. We like to identify the problem, and work with the corporate IT team in conjunction."

The consultant you choose to partner with should also have deep experience in high-performance computing projects. "It's important that the company you engage displays a real-time knowledge of the pros and cons, pitfalls and benefits of an HPC solution," Smith says. "It would certainly warrant an investment in a business process analysis of all aspects of your business, particularly if you are determining if you can use a cloud-based solution for all or part of your business."

Finally, check references, Khosla says. "And have a list of questions and concerns for the vetting process," he adds. "Look at the company's longevity, expertise, abilities and quality of service." **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 → **CAutodesk:** Autodesk.com

→ **eXcellence in IS Solutions:** X-ISS.com

→ **IMAGINiT:** IMAGINiT.com

→ **TotalCAE:** TotalCAE.com

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

Mobile Server to the Rescue

Gregory Carter owns and operates Applied Engineering Software Group (AESGI) in Green Bay, WI. The firm manages storage, network and server operations for third-party contractors for a variety of companies, primarily in the Internet service provider (ISP), telecommunications and logistics and supply industries. Secondarily, AESGI deals with enterprise operations, analytic measurement and problem resolution in the same areas. A lot of Carter's time is spent tracking down problems and redesigning issues of data center operations or cloud computing operations that are not working optimally.

"Sometimes a customer will become frustrated because of a vendor deadlock to resolve the problems, so they hire me to come in and look at a problem as a neutral party," Carter explains. "I then redesign the software, add software or debug the software to identify the problem so that the vendor, customer or I can fix/resolve it. The typical projects for my company tend to be fairly diverse, both large and small in the types of server, storage and networking and or software problems I have to deal with."

Carter was initially attracted to the Eurocom Panther 5SE because of its ability to handle a server class Intel Xeon processor. The 5SE is an all-in-one system that results in a 12-lb., 16.76 x 11.44 x 2.31-in. mobile server solution. AESGI is sometimes running anywhere from 50 to 100 virtual machines, and requires the ability to model complex network problems.

Before the Panther 5SE Mobile Server, Carter used to haul around a desktop class system in excess of 80 lbs. to and from customer sites. The problems didn't end there, though — sometimes simply plugging in the desktop machine would pull 600 watts and trip an uninterruptible power supply (UPS).

"When I saw an X79 class chipset in a laptop that could run a Xeon class server processor, I decided to have a look at the Panther 5SE," Carter says. "I found the machine could run all of the tool sets I use, just as well as the desktop server machine I would haul around."

MORE → deskeng.com/de/?p=20085

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

Designing On the Go Evolves

Construction and survey projects emerge as testing grounds for mobile CAD.

BY KENNETH WONG

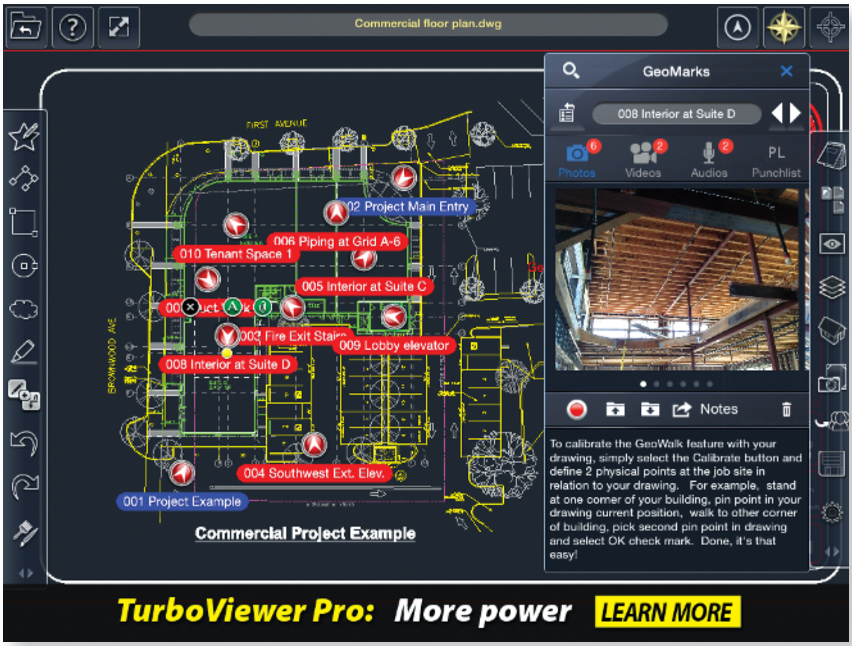

IMSI/Design's TurboSite (shown here in iPad) offers users a way to use the device's location-awareness to add geomarks, location-based notes, photo galleries and videos. *Image courtesy of IMSI/Design.*

At the Deutsche Bank headquarters in Frankfurt, Germany, Thomas Prinzen and two of his colleagues are starting their work just as the bank staff is heading home. With glowing mobile tablets strapped to their arms, they walk around the bank, taking measurements of the rooms and walls. Throughout the process, each of them is accompanied by an armed guard — a standard security protocol. By the time they go home, it's usually the wee hours of the morning.

Prinzen is part of the management circle of Graebert, best known for its ARES Commander CAD software. As director of Graebert's iSurvey Division, Prinzen oversees and participates in survey projects that involve capturing as-built conditions of existing architecture and interiors, ranging from Scottish schools and historic churches to McDonald's restaurants. In the case of Deutsche Bank, the Graebert iSurvey team has been asked to record the floor plans of the bank's headquarters, along with accurate positions of the furniture. The resulting DWG drawings are expected to offer facility managers valuable insight.

For the Deutsche Bank project, Prinzen's team uses Graebert's SiteMaster software, available for Windows 8 tablets. The measurements they make with their Leica Disto laser devices go straight to the software. As they work through the evening, the DWG files automatically develop in response to their measurements.

Currently, the use of lightweight devices equipped with mobile CAD apps is still largely confined to viewing, markup and approvals. But the iSurvey team's work is a harbinger to onsite work crews using lighter, nimbler mobile software and hardware to increase efficiency.

Betting on Mobile

In the late 1990s, when executives started expressing interest in exchanging their leather-bound appointment books for shiny PalmPilots, some brave companies tried to put CAD (by today's standard, a very limited function set) on those early personal devices. Cedric Desbordes, Graebert's marketing executive, recalls, "Graebert's first mobile CAD solution for Windows CE was released in the years 1999-2000 under the name FelixCAD Mobile. It was the first solution for Mobile CAD on Windows CE devices, and was later renamed PowerCAD CE."

Graebert's PowerCAD CE evolved into PowerCAD SiteMaster in 2003. "SiteMaster is now powered by ARES, and we recently released SiteMaster (Building Information Modeling, or BIM), adding the unique capability to export the 3D model of the building in IFC format for further modification with BIM software such as Autodesk Revit," Desbordes says. "The next milestone is to migrate SiteMaster to Android OS using the ARES Touch engine."

Desbordes says he believes engineers and designers won't work exclusively in mobile devices. Instead, they'll work in a "cross-device CAD experience," switching to different devices as circumstances demand. For example, they may approve a change from their phone, gather measurements onsite with a mobile tablet, and edit the DWG file back in the office on a proper workstation.

This workflow requires cloud storage, where the same project files and folders can be remotely accessed by several devices. According to Desbordes, "Sharing of projects for SiteMaster is currently done via email, but because the tablets run on Windows 8, it is also possible to install Dropbox or other cloud solutions, just like on your desktop. For ARES Touch, it will be possible to share files in the cloud using solutions such as Dropbox or Google Drive, and also by email. Of course, when SiteMaster migrates to ARES Touch, cloud sharing from the app will become available, too."

In June, Graebert announced its plan to develop a 2D-3D hybrid CAD program for Android devices. The new product is based on Graebert's ARES software, the engine behind Dassault Systèmes' DraftSight and Corel's CorelCAD.

"By 2015, tablets will outsell PCs by a significant margin; as such, the importance of making the full ARES CAD experience available on mobile devices became a top development priority," observes Wilfried Graebert, CEO and founder of Graebert.

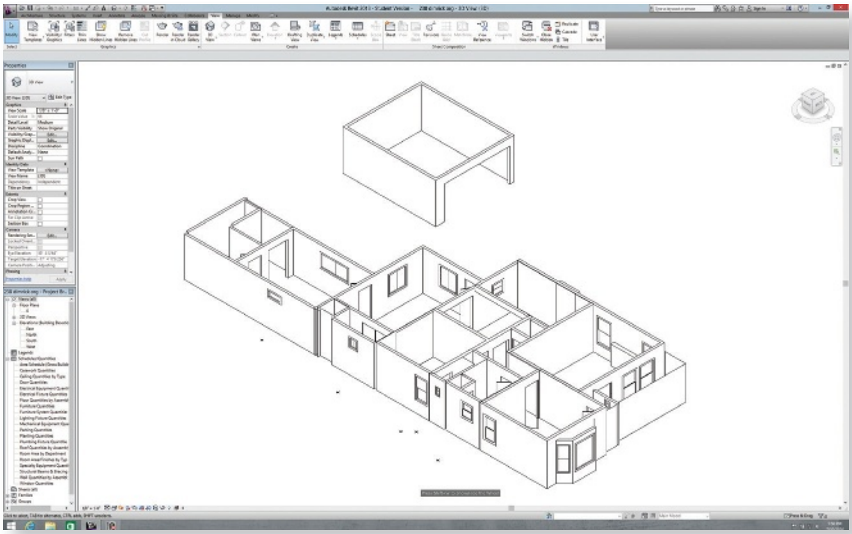

SiteMaster from Germany-based Graebert (shown here running on an iPad) is used with a custom tablet that can be worn around one's wrist or waist for onsite measurements. *Image courtesy of Graebert.*

Mobile CAD Landscape is Purpose-Built

Many mobile CAD apps are well tested and already established as an ideal way to view, annotate and share lightweight 3D files.

In cases where sharing 3D CAD assemblies may be overkill, users turn to mobile viewers like GstarCAD MC from China-based GstarCAD, eDrawings Mobile from SolidWorks, 3D VIA Mobile from Dassault Systèmes, TurboViewer from IMSI/Design, iXVL Player from Lattice Technology, and others. Some of these apps offer a more robust version (usually called "Pro"), with enhanced markup tools or additional storage for a modest upgrade fee.

Mobile apps have also become a viable alternative for quick computation based on established formulas. Autodesk ForceEffect for mobile devices and NEi Stratus belong to this class. In product lifecycle management (PLM) and collaboration, mobile apps like GrabCAD Mobile, Siemens PLM Software's Teamcenter Mobility, and Autodesk 360 Mobile offer project participants a way to quickly access and view shared files, track progress and respond to outstanding queries.

However, the idea to use a mobile app to create designs from scratch still remains a topic of debate among users. It's a need not easily filled by remote desktop apps, because the touch-sensitive navigation employed by mobile devices demand a drastically different way of interacting with engineering drawings and 3D models. Mirroring or replication the desktop CAD program's interface on a multi-touch device doesn't usually offer a good user experience. The answer may rest with a newer crop of mobile apps, developed specifically to create designs from start to finish.

— K.W.

"By 2017, there will be five touch devices (smartphones + tablets) shipped for each PC, and we fully intend to be the go-to solution for handheld devices long before then."

Graebert resellers are expected to offer bundled licensing in the near future: ARES Commander desktop license that offers ARES Touch, the mobile version, for free; and vice versa. At press time, the announcement was set to occur in Berlin, at the company's annual October meeting, Desbordes reveals to *DE*.

Integrating the Device's Strength

For IMSI/Design, which develops and markets the TurboCAD product line, the mobile device users are not just a secondary market to augment its desktop users. The company considers this to be the future, worthy of investing in an entirely new ecosystem.

In 2012, during the release of the free TurboViewer SDK, IMSI/Design's chairman and CEO Royal Farros said, "We are going to enable an entirely new generation of CAD developers, letting them create powerful, custom apps based on our TurboViewer series, the fastest 2D and 3D viewing platform in the mobile world. We think this can be a real game changer."

If users show sufficient enthusiasm for the SDK and are willing to develop programs on top of the viewer, IMSI/Design wants to nurture a mini-marketplace, an online zone where people can trade and buy specialized TurboViewer-based apps.

Within its TurboApps product line, the company also offers TurboSite, developed specifically for mobile work crews in manufacturing plants, factories and construction sites that must produce instant punch lists and highly customized PDF field reports. Among its many features is geomarking, the ability to mark a physical object's location using the mobile device's location awareness. The app also lets users take advantage of the device's built-in camera to attach galleries of photos, videos, dictations and text notes to a location to explain its relevance or importance.

Don't Duplicate the Desktop Interface

Anuj Kumar, a final-year architecture student, uses ZWCAD Touch on his iPad. The app was developed by China-based ZWSoft, which also develops and markets the desktop Windows-based ZWCAD+, ZWCAD Mechanical, and ZWCAD Architecture. He notices, "Right now, people rarely think about checking CAD files on their mobile devices. Some of my friends are unaware that they can even open CAD files or modify them on Android phones or iPhones."

ZWCAD and ZWCAD Touch both offer the ability to export DWG files, compatible with AutoCAD. That works well for Kumar, who uses AutoCAD 2014 and SketchUp on his desktop machine. Kumar says, "I think mobile CAD apps are a must-have for architects who are always visiting sites. While PDFs have helped me in past, CAD apps go much further; they create ultimate file-viewing experience on the move."

Kumar notes that for now, mobile CAD apps are better suited for viewing, editing and markup — not quite there yet for design creation from scratch. Developers "should not try to replicate desktop experience," he says. "They should realize most users will just be viewing or reviewing [on mobile devices]."

Replicating the desktop CAD program on a mobile device can now be done with apps like RemotePC (Android), Microsoft Remote Desktop (iOS), or Citrix's ShareConnect (iOS and Android). After a few registration steps and plug-in downloads, you'll find that you can remotely access your desktop machine, along with the CAD programs installed in it, from your mobile tablet or phone. But, to Kumar's point, you'll also find that it's extremely awkward to duplicate your usual mouse-and-keyboard commands (like the click-and-drag operations or the right-mouse menu items) using the mobile device's finger input.

Site Conditions

The most promising use of mobile apps at the present appears to be for work crews in remote locations. But to make the mobile experience optimal, the developer must take into account how the intended user might use the program, and provide additional accessories if necessary.

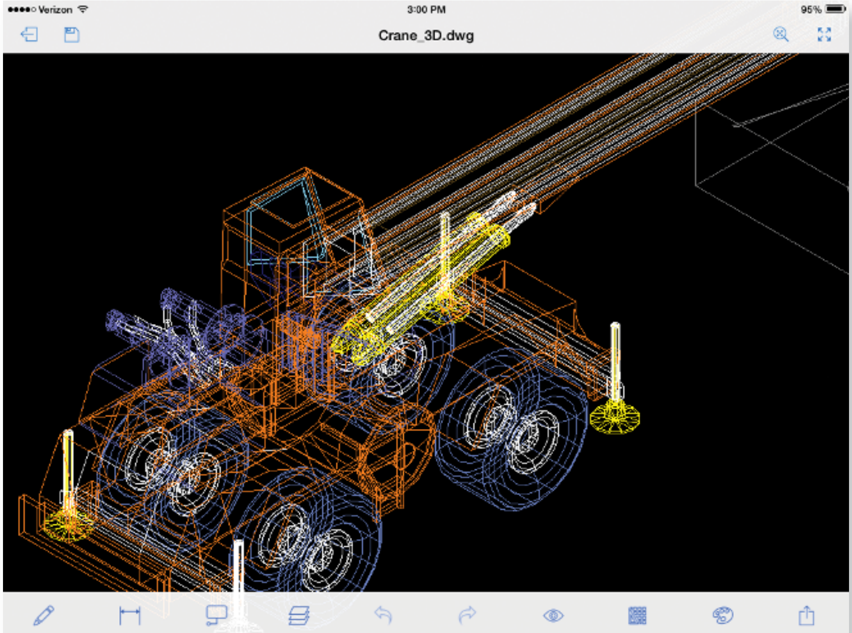

China-based ZWCAD offers ZWCAD Touch (shown here in iPad), a mobile version of its ZWCAD drawing and drafting software.
Image courtesy of ZWCAD.

"To measure buildings, we need to be extremely mobile — enough to be able to climb ladders, jump and down, and keep both hands free," Graebert's Prinzen points out. "You can't do that with a laptop."

This led Graebert to offer SiteMaster software on a PC-like device that can be worn around the wrist, or a tablet that can be worn around the waist. The company also offers a customized version of SiteMaster for kitchen design and survey, dubbed SiteMaster Kitchen. If there's sufficient demand, Graebert is said to be open to offering other specialized apps, like SiteMaster Plant or SiteMaster Bathroom.

Desbordes provides more examples of app developed with specific industries and workflows in mind. He says, "Connecting a Bluetooth laser telemeter with CAD software on a tablet is helping kitchen designers take or check measurements onsite about twice as quickly, and to get very easily and quickly the 2D and 3D geometry of the room without CAD experience."

Connecting Graebert's Mobile CAD solution SiteMaster Kitchen with the kitchen design solution from German-based company Carat is bringing up to 70% more productivity, he adds: "Before that, the kitchen planners had to enter and update manually the geometry in Carat, and an error of even 1cm had heavy financial consequences."

Moving Beyond Viewing and Markup

Perhaps one of the most robust mobile CAD apps in the market, and certainly one of the earliest to come out, is Autodesk's AutoCAD WS (recently renamed AutoCAD 360). The mobile version, available for free, includes features that let users edit and annotate DWG files.

In theory, AutoCAD 360 and a few mobile apps offer enough tools to create and complete a drawing from start to finish. But currently, most mobile device users appear to rely on CAD apps only for fleeting, spontaneous operations such as change approval, minor edits, redlines and annotations.

This may be partly due to the entrenched mouse-and-keyboard computing behaviors that are inseparable from the widely adopted CAD programs, and partly due to engineers' preference (sometimes for justifiable reasons, other times not) for a dedicated workstation and an environment free of distractions to do their primary design work. But in environments where a laptop is a burden and a desktop is nearly impossible to operate — such as on construction projects, survey sites or plant floors — mobile apps and mobile devices are breaking new ground. **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 ➜ Autodesk: Autodesk.com

➜ **Citrix:** Citrix.com

➜ **Corel:** Corel.com

➜ **Dassault Systèmes:** 3DS.com

➜ **GrabCAD:** GrabCAD.com

➜ **Graebert:** Graebert.com

➜ **GstarCAD:** GstarCAD.net

➜ **IMSI/Design:** TurboApps.com

➜ **Lattice Technology:** Lattice3D.com

➜ **Leica Geosystems:** Leica-Geosystems.com

➜ **Microsoft:** Microsoft.com

➜ **NEi Software:** NENastran.com

➜ **RemotePC:** RemotePC.com

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

➜ **SolidWorks:** SolidWorks.com

➜ **ZWCAD Software Co.:** ZWSoft.com

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

Two GREAT Images

The new BenQ BL2710PT and BL3200PT display monitors deliver great images at an amazing price.

BY DAVID COHN

The 27-in. BenQ BL2710 and 32-in. BenQ BL3200 offer feature-rich 2560x1440 resolution displays at an affordable price. ***Photo by David Cohn.***

At Autodesk University in December, we first saw an impressive 27-in. display from BenQ, a company originally spun off from Acer in 2001. While the Taiwanese firm is not particularly well known in the engineering market, it has been selling LCD monitors, digital projectors, digital cameras and mobile computing devices for years. Its lack of visibility among CAD/CAM users will likely change, however, judging from the two monitors we recently received for review.

In addition to the 27-in. BenQ BL2710PT, the company also sent us its new BenQ BL3200PT, the world's first 32-in. WQHD monitor aimed at CAD/CAM and 3D printing. It offers the same 2560x1440 resolution as the BL2710, but on a larger panel.

Both monitors are based on LCD panels manufactured by AU Optronics (AUO). While the BL2710 uses an in-plane switching (IPS) technology that AUO calls advanced hyper-viewing angle (AHVA), the BL3200 uses an advanced multi-domain vertical alignment (AMVA) panel. Both monitors offer wide viewing angles of 178° in both horizontal and vertical directions; their response times are as low as 4 milliseconds gray-to-gray (the time it takes to change a given pixel from gray to a different color and then back to gray). Faster response time reduces image smearing that can occur with moving images. The BL2710 has a .233mm pixel pitch, a brightness of 350 cd/m2, and a contrast ratio of 1000:1, whereas the BL3200 has a larger .276mm pixel pitch, a brightness of 300 cd/m2, and a contrast ratio of 3000:1.

Each monitor arrived neatly packed with its panel and stand wrapped separately. In addition to the power cord, BenQ also provides six cables: D-Sub (VGA), DVI-D, DisplayPort, HDMI and an audio cable to connect the computer's audio port to the monitor's built-in stereo speakers. It took just a few minutes to assemble each display by placing the panel face down on a

BENQ 27- AND 32-INCH MONITOR COMPARISON

	BENQ BL2710PT 27-IN. IPS DISPLAY	BENQ BL3200PT 32-IN. IPS DISPLAY
Price	$699 MSRP ($599 street price)	$799 MSRP ($699 street price)
Size	27-in. (diagonal)	32-in. (diagonal)
Display Type	IPS LED backlit	IPS LED backlit
Screen dimensions without stand (WxHxD)	25.12x15.00x 2.81 in.	29.12x17.12x 2.5 in.
Physical size with stand at highest setting (HxWxD)	21.62x25.12x10.00 in.	29.12x25.63x10.5 in.
Weight	17.7 lbs.	26.9 lbs.
Native Resolution	2560x1440 pixels @ 60 Hz	2560x1440 pixels @ 60 Hz
Horizontal frequency range	30 to 83 kHz	30 to 88 kHz
Vertical refresh rate	50 to 76 Hz	50 to 76 Hz
Aspect Ratio	16:9	16:9
Pixel Pitch	0.2331mm	0.276mm
Dot/Pixel Per Inch	108.79 dpi	91.79
Brightness	350 cd/m2	300 cd/m2
Contrast Ratio	1000:1	3000:1
Response time	4 ms (gray to gray)	4 ms (gray to gray)
Number of Colors	1.07 billion	1.07 billion
Color Gamut	100% sRGB, ~72% NTSC	100% sRGB, ~79% NTSC, 78% Adobe RGB
Power Conusmption	43 watts typical, >0.5 watts standby (43 watts max)	52 watts typical, >0.5 watts standby (97 maximum)
Video input ports	VGA, DVI-D, DisplayPort, HDMI	VGA, DVI-D, DisplayPort, HDMI
I/O Ports	USB 3.0 in, two USB 3.0 out, two USB 2.0 out, audio line-in, headphone	USB 3.0 in, two USB 3.0 out, two USB 2.0 out, one mini-USB out, audio line-in, headphone jack
Other Features	tilt/swivel base, portrait/landscape pivot, built-in 3-watt stereo speakers, Kensington lock slot	tilt/swivel base, portrait/landscape pivot, built-in 5-watt stereo speakers, Kensington lock slot, OSD controller, SD card slot
Cables Included	AC power cord, DVI, DisplayPort, VGA, HDMI, USB 3.0, audio	AC power cord, DVI, DisplayPort, VGA, HDMI, USB 3.0, audio
Warranty	Three years parts and labor	Three years parts and labor

flat surface, attaching the monitor stand to the monitor base and tightening the thumbscrew on the bottom of the base, and then aligning the stand arm with the monitor and locking it into place. Before attaching the monitor to its stand, we noted the standard 100mm VESA mounting holes that enable the panels to be wall-mounted or attached to other supports.

The 27-in. BL2710 weighs 17.7 lbs. (including the stand). The panel itself measures 25.12x15.0x 2.81 in.; it has a height adjustment range of 16.5 to 21.62 in., and needs a space 10 in. deep. The 32-in. BL3200 weighs 26.9 lbs., has a panel measuring 29.12x17.12x 2.5 in., a height adjustment range of 19.5 to 25.63 in., and is 10.5 in. deep.

The stand included with each monitor provides a stable support. It allows the displays to be swiveled 45° left and right, and to be tilted from -5° to +20°. The panels can also be pivoted 90° from landscape to portrait mode. Thanks to the supplied Display Pilot software, the image automatically switches between landscape and portrait orientation when you pivot the monitor. Cables can be neatly routed through a hole in the center of the lower portion of the stand.

Nice Touches Abound

Both displays sport panels with a 16:9 aspect ratio surrounded by a thin black bezel. A power button located in the lower-right corner glows white when the screen is active, and amber in standby mode. Five adjacent touchpad buttons with LED indicators let you access controls using the on-screen display (OSD). The LEDs in these touch-sensitive buttons light up as your finger approaches, the first of many nice touches we noted. The first three buttons on the left can be customized to access specific controls.

Both BenQ monitors also feature a full range of input ports, including dual-link DVI, DisplayPort (with DP 1.2 supported on the BL3200), D-sub (VGA), and HDMI video inputs. Each also includes a USB input port and an audio line-in jack. On the BL2710, all of these connections are located on the bottom of the rear panel — along with an AC power connector, a master power switch, two USB 2.0 ports and a headphone jack. These ports can be difficult to reach when the panel is in its landscape orientation, however. There's also a pair of USB 3.0 ports along the left side of the bezel, and the BL2710 includes a pair of built-in 3watt stereo speakers.

On the BL3200, the connections are a bit more spread out. The USB input port, a mini-USB port (for the OSD Controller), two USB 2.0 ports and an audio line-in jack, as well as the AC power connector and master power switch, are located on the bottom of the rear panel. But the four video inputs are located on the right side of this panel, where they are much more easily accessed. A panel on the lower-right edge of the bezel provides easy access to a pair of USB 3.0 ports, a headphone jack and an SD card reader. Other than the difference in size and port locations, the BL3200's vertical stand is silver aluminum whereas the BL2710's is black, and the built-in stereo speakers in the 32-in. monitor are rated at 5watts.

Both BenQ monitors also feature an ambient light sensor built into the center of the bezel below the panel, enabling the display to automatically adjust its brightness as the surrounding lighting conditions change. There is also a human motion sensor (BenQ refers to this as an "eco sensor") that can turn the screen off when it senses that no one has been in front of it for 40 seconds. The sensor automatically turns the display on again when it senses motion.

The BL3200 features an OSD controller, a small puck-shaped device that connects to the mini-USB port. Buttons on this controller duplicate and augment the bezel-mounted buttons. Although it's a bit of a gadget,

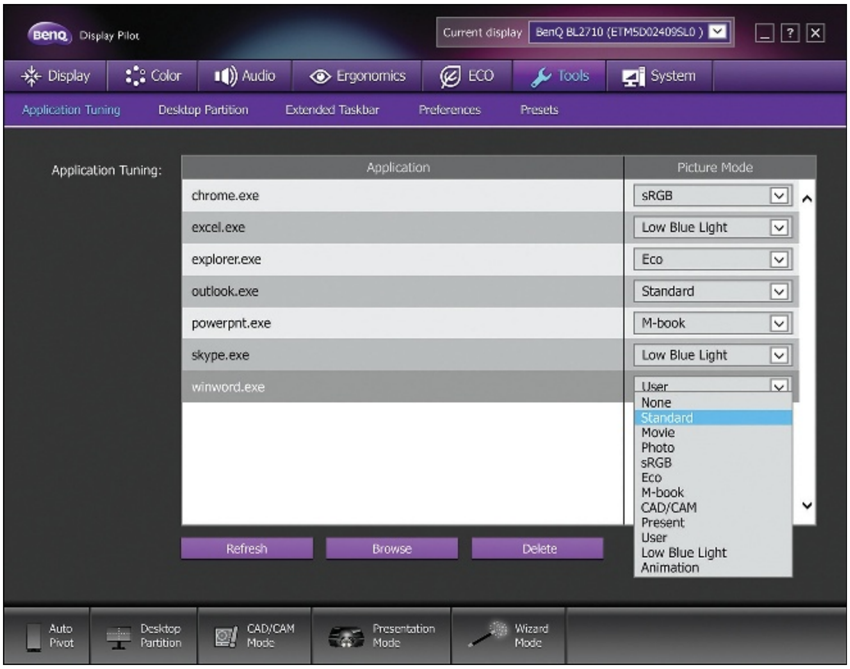

BenQ's Display Pivot software features a picture mode that automatically changes color temperature, brightness, and contrast settings to best match the program you are using. It features a CAD/CAM setting. ***Image courtesy of BenQ.***

we did find it easier to access and navigate the OSD using the controller than using the buttons on the bezel. When not being used, the controller fits into a round indentation in the monitor stand.

Useful Software

The supplied Display Pivot software proved to be one of the most powerful and useful programs we've seen yet for adjusting a display. In addition to automatically detecting when the display is pivoted between landscape and portrait mode and matching the image to the new orientation, it also features a picture mode that, once configured, automatically changes color temperature, brightness, and contrast settings to best match the program you are using.

For example, when using CAD software, the monitor can switch to CAD/CAM mode. Other choices include:

- sRGB (for matching colors to devices such as printers and digital cameras)
- animation
- presentation
- movie
- photo
- reading (for ebooks and documents)
- eco (to lower brightness and conserve power), and
- M-book (to minimize differences between the monitor and a connected MacBook).

You can also use the Display Pilot software to control the ambient light sensor and eco sensor; set a pop-up message to display at pre-defined intervals to remind you to rest your eyes; set a timer to automatically power off the monitor in power-saving mode; span the Windows task bar across multiple monitors; and display the recommended resolution whenever a new input source is detected.

Impressive Performance

We used DisplayMate from DisplayMate Technologies (displaymate.com) to help evaluate the visual quality of each monitor. DisplayMate uses a series of test patterns, both to help users fine-tune the image and to uncover any picture quality problems or video artifacts that might otherwise go unrecognized.

Both BenQ monitors displayed excellent color and gray scale, and showed absolutely no pixel defects. We were able to read text down to 6.8 points, even at different intensity levels, and the fast response time resulted in no image smearing when viewing full-motion video.

We were also impressed with the performance of both of these displays — and with their price. The BL2710 has a suggested retail price of $699, but is currently available both from BenQ and numerous retailers for $599, much less than many other 27-in. IPS displays. The 32-in. BL3200 has a suggested retail price of $799, but it, too, has a street price $100 lower, largely unheard of for a display of this size. Both monitors are backed by a three-year warranty that covers parts and labor. At these prices, we expect these BenQ monitors to start turning up on the desks of many engineers.

Although 27-in. monitors have become an optimum size, and the BenQ BL2710 is a great monitor, remember that both the BL2710 and BL3200 offer identical 2560x1440 WQHD resolution. In other words, if you find that text looks too small on a 27-in. monitor, the big BenQ 32-in. display may be an ideal alternative. **DE**

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

INFO → BenQ: BenQ.us

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

3D Printing's New Materials
Part 2: Metal

18k gold watch-case produced for Hoptroff London on the new PRECIOUS M 080 AM system from EOS and Cooksongold. *Image courtesy of Hoptroff.*

Materials and systems offer ever-increasing choices for one-off parts or hundreds of custom pieces.

BY PAMELA J. WATERMAN

When it comes to making parts out of metal, additive, subtractive and hybrid technologies may all offer a solution that meets your needs. Mindful of the old adage, "what the customer really wants is a hole, not a drill," system manufacturers and supporting businesses continually work to supply a tool that simply gets the end-job done. That need is increasingly being addressed by additive manufacturing (AM), sometimes combined with subtractive (traditional) processes.

AM is gaining acceptance across an ever-widening range of industries. Tim Caffrey, senior consultant at Wohlers Associates, reports that "unit sales of metal AM systems increased 75% from 2012 to 2013, growing faster than the rest of the industry." He notes that while medical applications form a large segment of this market, aerospace is definitely growing, referring to GE Aviation's 2013 acquisition of Morris Technologies with its EOS direct metal laser sintering (DMLS) technology.

"The GE Aviation fuel nozzle project is going forward a little bit slower than we thought, but as far as I know, there is no Plan B," Caffrey says. He adds that Airbus has been fairly visible about working with AM, along with Rolls Royce and Pratt & Whitney. Aerojet Rocketdyne is also working on 3D-printed rocket-engine components.

In this second article of a two-part series on the world of AM materials (see "3D Printing's New Plastics, Ceramics, Composites and More," *Desktop Engineering* October 2014), *DE* reviews developments in metals by longtime players and newer entries in the field, activity with standards organizations, technology resources, and market implications behind these processes.

News and Nuances in Metals for AM

Titanium alloys, nickel alloys, cobalt chrome and several stainless steels are the most common materials used for AM processing, but years of behind-the-scenes development are paying off with even more options. The following companies represent just a few of the metal AM systems; these all work with powder-bed formats, but create parts through different technologies.

In mid-September, EOS announced it has partnered with Cooksongold (a UK-based supplier of fabricated precious metals) to launch the PRECIOUS M 080 DMLS system. Equipped with a 100-watt fiber laser whose small spot-size allows extremely fine resolution, the PRECIOUS M 080 has a build volume of 80mm diameter by 95mm high, including the build platform. The system currently operates with a variety of gold alloys targeted to luxury watches and jewelry; additional materials are planned. Earlier this year, EOS announced a new corrosion-resistant grade of titanium TI64ELI and a highly ductile SS316L good for medical instruments and watch cases, bringing its own material count to 13.

3D Systems never stands still. In September, it acquired LayerWise of Belgium, whose laser-sintering approach adds another direct metal AM process in-house. Rachael Dalton-Taggart, 3D Systems' director of marketing communications, professional products, says, "The LayerWise acquisition brings new technologies, tools and services to our portfolio." Commenting on the company's expanding metal capabilities, she adds, "Our three ProX metals printers are a direct development from the technology acquired from Phenix Systems last year." ProX systems currently support more than 11 materials,

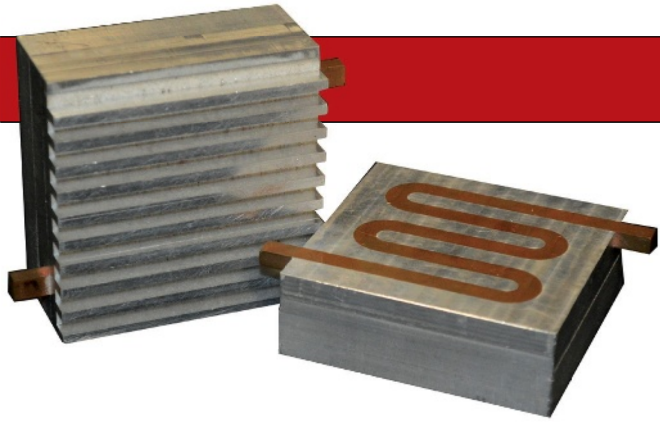

Thermal management device made of copper and aluminum ultrasonically bonded on a Fabrisonic SonicLayer 7200 system. *Image courtesy of Fabrisonic.*

including bronze, aluminum and Inconel 718. Dalton-Taggart says that 3D Systems is testing molybdenum, Inconel 625, pure copper, tungsten and other stainless/tool steels. Using ProX equipment, customers such as Metal Technology Inc. (MTI) are developing additional materials, including C-103, a niobium-based alloy widely used in space applications.

Nickel-based alloys are desirable for their resistance to heat and corrosion. Arcam, with its Electron Beam Melting systems, has added nickel-based Inconel 718 to its material list, joining EOS, 3D Systems, Concept Laser, Renishaw and SLM Solutions, among others, in offering this material for sintering/melting.

Taking a very different approach to producing fully dense AM nickel parts, ExOne, with its jetted-binder-on-powder technology, introduced Inconel Alloy 625 in May; parts display an impressive 99% density. This achievement brings the company a big step toward its goal of single-metal 3D-printed parts (compared to its stainless steel/bronze-infiltrated process).

ExOne continues to put extensive effort into materials development, saying its ExMAL group is on track to qualify a new industrial material every six months. It's also working with the University of Pittsburgh and Magnesium Elektron Powders on magnesium and iron-based alloys for biomed applications.

Materials about Materials

- AMPM2015 (Additive Manufacturing with Powder Metallurgy 2015 Conference, San Diego, May 17-19, 2015; AMPM2015.org) — A Who's Who of worldwide metal AM industry experts headlined last year's inaugural conference, sponsored by the Metal Powder Industries Federation (MPIF.org), a not-for-profit federation of related trade associations that has dealt with powdered metals since 1944. Greg Morris of GE Aviation and Joseph Strauss of HJE Co. kicked off the event, which included 12 sessions with 45 presentations. The 2015 conference is timed to allow attendance at SME RAPID 2015 (Long Beach, CA, May 18-21, 2015).
- An Introduction to Metal Additive Manufacturing / 3D Printing (Metal-AM.com/introduction_to_metal-additive_manufacturing) — Excellent overview of technologies, materials, finishing processes and the market, all in one place. This website launched in May 2014.
- ASTM standards for metals (ASTM F3049) — ASTM International Committee F42 on Additive Manufacturing Technologies is working to create standards for metal powders, especially to help with highly regulated industries such as aerospace and medical. A proposed new standard, ASTM F3049, "Guide for Characterizing Properties of Metal Powders Used for Additive Manufacturing Processes," will point users to existing standards that may be appropriate for AM applications. A companion standard covering the mechanical properties of metal parts, ASTM WK43112, "Guide for Evaluating Mechanical Properties of Materials Made via Additive Manufacturing Processes," is also underway.

— PJW

Targeted Deposition, Gradient Alloys

Fusing a bed of powdered metal represents one approach to metals AM. Another technique, direct-metal-deposition (DMD), involves injecting raw material into a heated melt-pool and either fusing it onto an existing part or free-forming it directly on a build-platform. DM3D Technology (previously POM) and Optomec have employed variations on DMD systems for a decade or more; RPM Innovations and Sciaky are more recent entrants in this field; and NASA's Jet Propulsion Laboratory has its own related research program.

Wohlers Associates' Caffrey says only DMD offers highly desired, multi-material options for metals. Bhaskar Dutta, COO at DM3D Technology, notes that it takes a team of design, material and process engineers to define the best approach to building such parts; his company has extensive experience with customizing blended microstructures — offering dual- and quad-hopper systems for combining materials such as metals and ceramics.

Applying metal AM to part repair is a topic gaining more attention through the activities of America Makes, the National Additive Manufacturing Innovation Institute program. As part of this effort, Optomec will lead the "Re-Born in the USA" project for repair of aerospace metal components for the U.S. Air Force. Optomec's powder-fed DMD technology (based on the LENS process licensed from Sandia National Laboratories) can add metal onto an existing substrate of almost any 3D shape.

"We will demonstrate the benefits of additive manufacturing over traditional welding techniques and enable a 'repair, don't replace' approach to critical part sustainment for high-value aerospace components," says Dr. Richard Grylls, Optomec LENS general manager and Re-Born in the USA project leader.

Inspired by customer requests for large-part volume repair and manufacturing, in 2009 RPM Innovations, a spin-off of RPM & Associates, developed its own industrial-grade Laser Metal Deposition (LMD) equipment that features a 5x5x7-ft. work envelope and multiple powder feeders. Company founder Robert Mudge says that writing their own motion-control software has been the key to running the equipment 24/7; the com-

SonicLayer 7200 Ultrasonic Additive Manufacturing (UAM) system from Fabrisonic. Ultrasonic welding process creates a solid-state bond between layers of similar or dissimilar metals. *Image courtesy of Fabrisonic.*

pany currently operates three such systems, doing jobs that can go for 1,500 hours.

Another company with a growing presence in the DMD world is Sciaky, founded in 1939 as a welding system manufacturer. In 2009, it introduced its Electron Beam Additive Manufacturing (EBAM) process that deposits metal into a melt-pool via wire feedstock. Depending on part geometry and material selected, the Sciaky VX-110 system can deposit near-net layers at a rate of 7 to 20 lbs. per hour in a build envelope of more than 9 ft. each in XYZ. The company has been awarded several U.S. Department of Defense (DoD) and America Makes contracts, working with such partners as Lockheed Martin Aeronautics and the Applied Research Laboratory at Pennsylvania State University.

In a Class of Their Own

AM is indeed an umbrella term for dozens of fabrication technologies. Fabrisonic, based in Columbus, OH, is re-energizing a version of layered AM formerly called ultrasonic consolidation, now termed ultrasonic additive manufacturing (UAM). The company began in 2008 as a joint venture between EWI (an Ohio-based non-profit) and Solidica (now defunct, which had developed the low-temperature, sound-based process). UAM builds objects by ultrasonically bonding layers of metallic foil, then performing precision computer numerically controlled (CNC)-machining on the solid object after each layer. The process forms 100% dense parts and supports welding dissimilar metals — without the formation of brittle intermetallics.

Fabrisonic's CEO, Mark Norfolk, says the company has enjoyed 70% growth since 2012, and offers both part production and machine sales. Its largest system, built for repairing existing parts, can even ultrasonically bond material over a curved surface; the company has been awarded funding to add a rotary axis for welding on cylindrical parts. An interesting option exclusive to the UAM process is that of inserting continuous fibers or stainless steel meshes between layers, creating metal matrix composites that add high strength at low weight for aerospace applications.

A hybrid of another kind is the LUMEX Avance-25, a combination selective-laser-sintering (SLS) and machining system from Matsuura Industries now available in the U.S. The equipment houses an end-mill supplied by a tool magazine that can accommodate 20 milling tools. System operation alternates laser-sintering a succession of 10 layers of 0.05mm-thick powdered metal, then milling the part as needed, producing a finished part with high dimensional accuracy and surface smoothness. Matsuura has targeted customers looking to reduce production time of complicated mold dies.

German manufacturer DMG Mori is marketing its own version of a hybrid AM system: The Lasertec 65 combines laser DMD technology with a five-axis milling machine. The system features a build volume of 23.6 x 15.7 in.; wall thicknesses of 0.1mm to 5mm are possible.

Two more deposition-type hybrid systems are now coming into production. Based on a four-year UK research project, Hybrid Manufacturing Technologies announced its AMBIT Multi-Task system, a docking system of robotically controlled tools. Supported functions include laser cladding, five-axis machining, touch-probe precision measurements, polishing, annealing and even cleaning, all in a single set-up.

Optomec has also introduced a hybrid component called the LENS Print Engine, which adds laser deposition capabilities to CNC machines. Working with new or existing tooling systems, the LENS Print Engine includes proprietary powder feeders, deposition heads, LENS process controls, fiber laser support, integrated tool-path generation software, and safety packages. The company's first announced customer is the Center for Remanufacturing and Resources Recovery (C3R) at the Rochester Institute of Technology (RIT).

Third-party Materials

Materials development by independent suppliers is expanding as more AM systems work with non-proprietary powders. Metalysis of Rotherham, UK, working with engineers at Sheffield University's Mercury Centre, has developed a new way to produce titanium powder that is cheaper compared to the usual multi-step, energy-intensive Kroll process. It extracts titanium from rutile sand (naturally occurring titanium ore present in beach sands) in a single electrolysis step. Dion Vaughan, CEO of Metalysis, says that the process could "reduce the price of titanium by as much as 75%." The company is also developing tantalum powder.

Last year Materialise, whose software solutions, consulting services and production facilities support the full breadth of the AM world, opened a Metal Competence Research Centre. Abbey Delaney, Materialise North American marketing manager, says the Centre was established to offer surgeons a wider range of products and services. They are working with pure titanium and titanium grade 23 materials, especially to create porous structures.

Thinking of developing your own materials, or at least doing R&D on the latest equipment? Linear Mold & Engineering is currently expanding its facility space and AM machine capacity.

"We see ourselves as a solution provider, offering not just part production but training sessions, consulting and, coming next year, leased 'pods' for private customer R&D," says Director Paul Parzuchowski. He says Linear Mold & Engineering will not only provide dedicated or time-shared equipment, but also maintenance and training. The goal is to help its customers develop new materials, techniques and products using high-end AM systems.

From Professional to Personal and Back

No discussion of metal AM news would be complete without mentioning the start-ups tackling desktop metal systems. Vader Systems, matterFab and Newton3D have all emerged this year as contenders in the (previously non-existent) personal metal-AM system space, and warrant continued attention.

Among the big-name players, AM developments happen daily. GE's Avio Aero division recently developed a new, higher-power electron-beam gun for Arcam equipment that in turn will allow use of lightweight titanium aluminide for turbine-blade production. JPL is working on depositing layers of metal on a rotating rod, creating parts with gradient metal alloys. And the buzz from the 3DPrint Show London 2014 is to watch for Stratasys to directly enter the metals market, rather than rely on third-party equipment. **DE**

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

INFO → **3D Systems:** 3DSystems.com

→ **Arcam:** Arcam.com

→ **Concept Laser:** Concept-Laser.de

→ **DM3D Technology:** DM3DTech.com

→ **DMG Mori:** DMGMori.com

→ **EOS:** EOS.info

→ **EWI:** EWI.org

→ **ExOne:** ExOne.com

→ **Fabrisonic:** Fabrisonic.com

→ **GE Avio Aero:** AvioAero.com/en

→ **Hybrid Manufacturing Technologies:** HybridManuTech.com

→ **Linear Mold & Engineering:** LinearMold.com

→ **Lockheed Martin Aeronautics:** LockheedMartin.com/us/aeronautics.html

→ **Magnesium Elektron Powders:** Magnesium-Elektron.com

→ **Materialise:** Materialise.com

→ **MatterFab:** MatterFab.com

→ **Matsuura Industries:** Matsuura.co.jp/english

→ **Metal Technology Inc.:** MTIAlbany.com

→ **Metalysis:** Metalysis.com

→ **NASA Jet Propulsion Laboratory:** JPL.NASA.gov

→ **National Additive Manufacturing Innovation Institute:** AmericaMakes.us

→ **Newton 3D:** Newton3D.com

→ **Optomec:** Optomec.com

→ **Pennsylvania State University, Applied Research Laboratory:** ARL.PSU.edu

→ **Renishaw:** Renishaw.com

→ **Rochester Institute of Technology, Center for Remanufacturing and Resources Recovery:** RIT.edu/gis/remanufacturing

→ **RPM Innovations:** RPM-Innovations.com

→ **Sciaky:** Sciaky.com

→ **Sheffield University, Mercury Centre:** MercuryCentre.org

→ **SLM Solutions:** SLM-Solutions.com

→ **Stratasys:** Stratasys.com

→ **Vader Systems:** VaderSystems.com

→ **Wohlers Associates:** WohlersAssociates.com

FEA and Composites PART 2

A look at some background theory to laminate stiffness calculations and review some of the more modern approaches to composite failure analysis.

BY TONY ABBEY

Editor's Note: Tony Abbey teaches live NAFEMS FEA classes in the US, Europe and Asia. He also teaches NAFEMS e-learning classes globally. Contact tony.abbey@nafems.org for details.

The calculation of the stiffness of a laminate or layup is an important step in evaluating a composite finite element analysis (FEA) element. Unlike an isotropic shell element, the stiffness depends on the number and stacking sequence of the individual plies. Each ply can be a varying thickness, angle and material property. The laminate stiffness is also important as a predictor of the response of the composite. Let's look at how the laminate stiffness is calculated.

The fundamental building block in a laminate is the ply. For a thin shell calculation, we assume a 2D orthotropic material can represent the ply. The important directions are with-fiber (direction 1) and transverse to fiber (direction 2). Consider a single ply loaded in the with-fiber direction (1) and free to contract in the transverse direction (2) so that transverse stress is zero, but there is a strain because of Poisson's effect. This is shown in Fig. 1A.

Now consider the ply loaded in the transverse direction (2) and free to contract so that the with-fiber stress (1) is zero, but again there is a strain due to Poisson's effect. This is shown in Fig. 1B. The equations show that two separate Young's moduli (E1 and E2) and two Poisson's ratios are used in the coupling.

Now consider both stresses applied simultaneously. The stresses and strains are shown in Fig. 2. The strain in each direction is coupled between with (1) and transverse (2) terms.

The stress-strain relationship in shear is shown in Fig. 3.

If all the stress-strain relationships are combined, the compliance matrix is produced as shown in Fig. 4. Note the relationship between Poisson's ratios and Young's moduli. The compliance matrix relates strains to stresses.

To develop a stress-strain relationship, we need to invert this to form a "reduced" stiffness matrix. This is shown in Fig. 5.

The act of transforming the ply to an off-axis angle causes a coupling between the shear and direct stress terms, and creates a full ply stiffness matrix. This is shown in Fig. 6, where ply k has been rotated theta degrees. I won't show the transformation terms, but they are just functions of angle theta.

What is the Matrix?

So, how do we get a laminate stiffness matrix? The final step is to take each 3x3 ply stiffness matrix, factor up by its thickness,

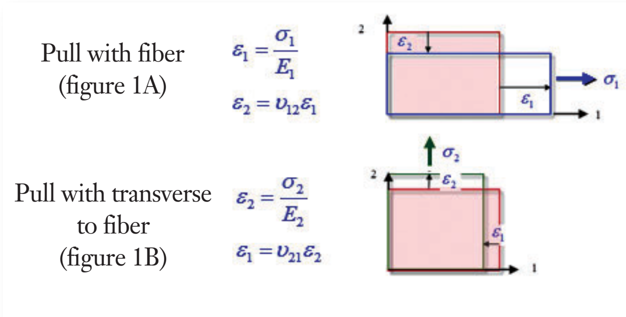

FIG. 1A AND 1B: Loading in either with (1) or transverse to (2), fiber direction.

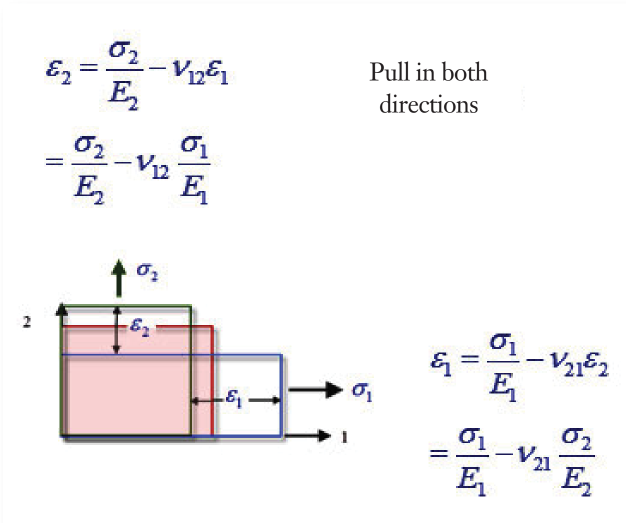

FIG. 2: Loading in both directions at the same time.

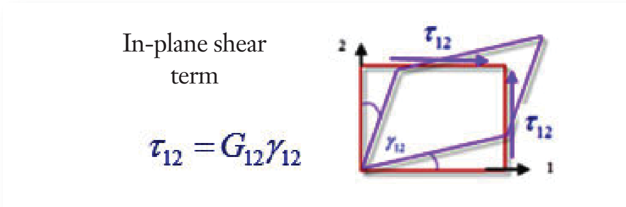

FIG. 3: In-plane shear term.

Compliance Matrix

$$\{\varepsilon\}=[S]\{\sigma\}$$

$$\begin{Bmatrix}\varepsilon_1\\ \varepsilon_2\\ \gamma_{12}\end{Bmatrix}=\begin{bmatrix}1/E_1 & -\nu_{12}/E_1 & 0\\ -\nu_{12}/E_1 & 1/E_2 & 0\\ 0 & 0 & 1/G_{12}\end{bmatrix}\begin{Bmatrix}\sigma_1\\ \sigma_2\\ \tau_{12}\end{Bmatrix}$$

Note it is symmetric as $\nu_{21}/E_2=\nu_{12}/E_1$

FIG. 4: Compliance matrix definition.

Reduced Stiffness Matrix

$$\{\sigma\}=[Q]\{\varepsilon\}\qquad \begin{Bmatrix}\sigma_1\\ \sigma_2\\ \tau_{12}\end{Bmatrix}=\begin{bmatrix}Q_{11} & Q_{21} & 0\\ Q_{12} & Q_{22} & 0\\ 0 & 0 & Q_{66}\end{bmatrix}\begin{Bmatrix}\varepsilon_1\\ \varepsilon_2\\ \gamma_{12}\end{Bmatrix}$$

$$Q_{11}=\frac{E_1}{(1-\nu_{12}\nu_{21})}\qquad Q_{66}=G_{12}$$

$$Q_{22}=\frac{E_2}{(1-\nu_{12}\nu_{21})}$$

$$Q_{12}=\frac{\nu_{21}E_1}{(1-\nu_{12}\nu_{21})}=\frac{\nu_{12}E_2}{(1-\nu_{12}\nu_{21})}=Q_{21}$$

Why Q_{66}?

The terms 3,4 and 5 are reserved for a full 3d orthotropic material definition (one direct stress in direction 3 and two interlaminar shears in 13 and 23)

FIG. 5: Ply reduced-stiffness matrix.

Ply Stiffness Matrix

$$\{\sigma\}_k=[\bar{Q}]_k\{\varepsilon\}_k\qquad \begin{Bmatrix}\sigma_x\\ \sigma_y\\ \tau_{xy}\end{Bmatrix}_k=\begin{bmatrix}\bar{Q}_{11} & \bar{Q}_{21} & \bar{Q}_{16}\\ \bar{Q}_{12} & \bar{Q}_{22} & \bar{Q}_{26}\\ \bar{Q}_{16} & \bar{Q}_{26} & \bar{Q}_{66}\end{bmatrix}_k\begin{Bmatrix}\varepsilon_x\\ \varepsilon_y\\ \gamma_{xy}\end{Bmatrix}_k$$

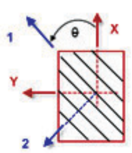

FIG. 6: Ply stiffness matrix.

and then add all the stiffness terms up. This provides a 3x3 in–plane stiffness matrix that relates applied load (in force/length) to resultant strains. This is called the A matrix. The influence of the A matrix can be seen by considering laminates that are balanced and unbalanced.

A balanced laminate has "balanced" ply angle contributions. For example a [5/-5/30/-30/-60/60] laminate has a complete set of +angle/-angle pairs. Note that 0/90 is also a balanced pair. The A matrix is shown in Fig. 7 (lower) and has no coupling between direct and shear stresses as shown by the highlighted zero terms. Even if [5/-5/30/-30/-60/60/60] is used, but the thickness of the single -60 ply is doubled, the laminate is still balanced. Fig. 7 (upper) shows an unbalanced A matrix resulting from [5/-5/30/-35/-60/60].

The physical implication of the A matrix is important; coupling between direct and shear terms means that any direct pull will also result in a shearing action. Any shear will result in a direct stress term. This can cause unwanted stresses in a loaded component — or warped components if thermal curing is used.

The A matrix is also important for calculating equivalent in-plane stiffnesses for hand calculation for rough design estimates, or checking results.

The out-of-plane stiffness matrix D can be calculated in a similar way, starting with each ply stiffness matrix. However, the stacking sequence plays an important role in determining the values. For example, a [0/90/s] laminate is compared to a [90/0/s] laminate. The resulting D matrices are shown in Fig. 8 — and clearly the 0° plies positioned furthest away from the centerline give the biggest stiffness. The D terms are used to check bending and buckling stiffness. Note the coupling terms are computed zeroes. This indicates there will be no coupling between pure bending and twisting. Estimation of this can be useful for either avoiding or promoting bending/twisting coupling in sheets; it is achieved by varying the balanced plies.

The B matrix is the final matrix that is calculated. It indicates the coupling between in-plane loading and bending response, and is probably the most dominant term in design. If a pre-preg sheet is to remain flat after curing in an autoclave, the B matrix must be forced to decouple by only allowing symmetric plies.

Fig. 9 shows a range of possibilities with balanced and symmetric plies. The full ABD matrix array is shown.

Failure Prediction

Progressive ply failure assumes the stiffness of the failed ply can be reduced in some manner, and the analysis continued. Ply failures can continue to occur with subsequent reduction in stiffness.

A progressive ply failure strategy requires:

- a failure criteria that can identify the mode of failure (Hashin, Puck, LaRC02, etc.); and
- a logical strategy for reducing element stiffness based on the mode of failure seen.

Progressive Ply Failure is sensitive to how the stiffness is reduced at each non-linear load step. If the drop is too great, instabilities can occur. Thus, a maximum percentage of stiffness in a particular orthotropic direction is typically used. The user can often elect to modify this to simulate more ductile composite materials. Certain failure modes, such as longitudinal tension, may be classified as final failure in their own right.

The Tsai-Wu failure theory described in Part I (*DE*, October 2014) is just one of many that were developed using known failure points and then interpolating in stress space using quadratic relationships. One of the limitations of this approach is that, although it is easy to program, all failure is based on a full and continuous interaction between stress states. It has been found experimentally that failure modes tend to be dominated by either fiber failure modes or matrix failure. There may be little interaction between them.

A class of failure theories has evolved that is sometimes described as "phenomenological," to indicate the nature of the failure is implicit in the theory. These fit well with the progressive ply failure methods, as they indicate the type of strategy that should be followed in reducing stiffness caused by damage.

Transverse compression is interesting, because the strength is generally higher than transverse tension. The matrix tends to act to stabilize the fibers until some form of shear cracking occurs. If in-plane shear or longitudinal compression is present, this will also affect the way the ply behaves. This behavior is not well understood in general, and is the subject of much manipulation of the failure theories. Hashin-Rotem failure criteria breaks up the assessment of failure into several sub-criterion, any one of which can give a critical failure index:

- tensile fiber failure
- compressive fiber failure
- tensile matrix failure
- compressive matrix failure
- through thickness failure

Fig. 10 shows the separate equations used in Hashin-Rotem, compared to Tsai-Wu.

Other Failure Models: Delamination

Two other FEA methods predict failure at the macro level:

- **Cohesive element methods** aim to model specific de-bonding or delamination situations by inserting a layer of special elements between the plies or materials. The behavior of the crack or delamination front is controlled by an energy rate law to allow tuning for different types material (such as brittle or ductile). The actual damage method is still using a stress-based approach, relating stress state to resultant strain across the delamination. Optionally, separation or erosion of the elements can occur. Implementations vary from continuous 2D or 3D elements forming the cohesive layer, to simple springs exhibiting the stress strain behavior.
- An alternative to using stress-based delamination methods is to use fracture mechanics, postulating a crack at the ply boundary. The energy required to create new crack length or delamination zone forms a strain energy release rate. The force and displacements are known in the FEA model of an open crack, adjacent to the crack tip. Thus the energy required to oppose this and close the crack is also known. This is equal to the energy required to produce the delamination or crack. The method is commonly known as the **Virtual Crack Closure Technique (VCCT)**. The big advantage of the method is that the near singular stresses at the crack tip do not have to be evaluated.

The rate of change of energy with respect to the crack growth rate is analogous to the stress intensity factor in isotropic materials. The strain energy release rate can be compared to the fracture toughness of the material to establish whether a crack will propagate. The original implementations required a regular-sized mesh; advances now allow an arbitrary mesh.

In both methods, the crack will follow the mesh shape no matter what, so take care to avoid "forcing" the crack. A very fine mesh with no directional bias is best.

Micromechanics

In micromechanics, the composite material is assumed heterogeneous, rather than homogeneous as in macromechanics. Fiber and matrix are treated distinctly. The general terminology for "fiber" is broadened to "inclusion," to allow for a variety of reinforcing media such as lamella (long fiber), ellipsoidal (short fiber) and spherical (beads, etc.). The inclusion is dispersed into

A Matix forms

7.31E+07	1.68E+07	1.50E+05
1.68E+07	3.65E+07	-8.93E+05
1.50E+05	-8.93E+05	2.08E+07

unbalanced

8.14E+07	2.39E+07	0.00E+00
2.39E+07	6.19E+07	0.00E+00
0.00E+00	0.00E+00	2.91E+07

balanced

FIG. 7: Balanced and unbalanced in-plane laminate stiffness A matrix.

Compliance Matrix

$$\{\varepsilon\}=[S]\{\sigma\}$$

$$\begin{Bmatrix}\varepsilon_1\\ \varepsilon_2\\ \gamma_{12}\end{Bmatrix}=\begin{bmatrix}1/E_1 & -\nu_{12}/E_1 & 0\\ -\nu_{12}/E_1 & 1/E_2 & 0\\ 0 & 0 & 1/G_{12}\end{bmatrix}\begin{Bmatrix}\sigma_1\\ \sigma_2\\ \tau_{12}\end{Bmatrix}$$

Note it is symmetric as $\nu_{21}/E_2 = \nu_{12}/E_1$

FIG. 8: High and low stiffness bending laminate stiffness terms in D matrix.

AB D Matrix

4.65E+07	1.28E+06	9.10E-11	0.00E+00	0.00E+00	0.00E+00
1.28E+06	4.65E+07	2.38E-09	0.00E+00	0.00E+00	0.00E+00
9.10E-11	2.38E-09	3.90E+06	0.00E+00	0.00E+00	0.00E+00
0.00E+00	0.00E+00	0.00E+00	5.02E-01	3.98E-02	4.93E-18
0.00E+00	0.00E+00	0.00E+00	3.98E-02	2.38E+00	1.29E-16
0.00E+00	0.00E+00	0.00E+00	4.93E-18	1.29E-16	1.21E-01

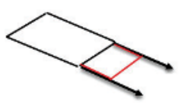

[0/90/90/0] - balanced and symmetric

AB D Matrix

4.65E+07	1.28E+06	9.10E-11	-3.08E+03	0.00E+00	6.94E-15
1.28E+06	4.65E+07	2.38E-09	0.00E+00	3.08E+03	1.81E-13
9.10E-11	2.38E-09	3.90E+06	6.94E-15	1.81E-13	0.00E+00
-3.08E+03	0.00E+00	6.94E-15	1.44E+00	3.98E-02	2.82E-18
0.00E+00	3.08E+03	1.81E-13	3.98E-02	1.44E+00	7.37E-17
6.94E-15	1.81E-13	0.00E+00	2.82E-18	7.37E-17	1.21E-01

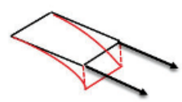

[0/90/0/90] balanced, not symmetric

AB D Matrix

4.66E+07	3.36E+06	2.62E+06	-3.13E+03	-1.36E+02	5.59E+01
3.36E+06	4.22E+07	5.62E+06	-1.36E+02	3.40E+03	-4.17E+02
2.62E+06	5.62E+06	5.97E+06	5.59E+01	-4.17E+02	-1.36E+02
-3.13E+03	-1.36E+02	5.59E+01	1.44E+00	5.58E-02	2.03E-02
-1.36E+02	3.40E+03	-4.17E+02	5.58E-02	1.41E+00	4.35E-02
5.59E+01	-4.17E+02	-1.36E+02	2.03E-02	4.35E-02	1.37E-01

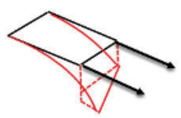

[0/70/5/90] neither balanced or symmetric

FIG. 9: Variations on balanced and symmetric plies in a laminate under axial loading.

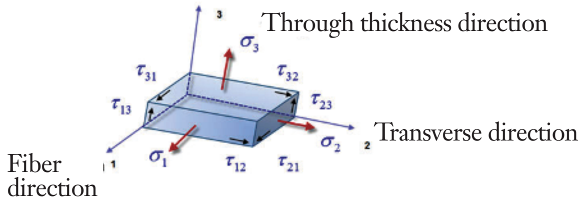

Tensile fiber failure: $\left(\frac{\sigma_1}{X_t}\right)^2+\left(\frac{\tau_{12}}{S_{12}}\right)^2=1$

Compressive fiber failure: $\left(\frac{\sigma_1}{X_c}\right)^2=1$

Tensile matrix failure: $\left(\frac{\sigma_2}{Y_t}\right)^2+\left(\frac{\tau_{12}}{S_{12}}\right)^2=1$

Compressive matrix failure: $\left(\frac{\sigma_2}{2S_{23}}\right)^2+\left[\left(\frac{Y_c}{2S_{23}}\right)^2-1\right]\frac{\sigma_2}{Y_c}+\left(\frac{\tau_{12}}{S_{12}}\right)^2=1$

Through thickness failure $\left(\frac{\sigma_3}{Z_t}\right)^2+\left(\frac{\tau_{23}}{S_{23}}\right)^2+\left(\frac{\tau_{31}}{S_{31}}\right)^2=1$

FIG. 10: Hashin-Rotem failure terms.

a continuous "phase' (the matrix). The microanalysis methodology includes classical solutions and FEA solutions.

The main objective is overall design and analysis of components and structures, using a more versatile range of micro-level stiffness and failure modeling techniques. Typical micromechanics methods assume a representative volume element (RVE) characterizes the extremely local effects. These can be integrated into a macro-level response to provide a form of progressive ply failure model.

Macro level methods for predicting undamaged stiffness for thin plates and shells are well established. Strength prediction is still a challenge in many stress states, but it does seem that the more modern methods are more able to adapt to known experimental behavior.

The extreme example of this is in the relatively new usage in FEA of micromechanics methods. However, tight correlation with known experimental evidence is required, as more reliance is placed on advanced failure methods.

The decision on what level of technology to use should be based on a sensible assessment of the degree of safety margin required, the available test correlation evidence and proven experience with a particular method. For preliminary sizing, the use of simple macro-strain-based criteria is still probably most effective. DE

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

Noise Neutralization

In July, in the heart of the U.S. automotive industry, global automotive supplier Faurecia opened its new North American headquarters. The 278,000-sq.-ft. facility, located about 30 miles north of Detroit in Auburn Hills, MI, encompasses a high-tech research and design center featuring an anechoic test chamber designed and installed by Eckel Industries of Cambridge, MA.

Chances are consumers don't give a thought to the noise and vibration emitted by the motor controls of the seats in their cars, unless they're loud or lurching. But automotive manufacturers invest in testing these parameters to ensure consumer comfort and meet industry standards.

Faurecia wanted to take its automotive seating noise, vibration and harshness (NVH) testing to the next level to achieve the highest industry standards. Because a large share of its global sales belong to premium automotive manufacturers, subjective "psycho-acoustic" benchmarks are crucial.

Having worked with Eckel in the past, Faurecia went directly to the Eckel team to design and install a hemi-anechoic test chamber for the Automotive Seating R&D Center in its new headquarters. It is the largest chamber in any of Faurecia's worldwide locations, and can accommodate an entire vehicle.

Eckel's team of designers and engineers worked closely with Faurecia's building contractor on the requirements for the room within which the chamber would be installed.

The design process began in June 2013. In February 2014, Eckel got the call to head to Michigan to install the chamber, which has been operational since March. Faurecia began conducting testing in April as it transitioned employees and operations to the new headquarters.

"Eckel's anechoic acoustic chamber design and installation teams were efficient and easy to work with," reports Christopher Kus, project engineer and NVH specialist for Faurecia. "They were able to quickly supply quotes during each altered revision of the sound chamber layout. Installation timing was exactly as defined, and the finalized acoustic chamber has world-class quality."

Faurecia Engineering Director Robert Parmann agrees. "Eckel worked very well with us and with our general contractor," he adds. "Their exceptional planning, organization and collaboration helped to keep our total project on schedule. The final test chamber was completed on time, and exceeded the ISO specifications and our expectations."

MORE → deskeng.com/de/?p=20076

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

The Sounds of Success

General Motors uses LMS Soundbrush to "visualize" sound and improve design.

BY MARK CLARKSON

Anyone who's ever spent a weekend trying to track down that annoying squeak in the family automobile knows just how difficult a task that can be. Unless you're a bat or a dolphin, you can't see sound.

LMS Soundbrush is a handheld scanning device incorporating a light, a camera and a 3D sound intensity sensor. You manually scan it over the object on which you want data, and the LMS software presents you with a color-coded 3D visualization of the sound field around that object — where sound is coming from, how intense it is, and the direction it's taking as it emanates from the source.

"It gives you a nice plot," says Margaret Behm, Engineering Group Manager for GM's Noise and Vibration Center. "We call it a 'sardine chart' because it looks like little fishes. It shows the direction of the sound and noise. You can plot it on top of a 3D model of your test object. You can manipulate the 3D views, or cut sections to determine the planar sound field wherever you're interested in it."

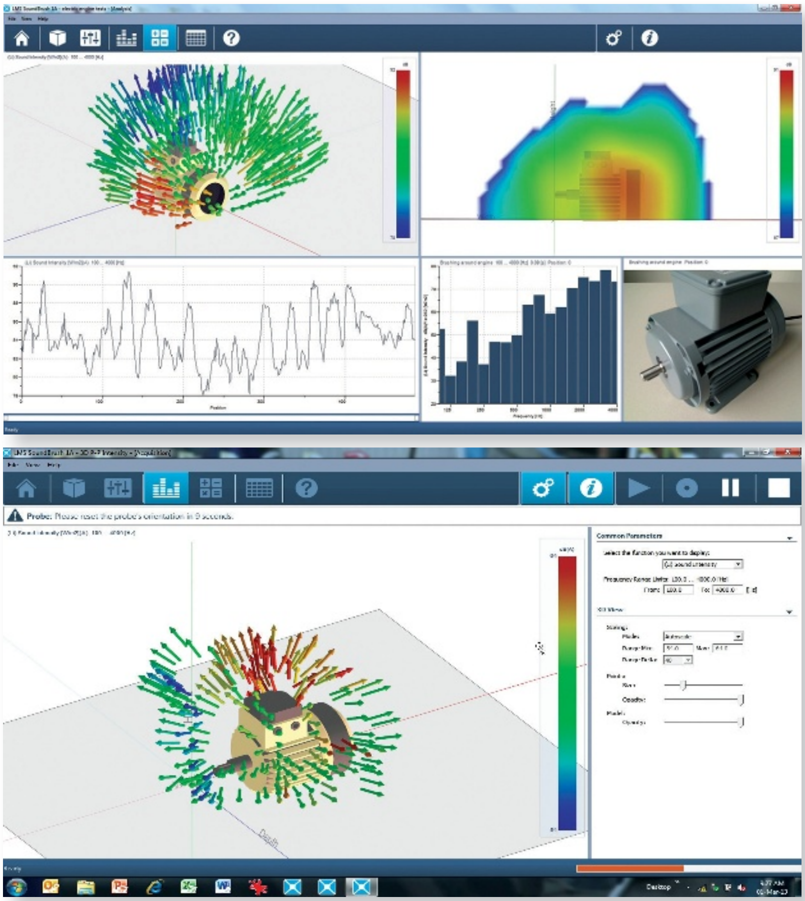

Examples of LMS Soundbrush software being used to visualize the sound coming from an electric engine. *Images courtesy of Siemens PLM Software.*

It's Like Wii Tennis for Unwanted Noise

To pull this off, the SoundBrush has to know where it's at in 3D space at all times. To that end, the system uses a simple motion tracking system incorporating inertial sensors and two cameras: one on the probe itself, pointed at the sound source, and one on a tripod pointed at the handheld probe.

The tripod-mounted camera uses the apparent size and position of an illuminated "ping pong ball" on the probe to determine the probe's distance and relative location, while inertial sensors in the probe track its orientation. Both plug into a PC or laptop, to provide real-time visualization of the sound in your environment.

"It's like [Nintendo's] Wii technology," says Behm. "The system knows where you are. The tracking camera knows exactly where the probe is in space, and the probe knows where it is in space relative to the camera and the object you're scanning."

As you're scanning the object, you actually see the sound appear on your computer screen in real time, in 3D. GM uses the SoundBrush to identify the sources of unwanted noise, as well as the path they travel to the driver and passengers. Once the noise is tracked down, GM uses sound dampening materials or — depending on the vehicle's stage of development — design changes to reduce or eliminate it.

The Time is Right

"A few years back, a bunch of companies came out with 3D acoustic cameras," says Behm. "I looked at the one that was the most popular at the time. It was about $250,000."

That, Behm decided, was just too much money. A 3D acoustic camera wasn't going to help GM's engineers that much, considering that there are plenty of much cheaper ways to locate the source and intensity of noise. Stethoscopes, for instance, cost a few hundred dollars at most.

"I consider the LMS Soundbrush tool a cheap version of that tool," says Behm. "It's faster to set up, it's cheap, and it does the same thing with less complication. So, five years later, when I saw the Soundbrush tool come out at a reasonable cost, I could justify buying it, because it would help my engineers to communicate the issues that they found, and that would help them get the changes that they needed."

Behm says users can show a picture to other engineers, suppliers or management, and communicate quickly what the problem is, where it is on the part, and what direction the sound is coming from to get necessary part changes. They can also use it as evidence to back up a business case to fund a necessary change.

"For instance, we just used it on a truck," says Behm. "We were having sound entering the cabin, and we wanted to determine where the sound was coming from. We scanned the back panel, and not only could we find the path that the noise was entering, but could quantify the intensity of the sound and visually show that on a color map, to communicate that issue to the people who had to make changes for us."

Mobile Data via SCADAS XS

LMS' SCADAS XS is a 12-channel data acquisition system which, at 170x114x23mm, is about the size of an external hard drive. Battery-powered and portable, it's intended as a grab-and-go replacement for a big, bulky multi-channel system.

"You don't have to pick up the stuff, put it in a cart and push it places," says GM's Margaret Behm. "You just put it in your pocket and go. As a matter of fact, it comes with a little carrying case."

Users can grab the XS, a few mics or vibration sensors, and jump in a vehicle with minimal fuss.

"[Engineers] want to be able to go out in a vehicle and assess the vehicle at that moment in time," says Siemens' Ian McGann. "What is the noise level? What is the vibration level?"

To better capture the experience, the XS comes with a headset with two built-in microphones to take binaural recordings. Later on, at your desk, you can play the sound back and it will sound exactly as it did when you were driving the vehicle.

"The tool has GPS," says McGann. "If you hear any unwanted noise, you can find out where you were on the track when you heard that noise. Is there anything specific about the track there? Was I on the Belgian cobblestones when it happened? What frequency range was this noise? Was it at 3,000Hz? 1,000Hz?"

Sometimes, the noise's frequency is sufficient to determine the source. Maybe that squeal matches the drive shaft's rotational speed, or the 10Khz tonal noise corresponds to the air conditioning fan.

"If you can't determine it using simple mathematics based on the speed," says McGann, "then one next step may be to use SoundBrush to see where that noise is actually coming from."

Inside Vs. Outside

The SoundBrush's two-camera motion tracking can become awkward when working inside of a vehicle. If the outside camera can't see the probe, the system won't work.

"This tool is much more conducive to doing stuff outside the vehicle," says Behm. "There was one report that said once you get in the interior, it's best to take out the seats. But once you're taking out the seats, you're wasting a lot of time, and the value goes down."

But that doesn't make it useless: GM has found uses for SoundBrush that LMS never foresaw.

"We write a lot of [Statements of Requirements] for acoustic materials such as dash mats and headliners," says Behm. Vendors deliver hardware that they insist meets those requirements, GM builds the vehicle, tests it and … noise.

"If we assume it's assembled correctly, and we assume it meets spec, then the engineers start chasing design changes," she adds.

But what was it Mom used to say about making assumptions? "Sometimes," says Behm, "we dig a little deeper and find out that what we thought was in spec wasn't in spec and, if [the vendor had] met our specifications, we wouldn't be chasing our tails on design changes. SoundBrush allows [engineers] to do a quick scan of, say, a dash mat, to make sure that we're dealing with parts that are in spec, before we begin assembling the vehicle."

If the parts don't meet spec, despite the supplier's claims to the contrary, GM can prove it more easily. Of course, you can't use the SoundBrush on a truck rolling over a test track, but you can scan, say, a rotating drive shaft.

"We were doing a test on a prop shaft, trying to compare the intensity scan to the CAE mode shape," Behm says. "So, real quick and dirty, they set up SoundBrush and did a scan of a prop shaft that was rotating at a constant 3,250 rpm. It matched up perfectly with the CAE mode shape, so that checked off another variable of concern: that your CAE model is telling you the right story.

"After that," she concludes, "there was no further conversation questioning the model." DE

Contributing Editor **Mark Clarkson** *is* Desktop Engineering's *expert in visualization, computer animation, and graphics. His newest book is Photoshop Elements by Example. Visit him at MarkClarkson.com or send e-mail about this article to DE-Editors@deskeng.com.*

INFO ➜ **General Motors:** GM.com

➜ **Siemens PLM Software:** plm.automation.siemens.com/en_us/products/lms/testing/soundbrush

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

A Smarter Dummy

WorldSID makes history as the first international regulation anthropomorphic manikin for crash testing.

BY MIKE BECKAGE

The World Side Impact Dummy (WorldSID) crash test dummy represents a major step in unifying regulations on a global level for the automotive crash test industry by allowing a single, universally accepted test device to be used for side impact testing around the world. The combined efforts of one ISO task group, 10 countries, 45 organizations, 1,000 tests and about $16 million have created the world's most biofidelic side-impact crash test dummy. The WorldSID 50th (50th percentile, or average male) and 5th (small female) have far-reaching ramifications not only for car manufacturers and test engineers to produce safer vehicle designs, but for the ultimate goal of reducing injuries in side-impact crashes.

Developing the WorldSID has been a slow, uphill road that has steadily driven the automotive crash test industry forward to setting universal test standards. In the early 1980s, the US developed a side impact dummy dubbed USSID. Following that, Europe designed its own version called EuroSID, and then a second generation, EuroSID2. The challenge for car manufacturers selling globally was that they had to go through two different, time-consuming and costly testing processes — one using the USSID and another with the EuroSID.

Car companies and the test world came together to unify efforts, and in 1997, an official ISO WorldSID Task Group was established to create the first SID that would be accepted for use in impact testing for any regulation around the world.

The first production WorldSID 50th system was completed in March 2003, with the first WorldSID 5th following a few years later. It has since undergone years of continuous development, testing and refinement. WorldSID is expected to go into regulation in the next one to two years, and will be the first dummy designed for worldwide regulation testing and with in-dummy DAS standard.

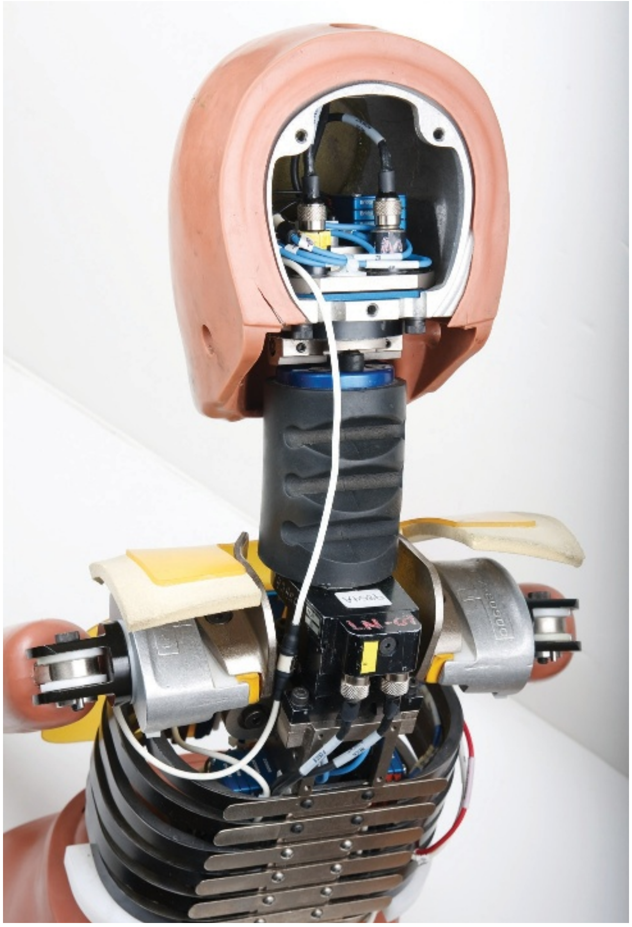

DTS offers three in-dummy DAS options: the new 36-channel SLICE G5 DAS, the original 32-channel TDAS G5 DAS, and SLICE NANO 3-channel blocks, shown in a child dummy with 9 channels in the head, 21 channels in the torso and 3 channels in the pelvis. ***Image courtesy of DTS.***

SID History

Crash test dummies have been around for a long time, with the very first crash test dummy created more than a century ago in Paris. However, mainstream crash test dummies really evolved in 1949 when "Sierra Sam," created by Sierra Engineering, became the first modern production anthropomorphic test device (ATD) developed for aircraft ejection seat testing.

Vehicle crash test dummies were the next step, and it was agreed that a single test dummy was not enough. This led to the development of different types and sizes of dummies, as well as different types of tests. As a result, frontal crash tests require a specific type of dummy, the Hybrid III; for side impacts, there is the new WorldSID; and a variety of component parts such as legforms and headforms are launched into the front bumper and hood of a car to simulate pedestrian injuries. The family of dummies includes average to large males, females, children and even dummies simulating pregnant females. A single dummy may take anywhere from 15 to 20 years to develop.

Originally cumbersome and limited to a single test measurement, crash test dummies have evolved into sophisticated test equipment capable of repeatability and true biofidelity. Today's dummy is instrumented with an average of 64 to100 channels; WorldSID is designed for up to 160 channels. A full-scale crash test may involve up to 500 instrumented channels, including vehicle, dummies and test fixtures.

The original WorldSID crash test dummy was developed over the course of six years by an international team. *Image courtesy of WorldSID.*

Internal Data Recording

Traditionally, each sensor embedded inside a test dummy requires a wire that is connected to an external data recorder. In the early days when there were only a few sensors and cables, it was a relatively easy task. Today's high-channel count tests require serious cable management to deal with the massive bundle of wires that can affect test results and reduce repeatability.

As the first regulation dummy with in-dummy DAS, WorldSID gives new meaning to the term "plug and play." The TDAS G5 DAS is the first data recorder designed to fit inside a dummy as part of a regulation standard (ISO 15830).

A data recorder was designed to fit in the chest, pelvis and legs of WorldSID, all without disturbing the critical dummy center of gravity (CG). Both WorldSID 50th and 5th use the same TDAS G5 in-dummy DAS technology.

"TDAS G5 dramatically changed the scope of testing by making in-dummy DAS part of the ISO specifications, and bringing the world of testing to a higher level of sophistication," says Diversified Technical Systems (DTS) co-founder and company President, Steve Pruitt. "In a typical vehicle testing scenario, where advanced test preparation can consume a week, cutting down that time by eliminating cables becomes a real benefit. It also removes potential inaccuracies caused by cable weight." DE

Mike Beckage *is Vice President/CTO for Diversified Technical Systems.*

ISO FYI

According to the ISO, the World Side Impact Dummy was developed to allow a single, universally-accepted test device to be used for side impact testing in any regulation around the world.

The design details are documented in ISO 15830:2005, "Design and performance specifications for a 50th percentile male side impact dummy (WorldSID)." The documentation includes nearly 500 pages, not including the 400 fabrication drawings and CAD files. It contains all of the design details, material specifications and performance standards required for the fabrication of the WorldSID, according to ISO.

Before WorldSID, there were three different adult male dummies commonly in use that were developed independently from one another. The ISO says WorldSID is intended to allow a single test procedure to be used for side impact in any regulation around the world.

INFO → **Diversified Technical Systems:** DTSWeb.com

→ **International Organization for Standardization:** ISO.org

→ **WorldSID:** WorldSID.org

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

Recipe for Efficiency

Whirlpool bakes a brick and mixes the results with multiphysics software to improve appliance performance.

BY KENNETH WONG

It's difficult to separate family life and work for Nelson R. Garcia-Polanco. His home in Milan, Italy, is equipped with Whirlpool ovens and washers — the ones he and his colleagues designed. Garcia-Polanco, an engineer at Whirlpool, sometimes find himself cooking pizzas with an oven from the same line he's been trying to redesign for energy efficiency.

For his family, an oven that can cook the same meal using less energy means a smaller electric bill. For Whirlpool, energy efficiency is part of its efforts to procure a higher-level European Union Energy label. In partnership with the Green Kitchen project, Whirlpool has begun an initiative to develop cutting-edge products that offer high energy efficiency and reduced impact on the environment. "[The] strategy of maximizing energy efficiency helps consumers reduce waste and achieve savings of up to 70% on energy bills," announces Whirlpool.

To figure out ways to reduce energy use by Whirlpool ovens, Garcia-Polanco and his colleagues Joaquín Capablo and John Doyle run physical tests and thermal simulations in parallel. In digital simulation, engineers could easily do what's usually time-consuming and costly in physical experiments. They could, for instance, alternate the materials and coatings in the oven walls to see how it affects the cooking process.

But before engineers can experiment with what-if scenarios on a computer, they must first cook a piece of brick, literally.

Marinate the Brick

To make an oven more energy-efficient, Garcia-Polanco and his Whirlpool colleagues must first study what happens inside the cavity of the appliance during the cooking process. The standard test for this in the European Union (EN 50304) is to cook a brick. It's a solid piece of heat-conductive material, capable of retaining water, and manufactured to closely resemble the characteristics of generic food — making it an ideal substitute for controlled testing of the heating process.

Eight hours before the test, Garcia-Polanco soaked the brick in cold water. The brick absorbed about 5 oz. of water overnight. "Most food contains water," Garcia-Polanco explains. "And that water evaporates during cooking. The evaporation also increases the heat transfer inside the oven."

The shape, size and functions of the oven are normally standardized. For example, the position of the broiler is often fixed (usually attached to the roof). So one of the strategies is to try out different wall materials and coating options to change the reflectivity, conductivity and emissivity of the oven cavity.

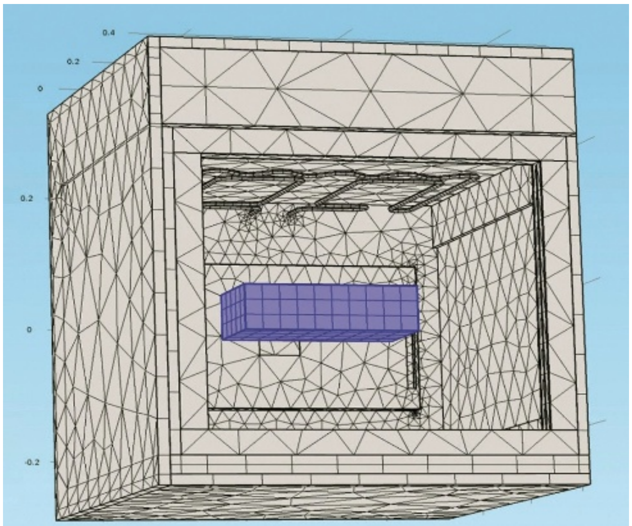

Whirlpool engineers created a "heat and mass transfer" module and a "transport of dilute species" module in their simulations to improve an oven's energy efficiency.

The engineers' plan was to use the brick test as reference, then run a number of simulations using COMSOL Multiphysics software. To obtain data from the physical test, the brick was rigged with thermocouples, feeding right into a data-logging device.

Running the Numbers

Taking detailed measurements from the brick test gave Whirlpool engineers a second-by-second account of what occurs inside the cavity. Engineers have a record of what happened inside the oven as the brick heated up to 392°F over the course of 50 minutes — including the velocity of the air flowing around the brick. "That data gave us the ability to set up a transient simulation, so we could recreate the cooking almost second by second, or for every 10-second interval," Garcia-Polanco says.

The two types of numerical simulation employed are "heat and mass transfer" and "transport of dilute species." The scenario involves non-linear, non-steady-state, partial differential equations. The oven model selected for the exercise is a European model with a 74-liter cavity, operated via a touchscreen.

The correlation between the behavior in the physical test and the digital simulation's output gave Whirlpool engineers confidence in the accuracy of the computation. For example, by measuring the weight of the brick before and after the experiment, engineers could account for water loss during cooking. That figure matched the values produced by software simulation.

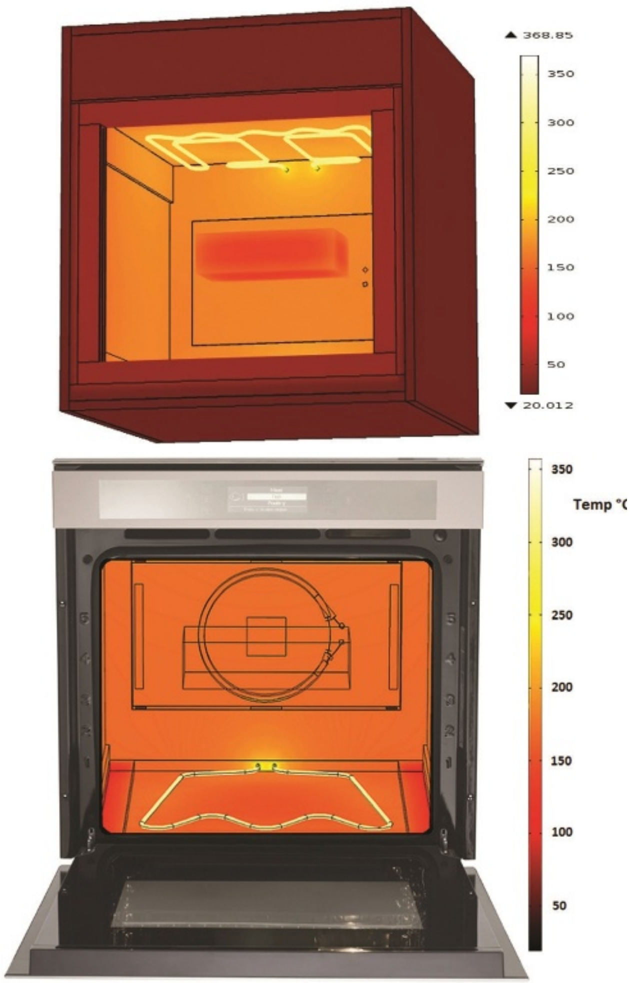

The digital simulation gave engineers temperature evolution over time. The results corresponded to the physical test conducted with a brick.

Checking the Glass

One of the culprits for heat loss during cooking turns out to be the glass door. "We know it's important for the cook to be able to look inside while the oven is operating," acknowledges Garcia-Polanco. "But the glass door is also responsible for radiation leaks, so a lot of heat is wasted."

Simulation in the COMSOL software points to three possible approaches to improve the oven's performance:

1. Figure out a way to change the glass's emissivity.
2. Experiment with the dimensions of the walls.
3. Try out different wall properties by using different materials.

"We are still working on reaching our 20% energy reduction goal," Garcia-Polanco reports. "We're about halfway there."

The next step in the team's plan is to employ parametric studies and design optimization, driven by simulation software.

In the paper written about the experiment ("Multiphysics Approach of the Performance of a Domestic Oven"), Garcia-Polanco and his colleagues note the advantages that come with "the option to use coupled physic models (as, for example, mass and thermal transport), the use of established boundary conditions or the possibility to estimate the prediction error." In this case, they were made possible by COMSOL Multiphysics software.

"Ovens are one of the worst appliances for energy efficiency in the kitchen," Garcia-Polanco concludes. "It's not easy to get to the A+ energy class."

If he succeeds at work, Garcia-Polanco will be able to cook his pizzas at home with approximately 20% less energy. **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 → COMSOL: COMSOL.com

→ Green Kitchen: iapp-GreenKitchen.eu

→ Whirlpool: WhirlpoolCorp.com/about/design/green_kitchen/

Advertising Index

★ SPOTLIGHT ★

Commentary

by Eric Stamper

5 Steps to Large-scale FE Models

It's now common for users solving finite element (FE) problems to use machines with multiple CPUs. With these more powerful computers, tasks such as creating a mesh and solving can be performed in parallel. Yet, as our FE models continue to grow in size by including more components in the total assembly, the limiting factor in the overall analysis process tends to become how quickly an engineer can set up the simulation.

The same idea of parallelization can also be applied to the modeling effort; however, instead of multiple CPUs, we use multiple engineers. This can be a tricky task to achieve in an efficient manner. But if the following five steps are followed, it can be an effective way to decrease the lead time for obtaining the overall result. This process can even evolve into a standard simulation plan for a company that regularly models large assemblies.

1. Divide up the assembly into modular components.

The goal of having multiple engineers building one FE model is to arrive at the finished assembly model faster. Therefore, each individual needs to work on one component or a subset of components in the assembly and follow the same setup process as everyone else. The end result will be to combine everyone's model (or database) to create the full assembly.

> **Clear and consistent communication is critical.**

2. Create component alignment.

Each engineer should build his or her component in the correct "modeling space" so that all the parts, once assembled, will align properly. A master file relating all the positions of all the components could also be used when combining the models, but this can easily be avoided in the upfront planning.

3. Tag attributes.

A tagged attribute is simply a way to identify certain features in the model that can be used in the setup when combining all the models into the full assembly. These attributes will aid in this process when it's time to assign the properties — material properties, boundary conditions, loads and contact pairs, for example. Depending on the software, it might be easier to assign materials, for example, within each database. In other finite element analysis (FEA) software, it might be easier to use the named attributes within each database to assign the materials once combined into the full assembly. Whichever approach is used, these properties, saved into each individual database, must transfer into the main assembly and will be used to define their respective conditions.

4. Implement assembly contact.

The individual databases will eventually be assembled and connected together in the full model. The use of bonded contact will be of great importance in assembly modeling, as it allows for all the components to be meshed independently by each engineer, without having to know what the mesh on all the other components will look like. Some FEA software, like ANSYS Mechanical, has the built-in capability to automatically recognize mating surfaces and create a bonded contact pair. If this capability is not available when combining the full assembly, this is when the consistent naming of surfaces (tagged attributes) facilitates manually setting up the contact between components. For example, a toy figure is divided into four components ("body," "left arm," "right arm" and "legs") in individual databases. Each component was built in the correct position in the modeling space, and the surfaces where the parts will "stick together" are tagged with the names. In the "left arm" database, some surfaces are tagged with "bond_left_arm_to_body;" in the "body" database, some surfaces are tagged with "bond_body_to_left_arm." Listing all the contacts that need to be bonded together in a master file will make it easy to automate the contact creation with a short script.

5. Follow guidelines and document, document, document.

While this step is listed last, it needs to be established at the very beginning before each engineer starts working on his or her model — and be continuously updated throughout the entire setup process. Documentation will allow the team to keep track of what's in each model. It ensures that everyone is building their model in a consistent way, naming all their attributes with a similar convention and notating how all the components will come together in the full assembly. Clear and consistent communication among all the engineers is critically important.

This modular approach to large assembly modeling helps decrease the overall lead time on setting up the analysis. Another added benefit of using this technique is that it makes it easier to swap new components in and out without affecting the rest of the assembly. This leads right into applying this modular simulation approach in a design of experiments to perform quick design changes on the component level and determine the full system's response to the new design. DE

Eric Stamper *is engineering manager at CAE Associates Inc. He has more than 10 years of experience in engineering simulation and analysis. Contact him via CAEAI.com/fea-services.*

So many combinations–
it might blow your mind.

[COLOR + RUBBER + TRANSPARENT + RIGID]

Introducing the Objet500 Connex3 from Stratasys, the world's only full-color and multi-material 3D printer. Select from incomparably brilliant and consistent colors, plus a full palette of transparent colors — the only 3D printer to offer such a wide array. And it's the only 3D printer that prints flexible materials in a broad range of shore values. All with ultra-fine detail creating the most true-to-life modeling possible. Stratasys is the proven leader in multi-material 3D printing. **For whatever your mind can imagine, visit stratasys.com/Objet500Connex3.**

Stratasys® | Objet500 Connex3