

DE

Desktop Engineering®

TECHNOLOGY FOR DESIGN ENGINEERING

November 2013 / deskeng.com

Working with
Large CAD Files P.42

Simulations on
the Cloud P.38

Rendering P.34



HIGH-PERFORMANCE PRODUCTIVITY

Workstations, clusters and
cloud computing speed
complex CAD, rendering
and simulation. P.24



METAL 3D PRINTING P.46

REVIEW: 3DBOXX 4150 XTREME P.30

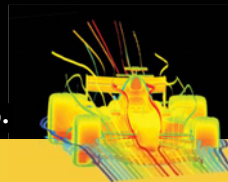
SIEMENS INCORPORATES LMS P.50

A Winning Formula

Infiniti Red Bull Racing and ANSYS



Victories by Infiniti Red Bull Racing. Simulation by ANSYS.



ANSYS®

Realize Your Product Promise®

Winning the Formula One World Constructors' Championship is a monumental accomplishment. Winning it three 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.



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



**GOES FROM ZERO
TO D-SIZE COLOR PLOT
IN 25 SECONDS**

And we
don't mean
draft mode.



THE NEW EPSON® SURECOLOR® T-SERIES

- Produces a presentation quality, color, D-size plot in 25 seconds
- Extreme plotting accuracy at resolutions up to 2880 x 1440 dpi
- Advanced pigment ink technology for truly durable, full-color plots
- High-capacity ink cartridges up to 700 mL for low printing cost
- Space-saving design with easy access front-loading paper and ink

Epson SureColor T3000 - 24" | \$2,995
Epson SureColor T5000 - 36" | \$3,995
Epson SureColor T7000 - 44" | \$4,995

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. For the SC-T3000, fastest print speed for a D-size plot is 28 seconds. 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. Copyright 2013 Epson America, Inc.

EPSON.COM/PLOTTERINFO



The Art of Engineering

It's easy to get completely caught up in the technical details of the latest software and hardware. There's always a new development, a better piece of equipment, or an innovative way of using technology to focus on in the specific language of megahertz, data points or gigabytes per second. But there's another side of design engineering that deals with less precise vocabulary like the look, feel and even emotional attachment people have to products. The design side of design engineering is more difficult to quantify, but often just as important as the technical aspects.

The National Endowment for the Arts recently attempted to define and quantify industrial design in *Valuing the Art of Industrial Design*, which it calls a "profile of the sector and its importance to manufacturing, technology, and innovation." The report, which spans such topics as industrial designers as inventors and entrepreneurs, hiring industrial designers as service providers, and how "design thinking" is being applied to all

Design thinking can make a positive impact on every stage of manufacturing.

manner of problems inside and outside of traditional engineering environments, is the first such attempt to analyze federal data and draw industrial design industry conclusions from it.

Industrial Design, by the Numbers

While the report is quick to point out the challenges it faced in defining and quantifying industrial design, it does establish parameters and then apply numbers from the Bureau of Labor Statistics, U.S. Census Bureau and the U.S. Patent and Trademark Office. According to the report:

- There are more than 40,000 industrial designers in the United States, 30% of whom are self-employed.
- Most salaried industrial designers — 11,730 of them — work in the manufacturing sector, with 7,570 working for what the report calls the "professional services sector," which includes specialized design firms, architectural or engineering firms, consulting firms, and research and development.
- However, industrial designer employment in the professional services sector is projected to leap by 29% when measured from 2010 to 2020.
- The Bureau of Labor Statistics' Office of Employment Projections expects the number of employed industrial designers in the U.S. to reach 45,100 by 2020.

- There are 1,579 industrial design businesses in the U.S.
- In 2007, the most recent year for which such data is available, industrial design firms earned more than \$1.5 billion in total revenue. About 94% of that came from sales of product design, model design and fabrication.
- The four largest industrial design firms generate 11% of the industry's total revenue; the 20 largest, 32%.
- The number of U.S.-awarded design patents per 100,000 population is at an all-time high: seven per 100,000 in 2012, compared with one per 100,000 at the turn of the 20th century.
- In 2011, \$9.2 million was pledged on the crowd-funding site Kickstarter to support design projects, and 319 successful projects were funded.
- California and Michigan, as hubs of the auto and aerospace industries, each employ more than 3,000 industrial designers.

Drawing Data Conclusions

The numbers are interesting, especially in light of the fact that this is the first time they've all been brought together to define industrial designers. But what do they mean?

They show how important the design side of design engineering is to innovation. Design patents are at a record high and, as the report points out, "approximately 40% of inventors named on design patents were also named on utility patents." Utility patents are those that are designed to protect how a product works, as opposed to how it looks. That percentage shows how critical the creative process behind designing the look and feel of a product is to invention.

That fact isn't lost on engineering firms, as the projected 29% leap in industrial designer employment in the professional services sector shows. While manufacturing declines are expected to reduce the numbers of industrial designers in that sector, it is still expected to employ the largest share of them by 2020.

Not all industrial designers being trimmed from manufacturing will move into professional services. With 30% of them already self-employed, the \$9.2 million in design funding on Kickstarter could serve as this era's "there's gold in them there hills" rallying cry.

Executives would do well to make the most of industrial design by keeping design engineers fully engaged in multiple aspects of the business, even if their company's products will never get a close-up in a slick advertising campaign. Design thinking — how design engineers creatively approach and solve problems — can make a positive impact on every stage of manufacturing. **DE**

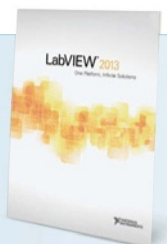
Jamie Gooch is the managing editor of Desktop Engineering. Contact him at de-editors@deskeng.com.

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

©2013 National Instruments. All rights reserved. LabVIEW, National Instruments, NI, and ni.com are trademarks of National Instruments. Other product and company names listed are trademarks or trade names of their respective companies. 12118





COVER STORY

Workstation, Cluster or Cloud?

24 Choosing where to run complex computational models has become a choice driven by more than just costs. Today, many factors come into play when choosing from among workstations, an onsite computing grid/cluster, or cloud services. This article by Frank Ohlhorst supports our issue focus on how today's engineering computing technology provides multiple paths to an optimized design cycle.

ON THE COVER: Design engineers have many high performance computing options from which to choose to help them speed up the design cycle.

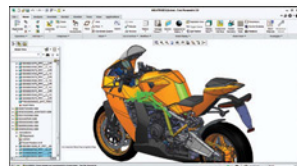
Buick rendering courtesy of RTT, front loader model screenshot courtesy of Autodesk, Inc., Background image courtesy of iStock.com.

DESIGN

42 Breaking through the Large Assembly Barrier

Increasing product complexity and cross-disciplinary design are adding to the size and scope of CAD models, but new technology advances are facilitating engineers' interaction with large assemblies.

By Beth Stackpole



PROTOTYPE

46 Metal 3D Printing on the Rise

From ordinary to exotic, AM machines and materials are getting the spotlight.

By Pamela J. Waterman



TEST

50 The Simulation-Test Loop

Siemens PLM Software's LMS business segment improves performance and system linkage with Test.Lab Rev 13.

By Randy Frank



FOCUS ON ENGINEERING COMPUTING

20 The Zombie Computer Survival Guide

Here are some things to consider when shopping for HPC for simulation.

By Kenneth Wong



30 Workstation Review: An Optimized Single-socket Solution

3DBOXX 4150 XTREME, the latest single-CPU workstation from BOXX Technologies, sets a new price/performance standard.

By David Cohn



34 Faster Rendering and Visualization

Today's technologies are making it easier than ever to speed up the design process.

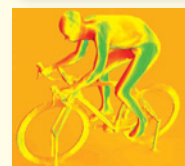
By Mark Clarkson



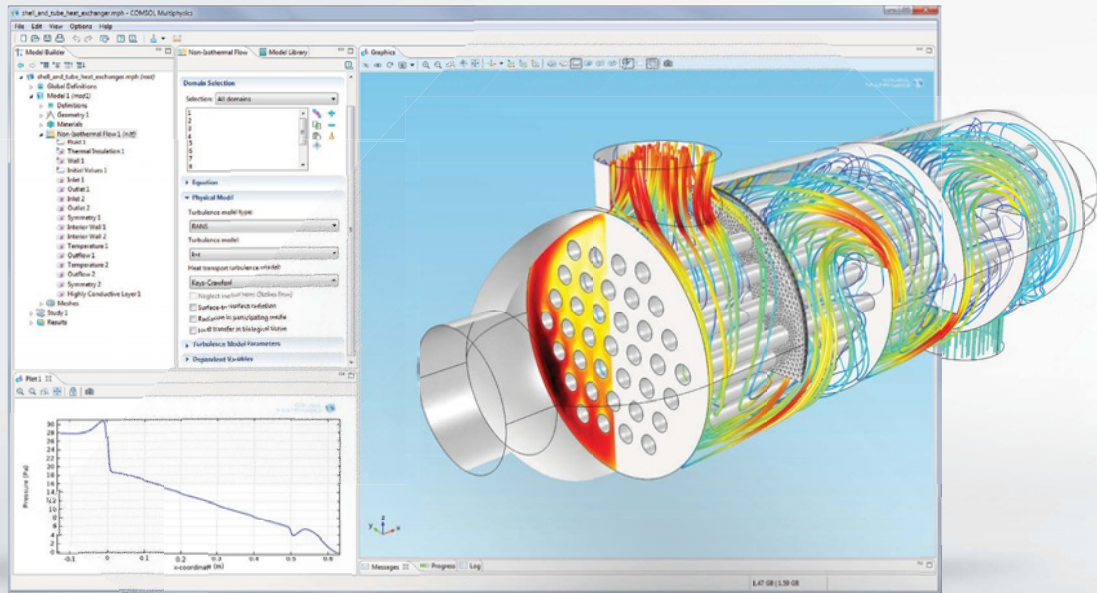
38 Real-life Results with the UberCloud CAE Experiment

Engineers are seeing higher computing power and ease of use in moving designs and simulations to the cloud.

By Peter Varhol



HEAT EXCHANGER: Model of an air-filled shell and tube heat exchanger with water flowing through the inner tubes.



Verify and optimize your designs with COMSOL Multiphysics®

Multiphysics tools let you build simulations that accurately replicate the important characteristics of your designs. The key is the ability to include all physical effects that exist in the real world. To learn more about COMSOL Multiphysics, visit www.comsol.com/introvideo

Product Suite

COMSOL Multiphysics

ELECTRICAL

AC/DC Module
RF Module
Wave 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
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
ECAD Import Module
LiveLink™ for SolidWorks®
LiveLink™ for SpaceClaim®
LiveLink™ for Inventor®
LiveLink™ for AutoCAD®
LiveLink™ for Creo™ Parametric
LiveLink™ for Pro/ENGINEER®
LiveLink™ for Solid Edge®
File Import for CATIA® V5



DEPARTMENTS

2 Degrees of Freedom

The art of engineering.

By Jamie J. Gooch

8 Virtual Desktop

Altair's HPC appliance, PTC acquires NetIDEAS, DesignSpark Mechanical is free, and Share3D improves collaboration.

By Kenneth Wong & Beth Stackpole



14 Engineering on the Edge

Japan simulates railways, Nissan readies autonomous cars, awards honor fuel cells, and nanostructure-level printing.



16 Rapid Ready Tech

Making 3D designs from 2D images, 3D-printed building structures, fold-out large-scale 3D printing and the International Wohlers Conference preview.



18 Editor's Picks

Products that have grabbed the editors' attention.

By Anthony J. Lockwood



45 Spotlight

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

49 Advertising Index

52 Commentary

Improving combustion modeling.

By Bernie Rosenthal, Reaction Design



deskeng.com

RAPID READY TECH BLOG

Read all about making digital designs physical @ rapidreadytech.com.

ENGINEERING ON THE EDGE BLOG

See the future of engineering technology @ engineeringontheedge.com.

MOBILE ENGINEER

Incorporate mobile technologies into your workflow.
@ mobileengineer.com.

VIRTUAL DESKTOP BLOG

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

NEWSLETTER REGISTRATION

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

DE MOBILE BETA

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

DE ON TWITTER

Follow us @ DEeditor

LIKE DE ON FACEBOOK

Visit DE @ deskeng.com/facebook



EDITORIAL

Jamie J. Gooch | Managing Editor
Kenneth Wong | Senior Editor
Anthony J. Lockwood | Editor at Large
Heather Pittinger | Copy Editor

CONTRIBUTING EDITORS

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

PUBLISHER

Steve Robbins | Publisher

ADVERTISING SALES

603-563-1631 • Fax 603-563-8192

Erich Herbert | Sales Executive (x263)
Jeanne DuVal | Sales Executive (x274)

ART & PRODUCTION

Darlene Sweeney | Director (x257)

A LEVEL 5 COMMUNICATIONS PUBLICATION

Steve Robbins | Chief Executive Officer
Thomas Conlon | President

ADVERTISING, BUSINESS, & EDITORIAL OFFICES

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

SUBSCRIBER CUSTOMER SERVICE

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

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

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

LIST RENTALS: For information on list rentals, contact Statistics, 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.

Each separate contribution to this issue, and the issue as a collective work, is copyright © 2013 Level 5 Communications, Inc. All rights reserved. Copying for other than personal or internal reference use without the permission of Level 5 Communications, Inc. is prohibited. Requests for permission should be addressed in writing to Desktop Engineering Permissions, 1283D Main Street, P.O. Box 1039, Dublin, NH 03444. PRINTED IN THE USA.



SHARE DESIGN FEEDBACK FASTER

From design review to fabrication, Bluebeam® Revu® provides digital workflow solutions to improve project communication and speed up productivity. Create 2D and 3D PDFs from AutoCAD®, Solidworks® or any IFC file. Annotate 2D and 3D PDFs with customizable markup tools, and compare drawing revisions to automatically identify design changes. Review and comment on the same set of drawings with project partners around the world in real time using Revu's integrated collaboration solution, Bluebeam Studio™.

Anything is possible.

www.bluebeam.com/designsmarter



Altair Transforms HPC into Turnkey Appliance

High-performance computing (HPC) has been associated with a high barrier for adoption — particularly among small- and mid-size companies that lack the resources to procure, manage and maintain large HPC clusters.

Recently, simulation specialist Altair began offering a solution that might level the playing field: The company has just released HyperWorks Unlimited, which essentially packs all the pivotal HPC technologies (cluster technology, CAE software and workload management capabilities) into a turnkey solution.

“From workload management to application software, we do all the different aspects,” notes Ravi Kunju, Altair’s vice president of strategy and business, Enterprise Solutions. “We’re trying to make the entire process of bringing on

HPC completely easy for the customer and eliminating an entire food chain to deal with.”

For the hardware piece of HyperWorks Unlimited, Altair is partnering with SGI. The pair is delivering an Altair-branded appliance that leverages Intel Cluster Ready technology, which unites Altair’s HyperWorks software with hardware based on Intel’s Xeon E5-2600 processor family. HyperWorks Unlimited is fully optimized to run massively parallel applications such as computational fluid dynamics (CFD) and finite element analysis (FEA) solvers (Altair’s and others). It also includes Altair’s Display Manager, a remote visualization web-based portal for accessing graphically intensive applications on HyperWorks Unlimited.

The combination of technologies

creates an optimized private cloud as the backbone of HyperWorks Unlimited.

The approach addresses three of the primary barriers to adoption that have long plagued HPC technology: scalability of application licensing; the long lag time associated with moving the vast amounts of data from clusters to local machines for post-processing and visualization; and the perceived data security concerns associated with public clouds.

HyperWorks Unlimited says it eliminates these issues. The offering includes unlimited use of all Altair software, including all HyperWorks applications and the PBS Works workload management tools, Kunju says, so there is no restricted access and no more concerns related to scalability of application licensing.

— B. Stackpole

PTC Makes Play for Managed Services

For more than a decade, NetIDEAS has provided secure hosting environments for PTC customers looking to offload the heavy lifting around product lifecycle management (PLM) deployment and day-to-day maintenance. For businesses that are not ready or willing to invest in the server and storage infrastructure to run Windchill PLM on their sites (along with the IT personnel to administer the enterprise application), NetIDEAS hosts, administers and supports Windchill, augmented with additional consulting services.

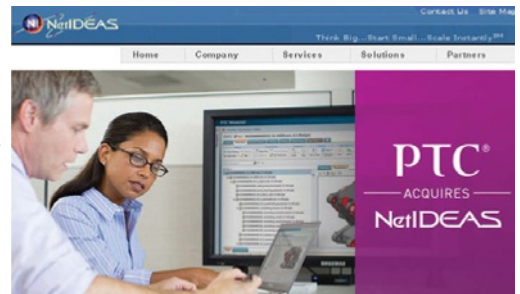
In late September, in a move that some might call a stepping-stone to the cloud, PTC acquired its long-time partner NetIDEAS. Sanjay Verma, PTC’s divisional vice president of global services, was adamant that the acquisition has nothing to do with PTC’s long-term vision for the cloud.

“This has nothing to do with Soft-

ware-as-a-Service (SaaS); this is a completely different model, and we’re not talking about the cloud at this time,” he says. As opposed to delivering a version of Windchill available as a monthly service over the public cloud, the PTC-NetIDEAS relationship makes the PLM software available as a managed service and with a traditional licensing arrangement.

Verma says the NetIDEAS technology and staff will allow PTC to broaden the managed services offering to small- and mid-sized companies that now recognize the value of PLM, but don’t necessarily want to invest in doing sophisticated PLM development.

“Companies recognize that in order to get the full value of PLM, they need to develop sophisticated applications and a managed delivery system in their own organization — and some have decided they don’t want to do that,” Verma says.

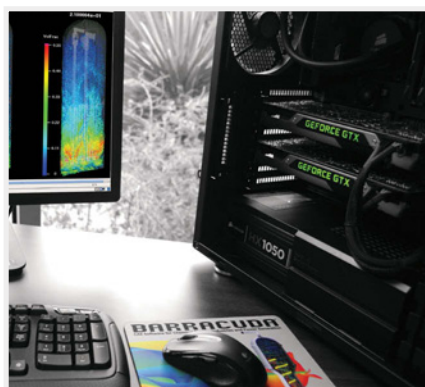


What NetIDEAS and the whole model of managed services bring to PTC is the opportunity to get closer to these customers and gain a deeper understanding of their pain points and needs. Part of that learning experience will be figuring out how to approach the cloud, which Verma admits is an inevitable transition.

“This will help us understand that market much better,” he says. “No doubt, the world is moving towards SaaS, but there are questions in the company as to when and how.”

— B. Stackpole

For Chemical Engineers,
Best-in-Class CFD for
Fluidized Bed Modeling
Now Available with
GPU Acceleration
for Desk-side Supercomputing!

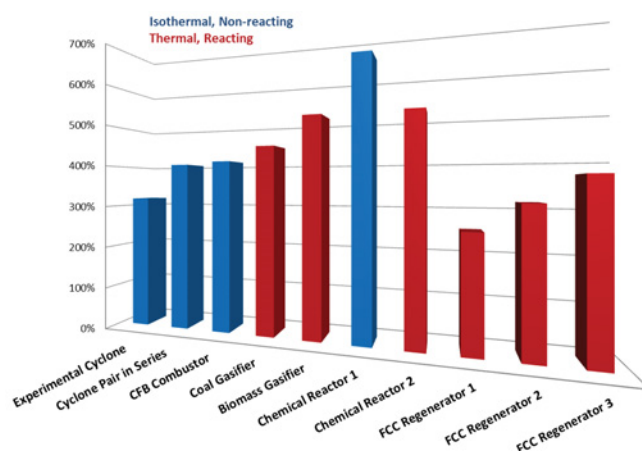


Optimize Fluidized Bed Reactors *with* Barracuda VR®

- Increase Yields
- Improve Reliability
- Extend Operating Cycles
- Control Emissions
- Scale Up Throughput
- Minimize Catalyst, Feed and Product Loss

The Right Answers, Right Now.

- Get Answers Up to 650% Faster
- Run on Low-cost GPU Cards
- Supports Unlimited CUDA Cores with a Single License



Relative Speed in % for Reacting & Non-Reacting Cases
Barracuda VR Series 16-GPU vs. Barracuda VR Series 15

Proud to run on  **NVIDIA**

To learn more, visit us at cpfd-software.com/de2013

cpfd COMPUTATIONAL
PARTICLE
FLUID DYNAMICS

DesignSpark Mechanical: A Free 3D Modeler Derived from SpaceClaim

In September, when Allied Electronics and RS Components (the trading brands of Electrocomponents plc.) released a free 3D modeling software called DesignSpark Mechanical, some of you might have found the product to be a bit familiar. That's because the modeling technology of DesignSpark Mechanical comes from SpaceClaim, a household name in 3D design.

"Allied and RS are partnering with SpaceClaim to launch DesignSpark Mechanical, which combines the power and ease-of-use of direct modeling technology from SpaceClaim with access to the massive library of standard parts from Allied and RS, trusted by millions of engineers around the world," says Rich Moore, vice president of business development for SpaceClaim.

Bringing Together Electronic and Mechanical Disciplines

Allied and RS are in the business of supplying ready-made electronic components, available

from online catalogs. Engineers and designers developing products that involve electric components usually incorporate digital counterparts of these components into their 3D design assemblies. The need to work with these components spans across mechanical design and electronic design — two disciplines that don't frequently work together, but should.

Both disciplines wrestle with design changes when collaborating, because traditional history-based mechanical CAD packages don't accommodate swift revisions. Mismanaged or miscommunicated changes — whenever the electrical engineer revised the printed circuit board (PCB) design or the mechanical designer adjusted the space allotted to the electrical components — can cause friction and headaches. That's the specific area Allied and RS are expecting to serve with DesignSpark Mechanical, based on SpaceClaim's direct-editing technology with a low learning curve.

Online Libraries Save Time

DesignSpark.com, the online portal where the free software is delivered, also offers a large library of approximately 38,000 readily downloadable components, identifiable by manufacturers, stock number and more. Allied and RS' partnership with TraceParts, another 3D content provider, gives software users millions of additional components from TracePartsOnline.net.

This ought to be an appealing feature to DesignSpark Mechanical users, because they won't need to create these components from scratch. Furthermore, the automated bill of material (BOM) extracted from the 3D model will include all the part numbers, manufacturers and purchasing info. This is also in the interest of component suppliers, as the BOM will eventually direct manufacturers to their online catalogs for purchasing.

Free Software

Like SpaceClaim, DesignSpark Mechanical can export STL files, which can be used to produce physical prototypes in 3D printers. Because SpaceClaim is a commercial product, however, it's reasonable to assume the freely distributed DesignSpark Mechanical will not contain all features and functions available in SpaceClaim.

To download the software, visit designspark.com/eng/page/mechanical.

— K. Wong



Based on SpaceClaim's direct editing technology, the free software DesignSpark Mechanical is expected to facilitate collaboration and mechanical and electrical designers.

This is what your competition is using.



"Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries."

www.bboxtech.com

Share3D from QuadriSpace: Another Desktop-Cloud-Mobile Collaboration Solution

Another bridge just went up in the collaboration space: Share3D is built by QuadriSpace, which offers a series of CAD-to-3D PDF publishing packages. The Share3D app, available for both iPhone and iPad for free, is designed to work in conjunction with the company's cloud-hosted offerings and the desktop product.

Several products from the Share3D line — the iPad app, the cloud services and the desktop client — come together to form a collaborative environment for project team members working on different devices in different places. This is the new workflow driven by what some are calling the BYOD (“bring-your-own-device”) policy now gaining acceptance in the design and engineering professions.

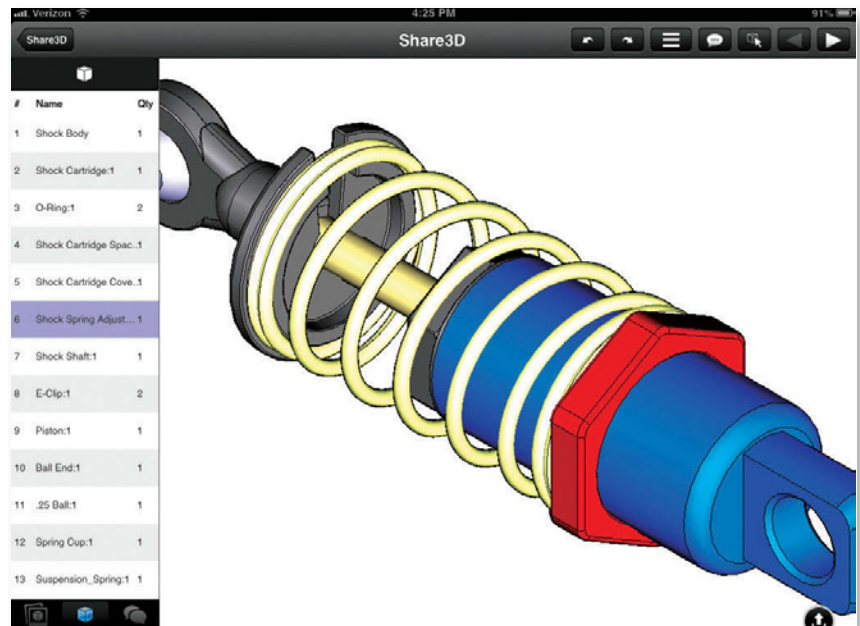
Other vendors like GrabCAD and CadFaster are also banking on this trend. GrabCAD skips the desktop client, but allows desktop machines to communicate with online content and iPad through standard browsers.

Share3D is offered at three different levels:

- Free: Three projects, up to five participants per project (\$0)
- Standard: Twenty-five projects, up to 25 participants per project (\$29 per month)
- Professional: Unlimited projects, unlimited participants (price unpublished)

There is, however, one distinct difference between the Share3D version and others: Share3D subscribers get access to the Free, Standard or Professional edition of Publisher3D that matches their plan. The Publisher3D product would allow Share3D subscribers to export 3D PDF files with markups and annotations.

Brian Roberts, founder and president of QuadriSpace, says that with Publisher3D, subscribers would be able to “publish step-by-step instruction manu-



QuadriSpace's Share3D, shown here as an iPad app, joins the multi-device collaboration market that spans desktop software, mobile devices and cloud services.

als, custom bill of materials, exploded views, and upload them to the cloud.”

Share3D comes with what the company describes as “email in the loop” — automatic email notices with screen captures sent to project participants when someone annotates or comments on a shared file.

“The real power of it is that users can simply reply to the automatically generated email and it will be shared just like any other comment,” Roberts explains. “The benefit here is that everyone can participate, using the same work process (email) that they have always used, and don’t even have to log into an app or the website.”

A similar email-integrated collaboration could also be found in Vuuch, a social collaboration plug-in available for Autodesk Inventor, SolidWorks and a few CAD software titles. The Share3D app

uses a lightweight 3D XML format to facilitate viewing and markup among collaborators. It’s not intended, however, to accommodate editing native CAD files.

“We never put native CAD files in the cloud,” says Roberts. “It’s added security, since we’re not putting your SolidWorks or Inventor files in the cloud.”

When lightweight desktop apps with viewing and markup features dominated the collaboration space, collaboration was confined to PC-to-PC communication. But the rising popularity of mobile devices and cloud storage has irreversibly changed the trend. The collaboration framework today needs to accommodate desktop software, mobile apps and browser-accessible cloud storage. QuadriSpace’s Share3D is another example of that multi-device workflow.

— K. Wong

Configurator Puts Time **Back** on Your Side

Use the online Workstation Configurator to determine the right engineering system for your workflow.

While most engineers understand the core components of workstation technology, matching the technical specifications of processors, GPUs, and memory to specific engineering workflows isn't exactly a straightforward experience. Any misstep can have significant financial ramifications, leaving an engineer stuck with a pricey, underpowered system that doesn't meet their needs or conversely, investing too much in a platform that is overkill for their day-to-day design tasks.

To help engineers sort through the confusion and make the best workstation match, Intel and SolidWorks collaborated on the development of a Workstation Configurator tool, an online tool designed to help engineers break out of the one-size-fits-all mold mentality and build a productive, custom workstation without overinvesting in technology. Working with engineering independent software vendors (ISVs), Intel developed a set of benchmarks that help identify workstation configurations optimized to provide the best overall design experience for specific engineering tasks.

How it Works

Visitors are prompted to choose their typical engineering workflow broken down into three choices: 1) 2D and 3D basic design for working with less featured parts in small assemblies that don't demand simulation; 2) advanced design and basic simulation for addressing the needs of 3D design and modeling, and simulation of complex parts and larger assemblies; and 3) advanced simulation and rendering, for work that requires flow simulation and photorealistic rendering of large projects.

For each workflow, the tool directs visitors to choose their most likely level of design complexity while serving up workflow requirements on screen to provide guidance to that choice. Once selected, the tool comes back with a recommended system configuration, with detailed descriptions to help the user navigate and fully understand the tradeoffs between important system components like processor cores,



The Workstation Configurator helps engineers determine their workstation needs based on their individual workflows.

memory, video cards, and the right storage medium. There is also the ability to print out the spec sheet in addition to links to specific OEMs so users can vet out their optimal workstation configuration and get real-world pricing.

Two Examples

Let's say, for example, a user mostly engages in simple 2D and 3D design—a workflow the tool equates with designing parts of only one to 25 features and little requirement for assembly modeling. For that specific workflow, the recommended workstation configuration is an Intel Xeon processor E3-1200 v3 family with four cores, 8GB of memory, an Intel HD Graphics P4600, and HDD storage.

Yet, someone looking to do advanced simulation and rendering work on parts with more than 200 features and assemblies with more than 2,000 unique components is directed to a very different system. In that case, the recommended configuration is stocked with an Intel Xeon processor E5-2600v2 based workstation with up to two processors of memory, a mid-range discrete graphics card, and SSD storage.

Of course, nothing is set in stone. Users can fine tune the recommended configuration to see how additional processor cores or the selection of a higher end graphics card will affect their system performance as well as the pricing.

To take the Intel Workstation Configurator for a test drive, visit: intel.com/content/www/us/en/workstations/workstation-configurator-tool.html.



Japan Boosts Railway Simulation Capability

The Railway Technical Research Institute (RTRI) in Japan has added more muscle to its supercomputer infrastructure



for complex railway simulations. The organization, which focuses on research and development of railway-related science and technology, has deployed a Cray XC30-AC supercomputer, a Cray CS300 cluster supercomputer and a Cray Sonexion storage system into production.

The RTRI's research spans everything from earthquake cluster prediction to improving the fatigue strength of wheel/axle fittings. The XC30-AC is the institute's primary high-performance computing (HPC) system, and will be used to run advanced simulations. The CS300, meanwhile, is used as a general-purpose application server.

The RTRI system has a peak performance of more than 100 teraflops. The Sonexion system includes 220TB of capacity and 10GB per-second of application horsepower.

MORE → deskeng.com/edge/?p=5503

Nissan Says Autonomous Cars are Coming in 2020

Nissan plans to have commercially viable autonomous drive technology in multiple vehicles by 2020. The company outlined its plans this summer, including the construction of a dedicated autonomous driving proving ground in Japan that should be completed in 2014. The facility will include mock townscapes that include real masonry buildings, to perform tests that go beyond what is currently possible on public roads.

At the Nissan 360 event, the company demonstrated Nissan LEAFs outfitted with laser scanners, around-view monitor cameras, and advanced artificial intelligence and actuators. The autonomous driving technology builds on the company's Safety Shield system, which provides 360° monitoring around the vehicle and offers warnings to drivers.

MORE → deskeng.com/edge/?p=5480

GreenTec Highlights Fuel Cell Technology

Innovative fuel cells took honors at this year's GreenTec Awards in Berlin. Airbus walked away with the award in the aviation category, thanks to its work in integrating fuel cells in commercial aircraft.

The company's research analyzed replacing the auxiliary power unit and ram air turbine with a multifunctional fuel cell. In the Airbus model, the waste products from the fuel cell could be repurposed elsewhere on the aircraft.



Australia's Ceramic Fuel Cells Ltd took the award in the energy category for its BlueGEN technology, a dishwisher-sized fuel cell "micro power plant" that runs on natural gas that can produce energy and heat for buildings. It can operate at up to 60% efficiency.

MORE → deskeng.com/edge/?p=5493



Research to Tackle Informatics of Making

Researchers from the University of Wisconsin-Madison and Drexel University hope to develop new computational tools that combine computing, materials and manufacturing advancements to better account for the complexity of new products that are manufactured in non-traditional ways using advanced materials.

Armed with a grant from the National Science Foundation, the team will attempt to tackle what they refer to as the "informatics of making."

The goal is to better match computational design tools with materials science.

MORE → deskeng.com/edge/?p=5467

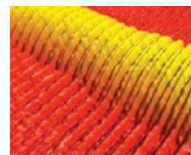
Higher-resolution Nanostructure Printing

Researchers at the University of Illinois at Urbana-Champaign, the University of Chicago and Hanyang University in Korea have combined advanced ultra-high resolution inkjet technology with self-assembling block copolymers to increase the resolution of intricate nanostructure fabrication — down to 15 nm from 200 nm.

The use of block copolymers helps mitigate the limitations of inkjet printing, which can typically only achieve a resolution down to 100 to 200 nm. Being able to create nanostructures from soft materials could potentially help create new classes of electronics and sensors.

Engineers create a topographical or chemical pattern using electrohydrodynamic printing (e-jet), then place a block copolymer on top of it, which uses the printed template to form patterns at a higher resolution.

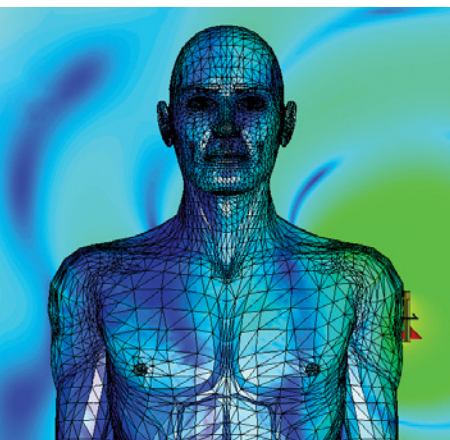
MORE → deskeng.com/edge/?p=5566





Make the Connection

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



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

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

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

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

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



3-Sweep Brings Ease of Use to 3D Modeling

Research by Ariel Shamir, Ph.D., from the Interdisciplinary Center at Herzliya, and Tel Aviv University's Daniel Cohen-Or, Ph.D., and Tao Chen, Ph.D. has revealed a method of pulling 3D designs from 2D images that has potential for improving



both current and future 3D design.

Dubbed 3-Sweep by the research team, the system uses both computer- and user-generated data to build 3D objects. (As anyone who has ever used Photoshop could testify, computers aren't great at grabbing the proper information from pictures without help.)

MORE → rapidreadytech.com/?p=5300

SpaceX Moves Toward Gesture-based 3D Design

Along with explosions and villainous schemes, the *Iron Man* movies have presented popular media with an updated example of what future gesture-based technology might look. Tony Stark's hands magically wave, pinch, pull and grasp information to design his suits. After all, in a visual medium like the movies, this sort of design process is much more interesting than watching him plug away at a CAD program for several hours.

Most moviegoers were simply entertained with the design spectacle, but Elon Musk, founder of SpaceX and Tesla Motors, thought the technology used by Stark might not be all that far out of reach. By combining a Leap Motion controller with the Oculus Rift VR headset, Musk and SpaceX developed a basic interactive and

International Wohlers Conference to be Held Alongside EuroMold

The International Wohlers Conference is a series of informative seminars hosted by Wohlers Associates, an additive manufacturing (AM) consulting firm. This year's International Wohlers Conference will focus on business and investment opportunities.

As 3D printing evolves from prototyping to end-use, the potential for gains and challenges to overcome grows along with the technology. The speakers at this year's Wohlers conference have been selected to discuss the research, development, business and investment aspects of AM. Speakers include a keynote address from Dr. Olaf Diegel, CEO ODD Guitars, who will provide examples of end-use products built through AM; Terry Wohlers, president of Wohlers Associates, speaking about the current global state of 3D printing; Wilfried Vancraen, CEO Materialise, who will break down the evolution of the AM market and forecast its future; and William J. Cass, Cantor Colburn LLP, with a presentation on intellectual property as pertaining to AM.

The Wohler's International Conference will be held on Dec. 5, and is a full-day event. EuroMold runs from Dec. 3-6 in Frankfurt, Germany.

MORE → rapidreadytech.com/?p=5553

design program that uses gesture control in a 3D environment.

According to Musk, it only took a couple weeks after SpaceX began experimenting with the Leap Motion controller and a standard computer display to move to virtual design.

MORE → rapidreadytech.com/?p=5292

AM Enters the Construction Yard

The Swedish-based construction firm Skanska will be using additive manufacturing as part of the construction of the Bevis Marks project in London. Topping the new construction will be an ethylene tetrafluoroethylene (ETFE) roof, supporting by 3D-printed cladding.

Because ETFE is much lighter than glass — similarly sized windows weigh around 1% when built with ETFE — the material gives architects a fair amount of freedom when designing decorative elements of a structure. 3D printing compliments this freedom by offering further options when it comes to lightweight designs.

In the case of the Bevis Marks project, the original plan called for a complex shroud of eight different cast steel nodes to hold up the ETFE roof. Building the



shroud through traditional manufacturing methods was considered to be either too costly or lacking in aesthetic appeal, so the company turned to AM.

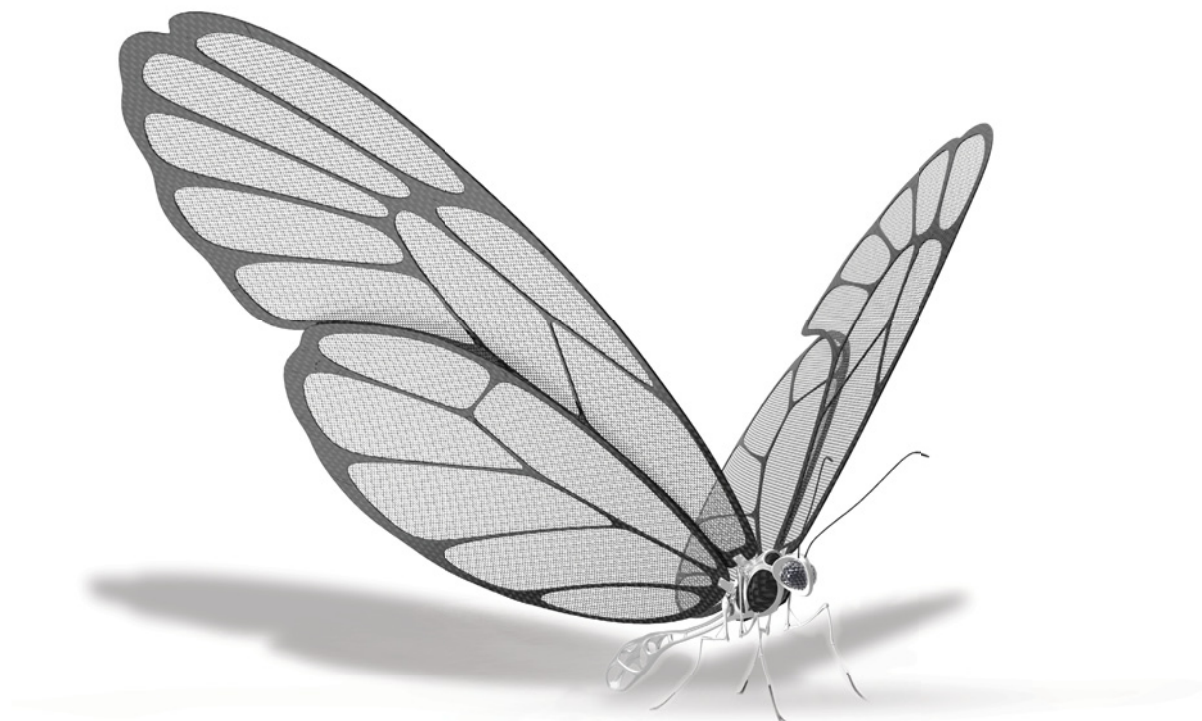
MORE → rapidreadytech.com/?p=5401

New Twist on Large AM

Researchers at MIT have developed a new method of printing large objects in limited-build envelopes that they call Hyperform.

The idea comes from Marcelo Coelho, Ph.D., and Skylar Tibbits. The research team designed, printed and constructed a chandelier that was approximately 4x4x4 ft. in a FORM 1, which has a build envelope of 4.9x4.9x6.5 in. The entire chandelier was printed out in a chain that included joints that are only capable of bending in one direction to assist with assembly. Once every joint was bent and snapped into place — in origami-like fashion — the chandelier took form.

MORE → rapidreadytech.com/?p=5378



Innovative design can make anything lighter.

Minimizing component weight with advanced engineering, technology and materials. Without sacrificing performance. It's what we do. **More at altair.com/lightweight.**





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.



3DBOXX 4150 XTREME with Intel Core i7, NVIDIA Maximus

Processor speed of up to 4.3 GHz and NVIDIA graphics said to deliver outstanding support for modeling, simulation, and visualization.

BOXX Technologies' 3DBOXX 4150 XTREME workstation starts at around \$2,800. So what do you get for that?

It starts with one of Intel's new 4th generation quad-core Core i7 processors that is overclocked to run at up to 4.3 GHz and is liquid-cooled. The Core i7 provides

Smart Cache for increased performance and hyper-threading as well as Intel Turbo Boost Technology 2.0. Next comes a 1GB NVIDIA Quadro K600 graphics card paired with an NVIDIA Tesla GPU (graphics processing unit) using NVIDIA Maximus technology.

MORE → deskeng.com/articles/aabkrh.htm



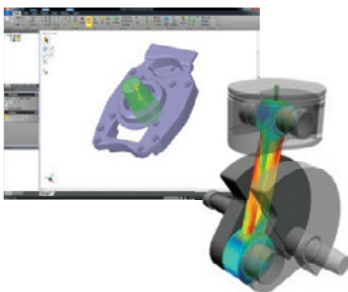
PlotWave Large-Format Print Systems are Cloud Integrated

New version of Océ PRISMPrepare all-in-one document preparation software also released.

The people at Canon Solutions America recently announced their new Océ PlotWave 340 and Océ PlotWave 360 large-format monochrome printing systems. Print speeds range from 10 linear feet per minute with the Océ PlotWave 340 to 13 linear feet per minute with the Océ PlotWave 360.

They both have a new 10.4-inch LCD touch-panel user interface called the Océ ClearConnect that allows for walk-up operation. You can print from a USB drive, your desk, a smartphone app, and print or scan via the cloud.

MORE → deskeng.com/articles/aabkwa.htm



3D Systems Releases Geomagic 2014

New simulation and analysis tools enable engineers to prove out designs.

3D Systems has released the 2014 version of its Geomagic suite of software tools for creating and 3D printing CAD designs, scanning objects, and inspecting articles.

A key element to keep in mind is that the modules in the suite are integrated. The main applications are Geomagic Design X for getting legacy scans into CAD; Geomagic

Design Direct, which integrates reverse engineering tools; Geomagic Verify for first part inspection; Geomagic Control, a 3D metrology solution; and Geomagic Freeform, a design platform. Making their debut are two new modules: Dynamics for Geomagic Design and Simulate for Geomagic Design.

MORE → deskeng.com/articles/aabkxx.htm



Open Platform for Embedded Control and Monitoring

National Instruments redesigns its CompactRIO platform for advanced systems.

The cRIO-9068 is a fully redesigned controller that maintains full backward compatibility with NI LabVIEW and I/O compatibility with the CompactRIO platform. The big thing is that cRIO-9068 features the Xilinx Zynq-7020 All Programmable system on a chip (SoC). The Zynq combines a dual-core ARM Cortex-A9 processor and Xilinx 7 Series FPGA fabric on

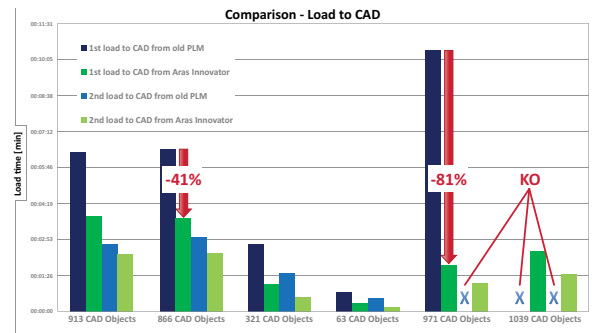
a single chip. And, finally, the cRIO-9068 controller also introduces the NI Linux Real-Time operating system. Yes, Linux and its community of possibilities.

NI says the cRIO-9068 performs up to four times faster than previous generations of the CompactRIO platform.

MORE → deskeng.com/articles/aabkzd.htm

PLM: Migrate to Modernize

Chart 1: Performance Measurement - Existing PLM System vs. Aras Innovator



Global heavy-duty truck manufacturer replaces legacy PLM and accelerates performance for large-scale multi-CAD management.

BY ROLF LAUDENBACH

Founded in 1850, the F.X. MEILLER Company is a global provider of heavy-duty trucking systems for the construction, waste management and commercial vehicle industries. To optimize the matching of vehicle chassis and body, the company develops its products in parallel as a system solution.

MEILLER continually improves its product development and supply chain processes, which recently caused the company to recognize the need to modernize its global PLM environment.

For several years MEILLER had been running a leading large-scale PLM system. The Unix-based system spanned three servers with 14 CPUs and 50GB of memory, making it challenging and expensive to maintain. Although the system managed MEILLER's extensive library of CATIA V5 CAD data, its high resource demands and operational complexity caused significant performance issues. The system offered little additional product or process automation capabilities for an increasingly global product development environment.

MEILLER wanted to improve performance, reduce costs and migrate away from the legacy Unix-based PLM system onto a PLM platform with a modern Web architecture on Microsoft. It also wanted to add new functionality to further streamline and improve global product development processes.

Working with T-Systems, a division of Deutsche Telekom and leading global IT systems integrator, MEILLER conducted an extensive review of currently available PLM solutions and performed benchmark prototypes. Aras was chosen for its modern Web-based platform and superior performance during benchmark scalability testing. Aras also offered greater functionality with a much smaller footprint, resulting in better cost efficiencies.

Under the guidance of T-Systems, MEILLER began an aggressive six-month implementation, migration and integration PLM project. The plan included: Basic PLM Platform, Engineering PDM, Document Management, Change Management, ERP Integration, Production Operations, Quality Compliance and Project Management.

Aras was installed, rights and roles were defined, and lifecycle

and workflow capabilities were enabled. By installing Aras and moving to an all-Windows Server-based Web environment, MEILLER reduced its server load from three servers to one and went from 14 CPUs down to four with greater scalability, performance and manageability.

The Engineering PDM phase included CAD integration to CATIA V5. During this phase, 2.8 million records were migrated from the legacy PLM system to Aras. This included 350,000 CATIA V5 CAD documents representing 650GB of 3D and 2D structured CAD data. Once complete, MEILLER's database was decreased in size from 32GB to 7GB, and the amount of memory required by the system went from 50GB to down to 20GB while performance improved.

T-Systems took the lead on CAD integration and data migration, system configuration and the implementation of key backend processes. Due to Aras's advanced Web architecture and ease of use, MEILLER was able to handle a large portion of the remaining project work themselves, which significantly reduced the overall project cost.

Today, MEILLER has eliminated the Unix / AIX infrastructure and consolidated its global PLM environment onto a single, modern platform running in a pure Windows and SQL Server Web environment. The company has more PLM functionality with faster performance. At the same time, MEILLER has reduced its PLM infrastructure maintenance costs and complexity while gaining the ability to customize Aras themselves.

Rolf Lauenbach is Director of Partner Services & Consulting at Aras.

Download the complete Performance Benchmarking report at aras.com/plm/002070.



THE ZOMBIE COMPUTER SURVIVAL GUIDE

Here are some things to consider when shopping for HPC for simulation.

BY KENNETH WONG

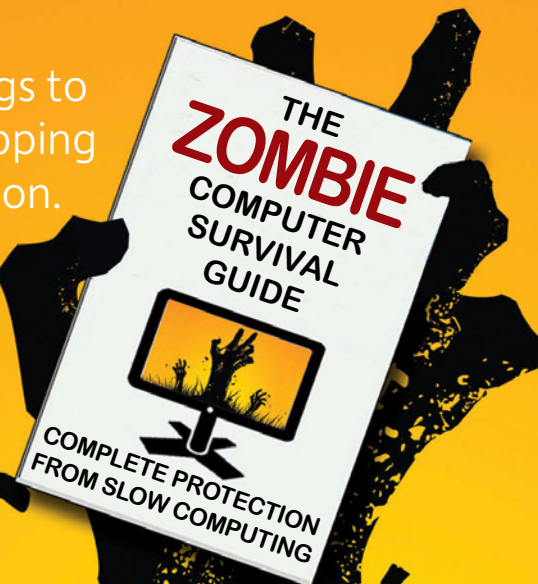


Image courtesy of Istock.

Sooner or later, advanced CAE software user will have to face the invasion of the zombie desktop computers. It's a common anxiety among those who rely on underpowered desktop computers to run complex simulation jobs. Fully engaged in number crunching, the computers come to almost a standstill, not quite dead but barely alive.

In a recent survey *DE* conducted on behalf of IBM Platform Computing, 39% of more than 1,000 respondents stated they're

running their applications exclusively on workstations. Only 11% are running their programs exclusively on clusters. The same participants also reported running a variety of simulation: complete systems (60%), large assemblies and models (58%), and parts and small assemblies (52%). The disparity between the volume of simulation executed and the number of clusters involved suggests a significant segment of the participants may have accepted the zombie workstation syndrome as a fact of life.

Two easy methods will prevent the return of the zombies:

1. Set up an in-house cluster to tackle the CAE jobs and free up the workstations; or

2. Offload them to a remote cluster or on-demand cloud computing vendor.

Our chat with server experts, custom cluster vendors and cloud computing pioneers sheds light on the sometimes not-so-obvious choices you must make with each option.

Open Source or Commercial Code?

In the survey, respondents listed what they felt were barriers to cluster deployment. The top ones include lack of budget for hardware and software (47%), lack of skilled IT staff (31%), and uncertainty associated with migrating desktop applications to the cluster environment (24%).

Budget concern, especially among smaller and midsize firms, may nudge high-performance computing (HPC) buyers toward open-source software. After paying for the hardware, the open-source cluster management software's \$0 price tag is attractive.

"We're big proponents of open-source cluster management for general HPC and select simulation cluster users," notes Brett Newman, Microway's HPC sales and marketing specialist. "We develop our Microway Cluster Management Suite package based upon Ganglia. It's robust software, and it's free with Microway clusters."

But Nick Werstiuk, a product line executive with IBM Platform Computing, points out that open source software is not exactly turnkey: "Open source generally requires [users] to pull together a variety of software, which they must integrate or support on their own. The expertise you need to integrate or pull together these solutions is something you need to consider."

Whereas cluster-management software exists both as open source and commercial code, the job-scheduling option may be best confined to commercial solutions sanctioned by the simulation software vendor. "ANSYS, for example, supports only Platform LSE, PBS Professional, and Windows HPC as third-party schedulers," Microway's Newman says, adding that some open-source schedulers, such as Torque or Grid Engine, "aren't wise choices for ANSYS users."

The Emergence of On-demand Clusters

Software-as-a-Service (SaaS) was yesterday's discussion. Today, the heated debate revolves around Infrastructure-as-a-Service (IaaS) or Platform-as-a-Service (PaaS). The emerging market is served by vendors who offer hardware resources augmented with cluster- and queue-management software as an integrated bundle to process simulation jobs. Like SaaS vendors, IaaS and PaaS vendors offer their products as on-demand solutions, remotely accessible to subscribers who pay a recurring fee.

STAR Global Conference 2014 MARCH 17 - 19



SIMULATING SYSTEMS



REASONS TO ATTEND THE BIGGEST CAE EVENT OF THE YEAR:

NETWORK WITH EXPERTS

The STAR Global Conference is one of the world's largest gatherings of simulation experts, it is also your opportunity to meet with CD-adapco Management, Software developers and Support engineers

LEARN

Over 50 leading industrial experts reveal the secret of their engineering success, exploring their simulation process and demonstrating in detail the tangible benefits of using STAR-CCM+

POWER UP YOUR ENGINEERING SKILLS

Build your own training day agenda, choose from 11 workshops, each of which is a deep-dive into a specific area of engineering simulation



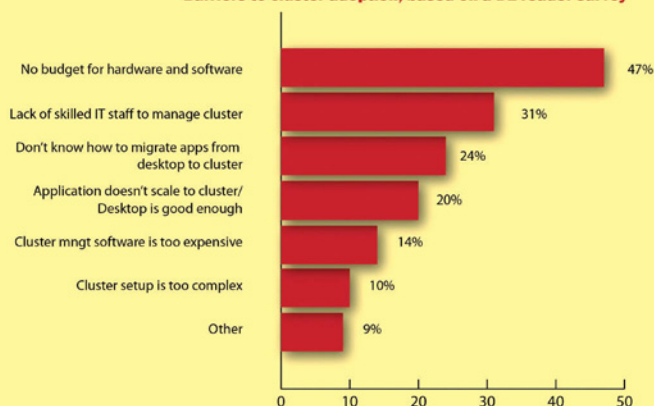
✉ info@cd-adapco.com

🌐 www.cd-adapco.com



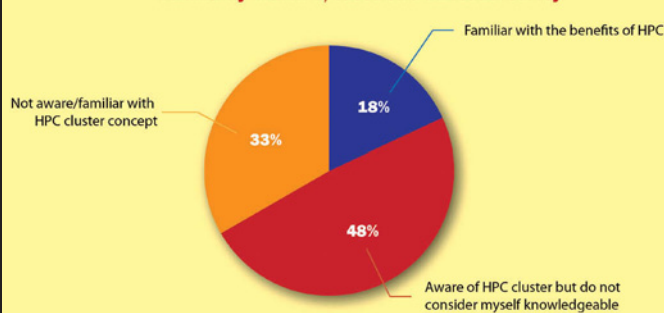
STAR GLOBAL CONFERENCE 2014 - AUSTRIA - REGISTER AT WWW.STAR-GLOBAL-CONFERENCE.COM

Barriers to cluster adoption, based on a DE reader survey



Respondents listed lack of budget for hardware and software as the top barrier to cluster adoption, followed by the lack of skilled IT staff to manage the cluster.

Familiarity with HPC, based on a DE reader survey



Only 18% of respondents stated they are familiar with the benefits of HPC.

Rescale, headed by former Boeing structural and software engineer Joris Poort, is among the new crop of vendors targeting the CAE market with PaaS offerings. The company certainly gets a fair amount of business from those who have no desire to invest in and maintain HPC hardware onsite. But Rescale has also found a new type of simulation users: those seeking scalability.

"Most of our customers actually have on-premise HPC clusters. Some of them say the reason they came to Rescale was because their businesses tend to peak and drop, but their onsite HPC capacity is flat," says Sunny Manivannan, Rescale's vice president of business development. "When their demand goes over the supply, they want to be able to tap into Rescale's infrastructure."

It's not just startups that are dipping their toes into the uncharted waters, however. Altair, an established name in simulation, has repackaged its CAE platform products, HyperWorks, as a cloud offering. Altair customers purchase pools of HyperWorks Units, which function like tokens, to get access to select Altair software titles. With HyperWorks On Demand (HWOD), customers may also use these units to pay for access to Altair's computing resources, available on-demand. Based on Altair's PBS Works suite, the integrated software lets you remotely submit, monitor and manage jobs from the web.

Linux or Windows?

The choice for cluster management software may also be dictated by the preferred cluster operating environment. Most workstation versions of CAD and CAE software are developed for Windows OS. This may lead first-time cluster buyers to conclude that supplementing the pre-existing Windows workstations with a Windows cluster makes the most sense. That, however, may not be the best way to decide the cluster's operating system.

"Some simulation software packages support more features in their Windows or Linux versions," says Microway's Newman. "We strongly recommend that customers consult their hardware and simulation software vendors to assess what's best for their needs."

"[Windows HPC Server] offers robust support for Windows hardware and applications, but not for Linux applications," notes IBM's Werstiuk. "A lot of these CAE applications do have a large Linux footprint, and clients want the ability to run those on Linux environment."

Rescale's own cluster setup offers additional clues. "When we started, we were pure Linux. Now we're looking at Windows options, too, because some small proprietary codes are written for Windows," says Manivannan. "Big commercial codes are written for both Linux and Windows, with occasional differences in features based on OS. Linux is more widely accepted as the standard for larger batch jobs."

Balancing Core and Node Counts with Interconnect Speed

If you throw in an excessive number of server units (computing nodes) in the cluster, the hardware suffers from increased chatter — the back-and-forth communication among the computing nodes to coordinate and attack your simulation job. If your interconnects, the pipelines that join the computing nodes, are not up to par, they could become a performance bottleneck.

"Purchasing a very small cluster of 2P [dual CPU] servers to run modest jobs can sometimes be a poor choice," says Newman, offering 32 total cores as an example. "There's latency involved in the internode communication. Small clusters are most likely to rely on Gigabit Ethernet with its much higher latency as well."

"CAE software users might have much faster results by forgoing a tiny cluster and purchasing a single big node with 32 cores instead," he says. "The entire model might even fit into memory."

For big CAE jobs, Newman says, there's strong benefits to choosing a fast interconnect: "CAE users want ample bandwidth to transmit job data with as little latency as possible. At up to 56GB per second and less-than-1 microsecond latency, Infini-Band allows for the strongest performance scaling."

The GPU Grid Defense

Graphics processing unit (GPU) maker NVIDIA offers another way to keep the zombies away: Put your machines on a grid. Launched officially during NVIDIA GPU Tech Conference 2013, the hardware premiered as the NVIDIA Grid visual computing appliance (VCA). Essentially, it's a GPU-powered cluster, capable of supporting eight to 16 users.

With this approach, users may use inexpensive laptops and consumer desktops as client devices to remotely access powerful virtual machines, each with its own designated virtual GPU. The virtual GPU will make a difference for those who routinely engage in real-time visualization or GPU-accelerated simulation.

Dawn of the Living

Complex CAE jobs will almost always test the limits of individual workstations. When these jobs are running, the burden on the CPU is significant. The ever-increasing size and complexity of the digital models — a consequence of manufacturers' increased reliance on digital simulation — suggests CAE clusters, both off-site and onsite, are bound to become a way of life for designers and engineers for the foreseeable future.

Like the backend IT, the front-end terminals are evolving, too. In the past, most CAD and CAE users had to be physically seated before their desktops, as they were the only machines powerful enough to perform the tasks. The emergence of remote servers, PC over IP, virtualization and cloud computing cuts the umbilical cord, allowing design software users to perform simulation from a mobile tablet, a phone or an ordinary laptop from a Wi-Fi-enabled café or an airport lounge.

As zombie workstations face their twilight, the formerly desk-bound professionals get a chance to operate beyond the cubicles — in sunshine and daylight. **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.

Free White Paper

IBM Platform Computing and Desktop Engineering have produced "5 Easy Steps to a High Performance Cluster," a white paper that is available as a free download from: <http://goo.gl/xMU7eW>

INFO → Altair Engineering: Altair.com

→ **IBM Platform Computing:** IBM.com/systems/technicalcomputing/platformcomputing

→ **Microway:** Microway.com

→ **NVIDIA:** NVIDIA.com

→ **Rescale:** Rescale.com

→ **DE-IBM "5 Easy Steps to a High Performance Cluster" white paper:** <http://goo.gl/xMU7eW>

→ **DE-IBM survey results download:** www-03.ibm.com/systems/technicalcomputing/platformcomputing/

→ **Podcast interview with IBM's Nick Werstuk:** deskeng.com/virtual_desktop/?p=7433

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

msi®

CERTIFIED BY >>

Adobe Creative Suite

AUTODESK

SolidWorks



THE PROFESSIONALS' CHOICE



**MSI MOBILE WORKSTATIONS
POWERED BY NVIDIA® QUADRO®**

BUILT FOR PROFESSIONALS

Featuring the latest technology from NVIDIA® Quadro® and certified by Adobe® Creative Suite 6, Autodesk, and SolidWorks, MSI's next generation GT70/GT60 Mobile Workstations set the new standard for enhanced performance and reliability.

To learn more about MSI's Mobile Workstations, visit: www.msimobile.com/workstation

AVAILABLE AT SELECT E-TAILERS >>



CUSTOMIZED SOLUTIONS AVAILABLE AT VALUE ADDED PARTNERS >>



Workstation, Cluster or Cloud?

Choosing where to run complex computational models has become a choice driven by more than just costs. Today, many factors come into play when choosing from among the traditional workstation, onsite computing grid or cloud services.

BY FRANK J. OHLHORST

Engineers today have more choices than ever before, especially when it comes to computing. A robust ecosystem of technology solutions that transcend traditional computing barriers have become readily available — and more importantly, affordable. Simply put, engineers can pick and choose where to run their simulations, animations and design applications based upon their needs, not their pocketbooks. Let's take a look at some of the options available today, and what a variety of computing professionals have to say about those options.

Workstation Workhorses

High-performance computing (HPC) workstations have long been residents of engineering and design departments across many types of businesses. These faithful adjuncts have supported CAD/CAM design applications, simulations and visualizations, while also serving as the repositories of complex calculations.

Truth be told, those very workstations have been a boon to engineering productivity, and have made it easier for engineers to delve into new frontiers. What's more, those engineering workstations have evolved — gaining more processing power while dropping in price. The result are advanced capabilities that are only slightly more expensive than traditional desktop PCs, allowing workstations to take on many advanced engineering roles.

Jon Wells, senior designer at the Morgan Motor Co., a British automobile manufacturer, notes that “the power of the latest generation of workstations are allowing us to do more and more design work from the desktop, eliminating the need for expensive, dedicated CAD/CAM systems.”

What's more, the growth in processing power and the



Modern desktop and mobile workstations are powerful enough to handle most of the tasks design engineers throw at them. Image courtesy of Lenovo, screen capture courtesy of Autodesk.

compact nature of the latest generation of workstations bring additional advantages with them, such as scalability and portability. Art Thompson, vice president of Sage Cheshire Airspace and leader of the team responsible for Felix Baumgartner's 24-mile skydive last year, reports that “most of our day-to-day operations are performed on workstations, especially since we can pack those up and bring them out into the field when needed for last-minute simulations or to capture data.”

With the latest enhancements and the incorporation of solid-state drives (SSDs), caches and ultra-high perfor-

mance video cards, Thompson adds that his team is finding that “we can standardize on a given vendor’s workstation to process our workloads and eliminate the need for specialized hardware.”

Going On-grid

Nevertheless, there are scenarios where underpowered workstations won’t cut the mustard, leaving engineers to seek alternatives that deliver even greater processing power. This has spurred a growing interest in grid computing, which is a way to enlist large numbers of machines to work on multipart computational problems such as circuit analysis or mechanical design.

“Certain simulations are beyond the scope of the typical workstation and require banks of CPUs to deliver a result, which is where the grid concept comes into play,” Thompson explains.

Morgan’s Wells agrees. “Some simulations, such as aerodynamics, mechanical component stress testing and 3D animations require much more power than a workstation can deliver,” he adds.

Leveraging grid computing has garnered significant attention among scientists, engineers and business executives. Grid computing excels at solving complex mathematical

problems, and is a technique that is one of the latest developments in computing, which has already delivered such advances as distributed computing, collaborative computing and the Web.

However, grid computing is not all that new. In fact, it has been in use for several years — allowing businesses to discover there can be high equipment and operational costs associated with the technology. In the past, a grid-based system may have been the only way to solve certain engineering problems, but others are now turning to hosted and scalable resources to maximizing productivity, while minimizing costs.

Grant Kirkwood, CEO of Unitas Global, a cloud services and hosting firm, gives a good example of seeking alternatives: “We have a film studio as a customer, which had set up a grid for animation and FX work in their studio. The grid grew in size, expense and created several heat and power problems. The film studio switched over to our hosted offering, and [as a result] eliminated the problems they were encountering and gained instant scalability.”

Perhaps scalability is only one of the keys to leveraging Platform-as-a-Service (PaaS) or Infrastructure-as-a-

Continues on page 28 ...

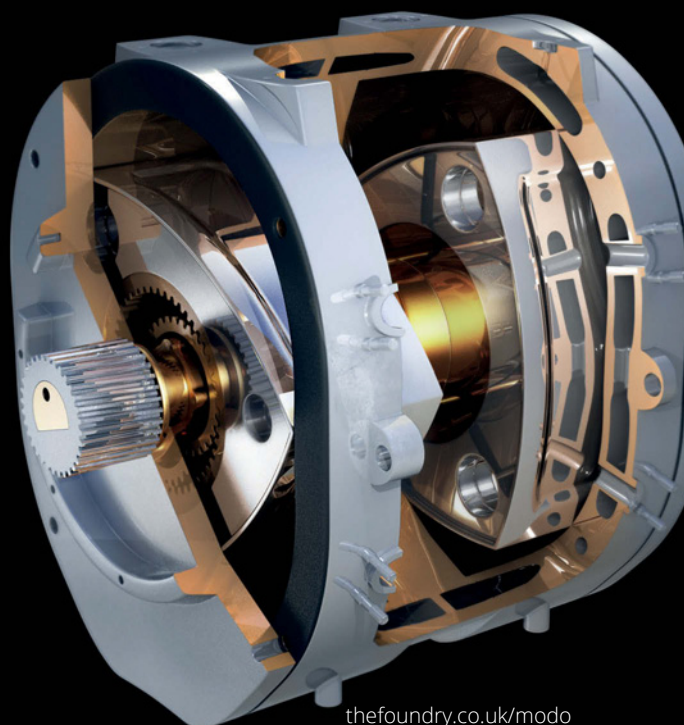


A creative tool for technical people

MODO is a true end-to-end 3D software solution that combines modeling, sculpting, shading and lighting, and rendering functionalities in a single integrated package.

**THE
FOUNDRY.**

Image courtesy of Igor Kachanov



thefoundry.co.uk/modo



Pros don't use computers.

They use workstations built for professionals.

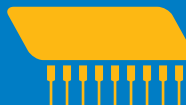
You're at the top of your game – so why let underpowered computing tools slow you down? It's time to accelerate your productivity beyond the PC with the pro-level performance of a workstation powered by the Intel® Xeon® processor.

You carefully balance your life to deliver consistent creative output ... and when you balance your investment in a professional workstation, you have the capabilities to make great things happen.



Processing

Accomplish more with accelerated simulation and render times, and get a burst of extra power when you need it.



Memory

Decrease your time to solution by investing in the amount of ECC memory you need.

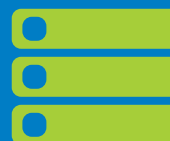
Graphics

Use professional, certified graphics that meet your CAD/CAE needs.



Storage

Maximize throughput by specifying solid-state drives.



When your most complex work demands a more capable machine, unlock your innovation with the reliability, stability and performance of a workstation based on Intel® architecture, and built for ... you.

Visit intel.com/go/workstations

... Continued from page 25

Service (IaaS) — at least that seems to be a common belief.

“German car company BMW recently moved its HPC operations to our Keflavik, Iceland, data center for a multitude of reasons,” notes Jeff Monroe, CEO of Verne Global, an international data center operator. “The most significant reason is that BMW expects to save around 80% of the power costs of running calculations, including crash test and aerodynamics simulations, as well as CAD/CAE calculations by using our data center.”

Echoing that thought is Morgan’s Wells, who notes that “offshoring HPC makes a lot of sense, as long as there are tangible savings and no latency problems.”

But savings may only be part of the story. BMW tested the network connections from Munich to Iceland, and Monroe reports that “the test results were a critical factor in their decision to place production systems in Iceland.” The move was also predicated by emissions concerns. With a big surplus and reliable long-term supplies of renewable energy, Iceland’s utilities offer very cheap deals and long-term contracts. Monroe says this is one of Verne’s core competitive advantages — and prices are guaranteed.

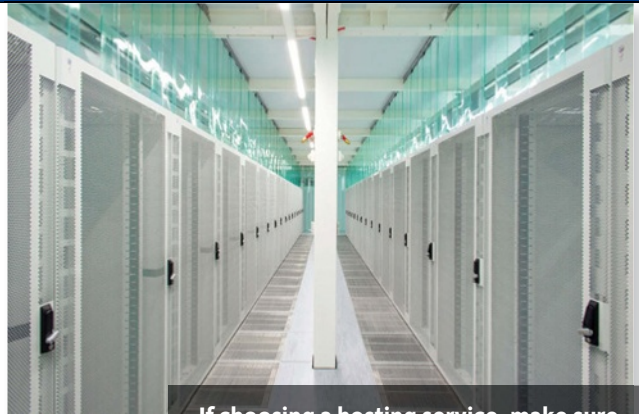
“We can offer customers a low, inflation-protected rate for up to 20 years,” he reports, noting that it’s a significant consideration, “in light of rising long-term electricity costs in Europe, the UK and US.”

Head to the Clouds?

Cloud computing seems to be evolving into an HPC user’s dream, by offering a reasonable metered cost, which includes unlimited storage and instantly scalable computing resources. Nevertheless, HPC cloud offerings require extensive due diligence, simply because remote HPC services can be based upon shared HPC clusters, hybrid cloud offerings, fully virtualized cloud environments or other technological combinations.



Make sure uptime is guaranteed by a host’s use of important subsystems, such as onsite generators shown here.



If choosing a hosting service, make sure the facility is neat, clean and secure.

“Before moving to cloud-based HPC, one has to consider latency, scalability and usage requirements,” advises Unitas Global’s Kirkwood. “Luckily, several tools exist to vet those concerns — and of course, there is always a SLA (service level agreement) that spells out the expectations of the service, holding the provider accountable for failing to meet service goals.”

The tangibility of on-desk HPC can sometimes trump cloud and grid offerings, according to Philip Ra, vice president of Yazdani Studio, an architectural and design firm.

“Working with CAD/CAM designs in real time in a conference room environment proves to be a powerful capability that fuels ideas and enhances the customer experience,” Ra notes. “With that in mind, portability becomes a major concern and we leverage portable workstations to make that a reality.”

That said, there is an extensive ecosystem of hosts that are ready to customize HPC offerings to fit the needs of any given company. Take, for example, Open Data Centers, an organization that offers carrier-neutral co-location options.

“Our ability to continuously evolve our data center architecture and control processes allows us to meet the ever-changing demands of customers,” says Open Data Centers’ CEO Erik Levitt. “With 8,500 sq. ft. of scalable data center space, a 24-hour, on-site Network Operations Center, and N2 infrastructure, we can offer the choice, flexibility and responsiveness of a more personalized data center.”

While Levitt’s comments firmly fit under the guise of marketing, he does make a reasonable point. Today’s hosts are more than willing to build custom offerings on top of existing infrastructures, helping to shift the provisioning and management of HPC to an external resource.

Nonetheless, the question still remains: Should engineering firms invest in workstations, grid computing or hosted offerings? While there is no easy answer to that question, there are several rules of thumb that can make the selection process more navigable. Questions to consider include:

- What applications need to be supported? For exam-

ple, CAD/CAM applications such as AutoCAD, BricsCAD, IntelliCAD and several others are designed for workstations running Microsoft Windows. Other applications may run under Linux, Solaris and so on, which has an impact on the type of computing environment needed.

- What type of output is expected by the application? Will the output be used to drive 3D printers, computer numerically controlled (CNC) equipment, plotters or presentation systems?

- How much processing power is needed? Does the processing power need to scale occasionally, frequently or never?

- Which fits the business model better, capital expenses or operational expenses? Each has its pros and cons, and are often decided on a project-by-project basis.

- Is there a baseline configuration used for each and every project? Or do HPC requirements change on a project-by-project basis?

- Is there sufficient staff on hand to support the computing environment? Grids need maintenance, and workstations need management. Is the business capable of handling those needs internally?

- What controls need to be put in place to guarantee uptime, meet business continuity needs or support disaster

recovery plans? Some businesses can survive a few hours of downtime, while others must have continuity. Answers to those questions will drive the design of the infrastructure.

- Do offsite operations need to be supported? Will engineers be working in the field? Will site offices be established? This drives the decision of whether a compute solution needs to function in isolation or requires some type of connectivity.

While the above considerations are just a fraction of what a complete computing design plan should include, they do offer a basic guideline that should help to narrow down what works best for any given HPC consideration. **DE**

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

INFO → The Morgan Motor Co.: Morgan-Motor.co.uk

→ Open Data Centers: OpenDataCenters.net

→ Sage Cheshire Airspace: SageCheshire.com

→ Unitas Global: UnitasGlobal.com

→ Verne Global: VerneGlobal.com

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

OUT NOW!

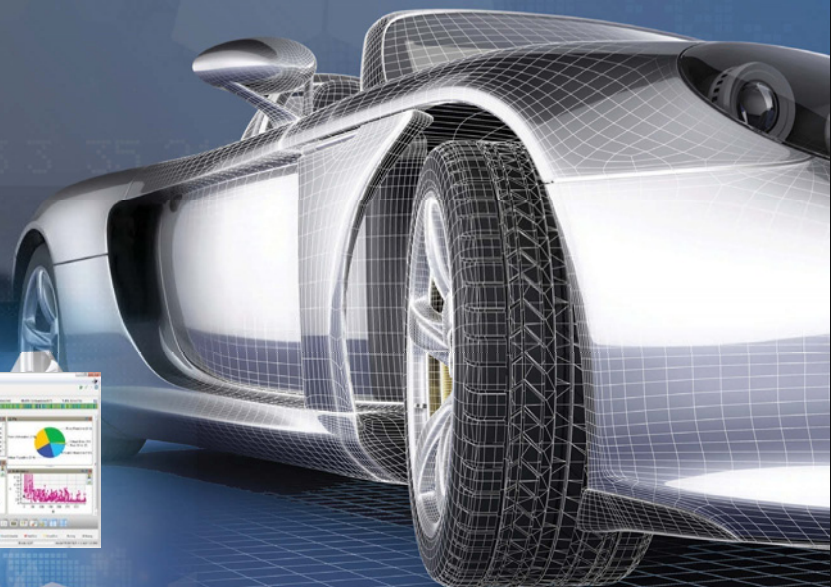
MORE
USABILITY
CONTROL
FLEXIBILITY
CONFIDENCE

www.esteco.com



mode FRONTIER
VERSION 4.5

EXPERIENCE THE ADVANTAGES
OF THE NEW RELEASE



An **Optimized** Single-socket Solution

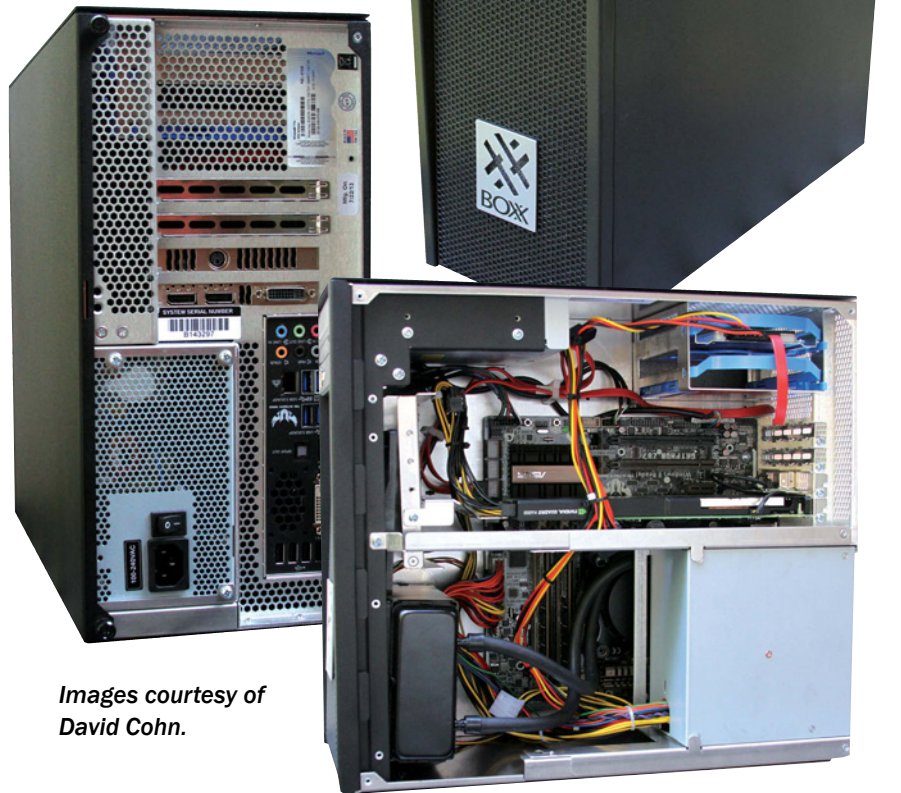
3DBOXX 4150 XTREME, the latest single-CPU workstation from BOXX Technologies, sets a new price/performance standard.

BY DAVID COHN

We're always excited to receive new workstations from BOXX Technologies. The Austin, TX-based company has been building computers since 1996, and its systems have consistently set new high-water marks for performance. The dual-CPU 8980 XTREME we looked at this summer (*DE*, July 2013) proved to be the fastest and most expensive system we've ever tested. In contrast, the new 3DBOXX 4150 XTREME is a much more affordable solution.

Like all of the other 3DBOXX workstations we've reviewed previously, the 3DBOXX 4150 XTREME came housed in a custom-designed aluminum chassis. But everything about the 4150, starting with the box it arrived in, was sleek and compact. Like its larger sibling, the 4150 XTREME's case is all black — except for the brushed aluminum BOXX logo on the removable front panel and the matching logo cut-out on the top of the case. That front grille conceals a pair of 4-in.-diameter cooling fans. The case itself measures just 6.85x16.6x14.6 in. (WxDxH), and weighs 19 lbs. This makes it much smaller and lighter than the 3DBOXX 3970 EXTREME, the last single-CPU BOXX workstation we reviewed (*DE*, January 2012).

In spite of the physical changes, the front panel still houses a single LG Electronics 20X super-multi DVD-RW



Images courtesy of David Cohn.

optical drive, as well as a panel containing two USB 3.0 ports, two USB 2.0 ports, audio jacks for headphone and microphone, a round power button with bright-white LED power indicator, a blue hard drive activity light, and a small reset button. But the smaller case does not provide any additional front panel drive bays.

A well-designed rear panel houses four additional USB 2.0 ports, four

more USB 3.0 ports (including one that allows for updating the system BIOS), an RJ45 network connection for the integrated Intel I217V Gigabit LAN controller, and both DVI and HDMI ports for the Intel CPU's integrated graphics. There are also six audio connectors (microphone, line-in, line-out, side, rear, and center/subwoofer), as well as an optical S/PDIF Out port.

Compact Interior

Removing the right side panel reveals a compact, but well-organized interior. In addition to the single drive bay with front panel access, there are two internal drive bays. In our evaluation unit, one of those bays contained an Intel 240GB SSD drive, adding \$250 to the base price. BOXX offers other drive options, however, including standard 7,200 rpm SATA drives ranging from the 500GB drive in the base configuration to a 4TB hard drive. The system's integrated drive controller supports up to six SATA drives, as well as redundant array of independent disks (RAID) 0, 1 and 10 configurations.

A 550-watt power supply with Gold-level efficiency nestled in the bottom rear of the case provides ample power for any expansion needs — but covers more than half of the ASUS Gryphon Z87 motherboard. That motherboard supports Intel's latest, fourth-generation CPUs and is based on a Z87 chipset.

We could just barely see the Intel Core i7-4770K CPU beneath its closed-loop cooling unit. A pair of black rubber hoses extend from the CPU's heat sink to a radiator case mounted behind the lower portion of the front grille. Those rubber hoses, as well as a stiffening panel running the full depth of the case, also make it a bit more difficult to access the motherboard's memory sockets.

The ASUS motherboard provides four memory sockets supporting 240-pin dual in-line memory modules (DIMMs). While the base 3DBOXX 4150 XTREME configuration comes with just 4GB of RAM, our evaluation unit was equipped with 16GB of memory installed as a pair of 1600MHz 8GB DDR3 modules, leaving room to expand system memory to its 32GB maximum. That extra memory in our evaluation unit added \$382 to its price.

New Intel CPU

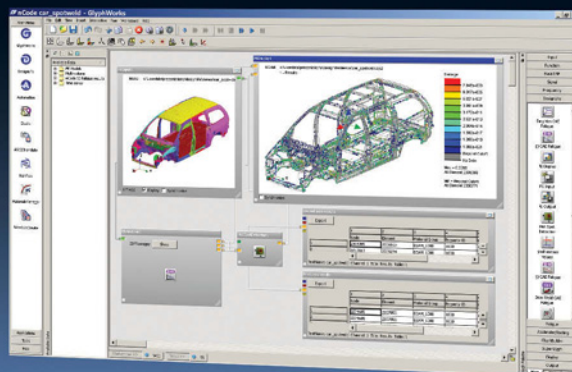
The 3DBOXX 4150 XTREME comes with an Intel Core i7-4770K quad-core CPU. This new "Haswell" CPU is the successor to Intel's "Ivy Bridge" architecture. Officially announced on June 4, 2013, these new processors use a 22nm process and provide a 6% increase in the number of instructions per clock, resulting in 6% faster single- and multi-threaded performance.

In its stock configuration, the CPU has a maximum turbo boost frequency of 3.9GHz while maintaining a maximum thermal design power (TDP) rating of 84 watts. But as we've come to expect, BOXX increases performance by over-clocking the CPU to a maximum turbo boost frequency of 4.3GHz. Our research indicates that the overclocked i7-4770K runs 10 to 15 degrees hotter than comparable over-clocked Ivy Bridge CPUs, but our evaluation unit remained cool and quiet

Achieve new standards through finite element based fatigue analysis.

As finite element models increase in size and fidelity, **HPC environments** are being used to improve simulation throughput. The largest computational problems can be broken down and solved many times quicker, but hardware is not the only requirement.

nCode DesignLife™ can take advantage of this scalability by distributing fatigue analysis tasks to open up the possibility of more **robust, up-front design through simulation.**



nCode DesignLife™

Learn more about **Distributed Processing and High Performance Computing** in nCode DesignLife at www.ncode.com/webinars

www.ncode.com | info@hbmncode.com

nCode® 
...an HBM brand

even under heavy computational loads.

Although the ASUS motherboard provides a pair PCIe 3.0 x16 graphics slots, as well as a third PCIe x16 slot with version 2.0 support, ASUS recommends using just a single graphics card with this motherboard (although the board does support CrossFireX and SLI). There is

also a single PCIe 2.0 x1 slot.

The base 3DBOXX 4150 configuration comes with an NVIDIA Quadro K600 graphics card. For us, BOXX included a more powerful Quadro K4000 GPU with 3GB of GDDR5 memory and 768 compute unified device architecture (CUDA) parallel processing

cores. That upgrade added \$891 to the overall system price.

BOXX also offers other graphics options, including dual NVIDIA GeForce cards, Quadro Tesla boards, and AMD FirePro. Customers could also save money by opting to use only the integrated Intel HD graphics provided by

Workstations Compared

| | | Single-Socket Workstations | | | | Dual-Socket Workstations | |
|----------------------------------|---------|---|--|--|--|--|---|
| | | BOXX 3DBOXX W4150 XTREME workstation (one 3.5GHz Intel Core i7-4770K quad-core CPU over-clocked to 4.3GHz, NVIDIA Quadro K4000, 16GB RAM) | Ciara Kronos 800S workstation (one 3.5GHz Intel Core i7-2700K quad-core CPU over-clocked to 5.0GHz, NVIDIA Quadro K5000, 16GB RAM) | Lenovo E31 SFF workstation (one 3.3GHz Intel E3-1230 quad-core CPU [3.7GHz turbo], NVIDIA Quadro 400, 8GB RAM) | Lenovo S30 workstation (one 3.6GHz Intel Xeon E5-1620 quad-core CPU [3.8GHz turbo], NVIDIA Quadro 4000, 8GB RAM) | BOXX 8980XTREME workstation (two 3.1GHz Intel E5-2687W eight-core CPUs over-clocked to 3.82GHz, NVIDIA Quadro K5000, 64GB RAM) | HP Z820 workstation (two 3.1GHz Intel Xeon E5-2687W eight-core CPU, NVIDIA Quadro 5000, 32GB RAM) |
| Price as tested | | \$4,273 | \$5,714 | \$1,093 | \$2,614 | \$13,454 | \$9,984 |
| Date tested | | 7/31/13 | 5/31/13 | 12/29/12 | 8/18/12 | 5/9/13 | 7/16/12 |
| Operating System | | Windows 7 | Windows 7 | Windows 7 | Windows 7 | Windows 7 | Windows 7 |
| SPECview 11 | higher | | | | | | |
| catia-03 | | 72.37 | 96.39 | 18.15 | 48.21 | 78.01 | 51.69 |
| ensight-04 | | 49.20 | 83.26 | 11.08 | 32.18 | 80.25 | 44.13 |
| lightwave-01 | | 100.78 | 103.15 | 46.79 | 64.47 | 77.07 | 59.02 |
| maya-03 | | 131.31 | 153.01 | 40.36 | 84.50 | 125.16 | 101.67 |
| proe-5 | | 24.74 | 22.87 | 10.29 | 11.93 | 16.14 | 11.72 |
| sw-02 | | 78.27 | 84.51 | 31.54 | 53.53 | 67.16 | 57.48 |
| tcvis-02 | | 55.73 | 77.82 | 16.53 | 37.66 | 71.58 | 44.52 |
| snx-01 | | 53.95 | 83.21 | 13.25 | 33.87 | 81.35 | 44.86 |
| SPECapc SolidWorks 2013 | Higher | | | | | | |
| Graphics Composite | | 5.25 | 3.89 | n/a | n/a | 2.69 | 2.15 |
| RealView Graphics Composite | | 5.38 | 4.1 | n/a | n/a | 2.86 | 2.37 |
| Shadows Composite | | 5.36 | 4.1 | n/a | n/a | 2.86 | 2.36 |
| Ambient Occlusion Composite | | 5.63 | 8.37 | n/a | n/a | 6.26 | 5.19 |
| Shaded Mode Composite | | 5.12 | 3.79 | n/a | n/a | 2.62 | 2.27 |
| Shaded With Edges Mode Composite | | 5.38 | 3.98 | n/a | n/a | 2.77 | 2.03 |
| RealView Disabled Composite | | 4.74 | 3.15 | n/a | n/a | 2.11 | 1.45 |
| CPU Composite | | 4.07 | 4.92 | n/a | n/a | 4.84 | 4.5 |
| Autodesk Render Test | Lower | | | | | | |
| Time | Seconds | 42.00 | 58.33 | 64.00 | 63.80 | 38.00 | 41.00 |

Numbers in blue indicate best recorded results. Numbers in red indicate worst recorded results. Results are shown separately for single- and dual-socket workstations.

the Core i7 CPU. But while tests indicate that the integrated HD4600 graphics performance is approximately 20% faster than the HD4000 graphics included in the older Ivy Bridge CPUs, intensive engineering computing performance is better using a discrete graphics board rather than the integrated graphics.

Record-setting Performance

We always expect great performance from BOXX workstations, and can report that the 3DBOXX 4150 XTREME delivered. On the SPECviewperf test, which focuses only on graphics performance, the 4150 did extremely well, scoring near the top on almost every test. It even out-performed the Ciara Kronos 800S workstation we recently reviewed (DE, August 2013), even though that system was equipped with a more powerful GPU.

On the SPECapc SolidWorks 2013 benchmark, the 4150 XTREME beat every system we've ever tested on most portions of this test, lagging only behind the Ciara workstation and the much more powerful 3DBOXX 8980 XTREME on several of the composite scores.

On the AutoCAD rendering test, which clearly shows the advantages of CPU speed, multiple CPU cores and hyper-threading, the 3DBOXX 4150 XTREME completed the rendering in 42 seconds. This is a record for a system equipped with just a single CPU.

BOXX rounded out the system with a Logitech K120 104-key USB keyboard and M500 USB laser mouse. Windows 7 Professional 64-bit came preloaded. Windows 8 and Linux are also available. BOXX Technologies backs the system with a 1-year premium warranty, 24/7 telephone support and next-business-day onsite service, followed by two additional years of standard warranty service. Premium support can be extended for the second and third years at the time of purchase for an additional charge.

Not only does the 4150 XTREME pack plenty of performance, it does so at an attractive price: The 3DBOXX 4150

XTREME has a starting price of \$2,800, which gets you the over-clocked Core i7-4770K CPU, 4GB of RAM, NVIDIA Quadro K600 graphics, 500GB 7,200 rpm SATA drive, 20X DVD-RW drive, and Windows 7. As configured, our evaluation unit priced out at \$4,273, making it the new price/performance leader.

Although the 3DBOXX 4150 XTREME is more expensive than entry-level workstations, it represents an optimized solution for CAD/CAM and other engineering applications that don't really benefit from more expensive, multi-CPU configurations. **DE**

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

INFO → **BOXX Technologies, Inc.:**
BOXXTech.com

BOXX 3DBOXX 4150 XTREME

- **Price:** \$4,273 as tested (\$2,800 base price)
- **Size:** 6.85 x 16.6 x 14.6-in. (WxDxH) tower
- **Weight:** 19 lbs.
- **CPU:** one Intel Core i7-4770K (quad-core) 3.5GHz (over-clocked to 4.3GHz in turbo mode)
- **Memory:** 16GB DDR3 at 1,600MHz (up to 32GB supported)
- **Graphics:** NVIDIA Quadro K4000
- **Hard disk:** Intel 240GB SATA SSD (two internal drive bays)
- **Optical:** LG Electronics 20X DVD+/-RW
- **Audio:** onboard integrated high-definition audio
- **Network:** integrated 10/100/1000 LAN with one RJ45 socket
- **Other:** Two USB 2.0 and two USB 3.0 on front panel; four USB 2.0 and four USB 3.0 on rear panel; integrated DVI and HDMI video ports
- **Keyboard:** 104-key Logitech K120 USB keyboard
- **Pointing device:** Logitech USB Laser Mouse

how to: PRINT IN 3D

The only ship today
provider of 3D prints!

1. visit **zoomrp.com**

We're your one-stop-shop for 3D printing. We offer PolyJet, SLA, SLS and ABS.

2. ^{securely!}upload file

You can upload as many files as you're ready to print. Review quantities and prices and select your shipping option for all your parts at once.

3. ^{Place Order}push button

It's really that easy. Our returning customers come back for simple convenience. We take the overhead, the guesswork, and the wait out of production every time.



ZoomRP  .com

FASTER RENDERING AND VISUALIZATION

Today's technologies are making it easier than ever to speed up the design process.

BY MARK CLARKSON

A rendering of a Toyota Scion.
Image courtesy of RTT.

When it comes to rendering and visualization, fast is never fast enough. So how to go ever faster?

The old answer was to buy the fastest processors and the most, fastest RAM you can afford. And that still holds in most cases: Grab a speedy Intel Core i7 processor (or better yet, several of them), coupled with vast quantities of DDR3 RAM, and you can work at a speed that would have made your head explode just a few years back.

Consider KeyShot, which I reviewed back in the February issue. KeyShot is a CPU-only renderer that is very, very fast when running on any semi-current computer. It will give you decent-looking renders in a few seconds. Give it plenty of RAM and plenty of CPU cores, and watch it fly.

Good Ol' CPUs

In fact, says Intel Workstation Segment Manager Wes Shimanek, "most of the high-end ray tracing algorithms out there are being developed on CPUs, because it gives them access to immense amounts of memory and immense amounts of compute capacity.

"Autodesk, for example, has a number of rendering technologies," Shimanek continues. "But their most advanced ray tracing technologies, based on VRED, are all CPU-based. The features and functions you get with a CPU-based ray tracer are superior."

If your CPU isn't up to the task, you can always add an Intel Xeon Phi coprocessor. Intel Xeon Phi coprocessors provide as many as 61 cores and up to 16GB of RAM, for more than a teraflop of performance.

"The same software that runs on your CPU runs over on the Xeon Phi coprocessor," says Shimanek. "From a software developer's perspective, that's great: Write it once, and it runs on mul-

tiples platforms. The CPU and coprocessor both look and smell just like Intel X86 architecture. This gives you greater flexibility."

What's in the Cards?

Caustic's ray tracing acceleration boards speed things up by moving the grunt work of tracing rays onto specially designed processors. Not to be confused with the graphics processing units (GPUs) found on many graphics cards, Caustic's ray tracing units (RTUs) are designed for ray tracing. By using the system CPU and keeping textures in system memory, the boards can render larger scenes than most GPU boards, while consuming significantly less power. (*Editor's Note: For more about Caustic's boards, see "Seeing Ray Tracing in a Different Light" in the April issue of DE.*)

GPU-based cards can make for some screaming-fast renders but, as hinted at above, a key limitation is that GPU-based rendering requires all of your scene's geometry and texture information to be loaded onto the GPU card itself, to avoid passing chunks of data across the painfully slow system bus.

In response, NVIDIA's new Quadro K6000 packs 12GB of RAM onto a single card — triple that of the previous flagship card, the K5000. It's not just more RAM, of course. With 2,880 compute unified device architecture (CUDA) cores on tap, and a 288GB/second memory bandwidth, the K6000 will push more than 2 billion triangles per second onto the screen. But the RAM itself is monumentally useful for loading really large scenes.

No More Data Preparation?

"Because engineering data has back sides, substrates, fasteners, harnesses, cables and hoses, that data is extremely heavy," says Ryan Schnackenberg, head of design and engineering

solutions at RTT. “In the past, we have relied heavily on a process flow called data preparation.”

Schnackenberg recalls how it could be the work of half a day, a day or even more, to organize, tessellate, delete, simplify and otherwise prepare CAD data to fit into available RAM, and to allow the application to manipulate it at a reasonable frame rate.

“Now that we’re seeing more advances in GPU technology, we see a move away from data preparation,” he adds.

At the recent SIGGRAPH 2013, RTT wanted to show some completely unprepared data from the Nissan Motor Co. The “data” was, in fact, an entire Nissan Pathfinder — comprised of more than 40 million polygons. Historically, Schnackenberg says, the team would have had to strip quite a bit out of the model to create something that would work in near real-time. But by using a 12GB K6000 card, RTT was able to load the entire car into its DeltaGen application with very little preparation. The result was a scene with high visual fidelity, running on a workstation in real-time, with full global illumination.

Plus, there are advantages to leaving all the fiddly bits intact rather than stripping them out. “In a design and engineering scenario,” says Schnackenberg, “designers want to peel the layers from the onion; they want to see the wires and the hoses.”

The less you have to change the model along the way, the easier the lifecycle management process is, he adds: “We try to re-leverage the CAD data throughout its lifecycle. As it goes to

manufacturing and then to launch, you can reuse the content again for sales and marketing, online configurators, kiosks, etc.”

Remote Clusters

Another technology on display at SIGGRAPH was a cluster of 232 NVIDIA GPUs. “The K6000 does an awesome job,” says Andrew Cresci, NVIDIA’s vertical marketing general manager. “But what a cluster does is nothing short of astonishing. It’s like having a photograph you can move around in 3D.”

It can be used for rendering, or it can be used for NASTRAN or computational fluid dynamics (CFD). It doesn’t have to be a massive cluster, either, Cresci notes: “A high-end workstation these days can have three or four of our GPUs bolted into them.”

In fact, the cluster at SIGGRAPH was actually running hundreds of miles away, in Santa Clara. “We were remot-ing down to Anaheim,” says Cresci, “running off a workstation at the trade show with full interactivity. In the last nine months or so, remot-ing has really become commercially deployable for the first time.”

By having a centralized compute capacity next to the file server on a high-speed link, you can load data onto the cluster very quickly. The central computer does the heavy lifting. Your access is a relatively low-bandwidth H.264 video stream.

“The visualization comes up quickly because the data loads incredibly fast,” says Cresci. “You get very high performance computing on the central cluster, and you can access it from any-

opera 16

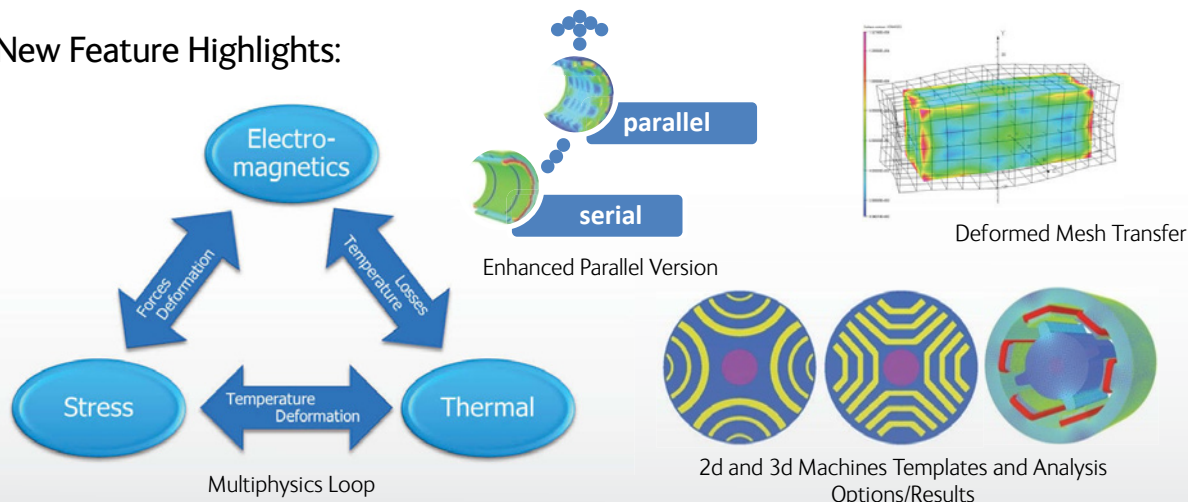
simulation software

COBHAM

The most important thing we build is trust

Magnets • Motors • Magnetic Signatures • Generators • Actuators • Transformers • Lightning Protection Systems
NDT • Magnetic Shielding • Magnetron Sputtering • Charged Particle & Plasma Devices • Superconducting Magnets

New Feature Highlights:



<http://linkd.in/RA7ib0>

For more information visit:
operaFEA.com





Renderings of a rotor (top) and a Ferrari (bottom) from Lagoa, and of an aircraft interior from RTT.



DeltaGen 12 Recently Released

RTT has released an upgrade of its flagship 3D visualization software, DeltaGen 12. According to the company, the new release includes more than 60 enhancements and new features.

It includes new artistic composing and post production features, such as motion blur, lens flare, sun shafts, render passes or flexible real-time light and shadow settings. The new camera model enhances a real photography approach through provision of sensor sizes, focal lengths and shutter speeds.

The new release also helps save time by offering a performance increase and high efficiency with increased reusability of data, the company says. It reduces time-to-asset through unseen performance increases for offline rendering and real-time ray tracing, including Global Illumination (GI).

The company also introduced RTT Xplore DeltaGen, a multi-touch navigation tool that facilitates the presentation set-up and improves scene interaction. RTT DeltaGen 12 for Teamcenter combines 3D visualization with PLM. The new version comes with an enhanced metadata concept, quick update processes and customized queries for search, filter and sorting.

where. The ability of guys to sit in their office and have access to this phenomenal capability, without long load times and crunching times, is huge. And if you go out to a supplier site, you just log on and everything comes up.

"The data never leaves your corporate site. All you've got is a video stream so there's no IP risk."

On a smaller scale, NVIDIA has a sort of cluster-in-a-box called the visual computing appliance (VCA). Plug it in, hook it to the Internet/intranet, and install your software. Now you can run, say, SolidWorks from your desk, from home, from the conference room or from the offices of your outside suppliers.

"You get staggering performance with SolidWorks and RealView," says Cresci.

Cloud Rendering

If this approach is appealing, but you don't want to actually build and maintain your own cluster, you can always go to the cloud. Autodesk's 360, for example, is a cloud rendering solution currently available for AutoCAD, Revit and Fusion.

"We're trying to shift people from just rendering for presentation purposes, when the design is finished, to thinking of rendering as a process that occurs throughout the design," says John Hutchinson, senior software architect and SWD manager. "We spent a good deal of time validating [our renderer's] accuracy from a photometric standpoint. The renderer is accurate, so using it early on in the process can inform many aspects of the design."

MSI Mobile Workstations – Built for Engineers

Multi-monitor capable mobile workstations ready for compute-intensive engineering.



As an engineer, you know that sometimes you have to take work home or to the client. You dread it because your home system and your business notebook

are not up for the job. Introducing the all-new MSI GT70 and GT60 Mobile Workstations from MSI Computer Corp. With the power and flexibility to serve as a multi-monitor desktop workstation running compute-intensive engineering applications, the GT70 and GT60 Mobile Workstations are built and designed so that you can take your engineering workstation with you wherever you need to go.

Based on the award-winning GT Series notebooks, the GT70 and GT60 next generation workstations are fueled by NVIDIA Quadro K3100M and K2100M mobile graphics cards and Intel Core i7-4700MQ processors. Fully certified and optimized for such industry-leading applications as Autodesk, SolidWorks, and Adobe, the GT70 and GT60 leverage NVIDIA Quadro technology empowering you to run design and rendering, structural analysis, or fluid analysis without any downtime or delays.

The GT70 and GT60's Matrix Display technology lets you switch from a mobile workstation to a multi-monitor desktop workstation easily and efficiently. Matrix Display supports up to three external independent displays simultaneously, enabling you to render 3D images, design, and check your reference files at the same time.

The GT70 and GT60 sport a SteelSeries keyboard. Originally designed for gamers, this high-quality, highly accurate backlit keyboard offers great key feedback. Its ability to detect up to 10 simultaneous key hits makes the SteelSeries an ideal keyboard for professionals taking on multiple tasks rapidly.

The GT70 and GT60 Mobile Workstations are also designed with exclusive networking optimizations — the Killer DoubleShot high-performance Ethernet and Wi-Fi networking solution from Qualcomm Atheros. Offering up to 5x better latency than similar solutions, Killer DoubleShot gives you the horsepower to handle multiple in-coming and out-going data streams at once. For superior bandwidth control, Killer DoubleShot's Visual Bandwidth Control lets you see and manage your data streams in real-time.

To learn more about the MSI GT70 and GT60 Mobile Workstations for engineers, go to:
www.msimobile.com/workstation.

msi

With AutoCAD and Revit, the application runs and stores data locally, then shoots a compressed soup of triangles to the cloud for rendering — obfuscating your intellectual property.

"This is shifting as more of our products become 'cloudified,'" says Hutchinson. "Fusion 360 is an example of that. While you do a desktop install, all the data is stored in the cloud. The storage is much closer to the compute in that scenario."

While Autodesk 360 is tied to a handful of Autodesk applications, Lagoa's cloud-based renderer will import from a variety of file formats to render in your browser window.

"You can load a full CATIA file of a BMW 3 series car, and it will render in less than a minute," says Lagoa Vice President Chris Williams. "Every detail — even the stitching inside the car is modeled in 3D."

To accomplish this, Williams says, "you have to build a lot of infrastructure, from things like decimation of geometry, compression routines to stream rendering, components to make a browser-based component run like an application, version control and asset management. We've got a very beefy infrastructure on the back end that dynamically scales."

Lagoa's approach presents other possibilities. "We can deliver our platform not only as an application, but also as a set of APIs," Williams says, referring to application programming interfaces. "About a third of our inbound interest today is people looking to improve the visual experience in a web environment. Most of these are web configurators."

With a few hundred lines of code, you can give your customers, both internal and external, access to your cars, shoes or desktop speakers, letting them repaint and resurface them and view the results in full 3D, in real-time.

"Five or 10 years ago, if you'd come out with this really sophisticated rendering engine, it would have had a hard time getting traction," Williams concludes. "But what we're seeing in the market is that people have become more sophisticated, and they're looking for more." **DE**

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

INFO → Autodesk: Autodesk.com

→ **Dassault Systèmes:** 3DS.com

→ **Imagination Technologies Ltd.:** Caustic.com

→ **Intel:** Intel.com

→ **KeyShot:** KeyShot.com

→ **Lagoa:** home.Lagoa.com

→ **NVIDIA:** NVIDIA.com

→ **RTT:** RTT.ag

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

Real-life **Results** with the UberCloud CAE Experiment

Engineers are seeing higher computing power and ease of use in moving designs and simulations to the cloud.

BY PETER VARHOL

In July 2012, the open, voluntary and free UberCloud CAE Experiment got underway with the goal of providing cloud alternatives for CAE among small and mid-sized manufacturing entities, or SMEs. This project recruited several teams that included high-performance computing (HPC) specialists, design engineers, and design and simulation software vendors. *Desktop Engineering* reported on this experiment (deskeng.com/articles/aabkzs.htm), as well as some of the projects that were using it in creating new designs.

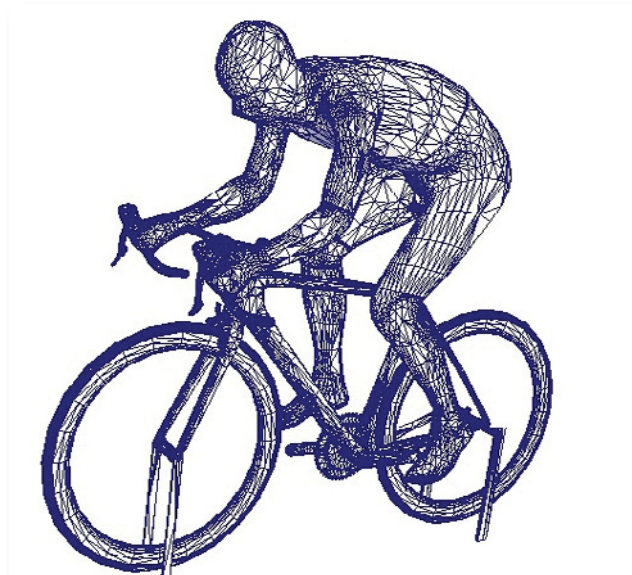
Now that some of that design work has reached an advanced stage, *DE* is following up with a few of the participants to determine what kind of design work they did in the cloud, along with the challenges and benefits of working

with cloud-based designs and simulations. We explore how they fared, including why they decided to move into the cloud, and what their challenges were in doing so.

The two companies profiled here are Trek Bicycle Corp. and Simpson Strong-Tie, two of among the 500 organizations and individuals from 48 countries that have participated in the UberCloud HPC Experiment so far. These teams consisted of an industry end user with one or more engineering design problems, and a software provider, as well as a resource provider and an HPC expert.

Trek: A Bicycle Built for New

Trek Bicycle Corp. started out with the goal of building the best bikes in the world. The company has since expanded its



UberCloud HPC Computing Cloud infrastructure was used to analyze the airflow around bicycle design iterations from Trek Bicycle using a simulation of a rider and bicycle in a wind tunnel.

mission to include providing simple transportation solutions to complex problems. In the past, Trek's designs have undergone iterative design processes including multiple prototypes and wind tunnel testing to create highly aerodynamic designs.

The UberCloud HPC Computing Cloud infrastructure was used to analyze the airflow around bicycle design iterations from Trek Bicycle. The CAPRI to OpenFOAM Connector and the Sabalcore HPC Computing Cloud infrastructure were used to analyze the airflow around bicycle design iterations. According to engineer Mio Suzuki, the project was straightforward, using the UberCloud to set up a simple simulation of bicycle and rider in the wind tunnel.

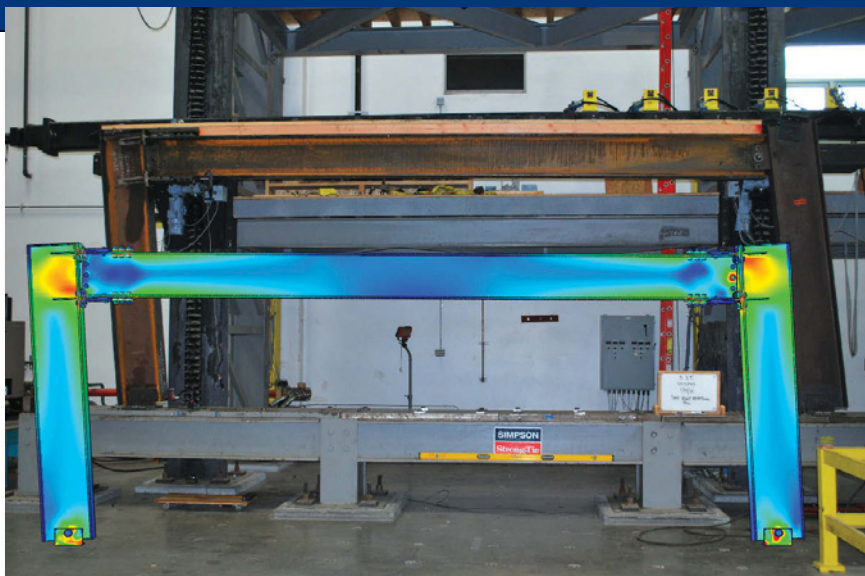
"This type of simulation is typical of the work I ran at Trek to benchmark computational fluid dynamics results to Trek's wind tunnel data, and I used the results to guide the future design direction," says Suzuki.

Trek didn't have significant issues in moving to the UberCloud, Suzuki reports; in fact, it was fairly straightforward to port the simulation into the cloud. "It didn't take a lot of effort," she says. "I believe that's the beauty of CAPRI software provided by CADNEXUS. The software helped to streamline the model importing process, running simulations, and getting data. I've been using cloud service for design/analysis for quite some time now, so I knew what to anticipate should there be any kinks in the process. And I'm always on the lookout for streamlining the process as much as possible."

Suzuki acknowledges that the introduction to the CAPRI software was a great reminder that there are many well-thought-out products and services that can make end users' lives much easier. For this team, though, the exposure to a greater range of tools will almost certainly be a long-term benefit to their experience and knowledge. It also gives the team and the engineering organization more options to consider in future engineering projects.

Suzuki says that there were several significant benefits Trek obtained by working in the cloud for these designs — not the least of which was a gain in computational speed. Trek has in-house HPC capabilities that are sufficient for small-scale runs and for testing new ideas and concepts. But large parallel computation demands additional resource power from the cloud-based system. It simply doesn't make sense for Trek to maintain that kind of computational capability in-house.

In addition, the cloud has helped Trek's design team diversify and accelerate the pace of aerodynamics research.



Simpson Strong-Tie needed cloud bursting for its heavy-duty Abaqus structural analysis of physical designs.

"It may not be widely known, but I think there are many aspects of aerodynamics one can observe and study with a bicycle and/or bicycle and rider," Suzuki says. "Having additional large-scale resource means I can study these details simultaneously."



DeltaGen 12 available now. Be Bold. Be Fast. Be everywhere.

Your needs evolve at a fast pace. And where real change is needed, simple tweaks will not do. DeltaGen 12 completely redefines high-end 3D visualization and realtime interaction. This breakthrough version gives you a broad suite of robust new features to truly revolutionize the way you work and help increase visual quality, speed, and flexibility.

www.rtt.ag/deltagen12



IS LICENSING ON-DEMAND TECHNOLOGY THE FUTURE?

If last month's Altair Technology Conference is any indication, the way people prefer to acquire and use engineering technology may be heading for a drastic change: "Cloud is driving the enterprise solution and new business models," said James Scapa, chairman and CEO of Altair Engineering during his address.

The Era of "What you need, when you need it" has prompted many tech companies to rethink or refine their strategy. PTC's rebranding of the all-inclusive Pro/E as modular Creo pieces is one example. Autodesk's swift shift to cloud-hosted and cloud-augmented products is another. Siemens' move to industry-specific solutions and subscription model on some titles, and a industry-focused approach taken by Dassault Systemes may also be prompted by the same trends.

Perhaps the most significant product from Altair to address the changing consumer behavior is the launch of a three-pronged HyperWorks product line: HyperWorks Unlimited, HyperWorks On-Demand, and Simulation Cloud Suite.

HyperWorks Unlimited, dubbed "a private cloud solution," is high-performance computing (HPC) in a box, an appliance configurable in your desired computing core-count. The hardware is preloaded with cluster management and job queue control software. The appliance is leased, not sold. It comes with unlimited HyperWorks Units — a kind of tokens that can be used to access software from Altair and Altair partners.

If your in-house clusters, workstations, or IT resources prove insufficient to tackle your simulation jobs, Altair's HyperWorks On-Demand could be the reinforcement you call on. It's described as "a public cloud solution," but in terms more familiar to the on-demand industry, it may be called Platform as a Service (PaaS) or Infrastructure as a Service (IaaS). The offering gives you access to Altair's data center to process simulation jobs.

For more, visit deskeng.com/virtual_desktop/?p=7627.

— Kenneth Wong



Trusted and Refined Productivity Tools for Complex 3D File Conversion, Optimization, Viewing, Rendering & More

'Umicore Solar Team' car model converted & optimized for Maxon's Cinema-4D by Okino's PolyTrans|CAD



(c) 2013 Solarteam & Renderhouse BVBA (<http://www.renderhouse.eu>)

Okino's PolyTrans® Pro Translation System is the most extensive and accurate 3D CAD, NURBS, DCC/Animation, and scene translation system with core functionality that includes scene optimization, data reduction, material editing, rendering, viewing & publishing.

Okino Computer Graphics, Inc.

Tel: (Toll Free) 1-888-3D-OKINO, (1-905) 672-9328

WEB: <http://www.okino.com>, Email: sales@okino.com

All products mentioned are registered trade names of their respective holders.

Simpson Strong-Tie: Moving Structural Analysis to the Cloud

Simpson Strong-Tie used the UberCloud for heavy-duty Abaqus structural analysis of physical designs. The applications used by this company range from solving anchorage tensile capacity and steel and wood connector load capacity, to special moment frame cyclic push-over analysis. The HPC cluster at Simpson Strong-Tie is modest, consisting of 32 cores. Therefore, when emergencies arise, the need for cloud bursting is critical. Also challenging is the ability to handle sudden large data transfers, as well as the need to perform visualization for ensuring that the design simulation is proceeding along correct lines.

Engineer Frank Ding from Simpson Strong-Tie led the engineering design effort on this project. In addition, there were UberCloud team members from software provider SIMULIA, a resource provider from Nimble, and the team expert.

Simpson Strong-Tie typically uses HPC resources in solving problems involving structures designed to support frames, traffic or other structures. Ding's particular HPC project dealt with customized concrete anchorage reinforcement design, a typical problem domain for finite element analysis (FEA).

Ding reports that there were few difficulties in moving the project over to the UberCloud. He did, however, note a couple of issues. For example, he found that "FEA result data movement is a bottleneck for the smooth cloud computing workflow." The analysis produced so

much data that getting the data on and off the cloud-based system slowed down the potential turnaround times of newer incremental designs and analysis. With higher bandwidth, this problem may have been able to be reduced or eliminated.

While the UberCloud easily brought HPC computing resources to bear on the analysis problem, Ding also identified several other actual or potential issues in his experiences. For example, while his project gained free use of the software licenses and time needed to complete the project, he notes that application software licensing for the cloud model is still in its infancy, and may make a project more expensive, depending on the computing resources required.

In addition, for many organizations looking to move to the cloud, intellectual property (IP) and security remain question marks. Concerns about protecting IP can make cloud computing a non-starter for some projects.

Still, there were also significant benefits to working in the cloud, Ding says. He identifies several significant benefits in designing and analyzing in the cloud, including:

- removal of HPC infrastructure support;
- the availability of scalable computing resources; and
- a simple pay-for-use model.

For engineering teams with limited budgets for HPC support, he concludes, the ability to rent time and scale quickly can overcome many of the disadvantages.

Building Better Designs

Overall, it is clear that the UberCloud experiment — and perhaps cloud-based HPC in general — isn't a big barrier to building great designs. With appropriate tools support and HPC expertise, doing design and simulation in the cloud can be a fast and cost-effective way to bring new products to market. Even projects that aren't using the cloud today can be migrated in a relatively short period of time.

But it's not going to happen overnight. Engineering organizations have to understand and face tradeoffs in vendor selection, software licensing, data movement and security. Not all of these will be issues on all projects, but teams need to consider these and other potential drawbacks on a project-by-project basis.

Software plays an important role in moving to the cloud. If existing toolsets support operation in the cloud, the design team can avoid a technology learning curve. If new design and analysis tools are needed, there may be training expenses and a period of time for engineers to get up to speed on how to use the software and the results it generates.

But it's also clear that the cloud enables engineering organizations to build better designs more quickly. Thanks to techniques such as iterative design, simulation and analysis, fewer physical prototypes have to be built. Because

larger parts of the design, or even entire designs can be simulated, it can enable further refinement of designs or faster times to market.

The UberCloud project is free, and any use of a commercial HPC cloud will entail costs for computer time and software license rental or purchase. But whatever the cost, there will likely be savings over in-house HPC resources. **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 → CADNexus: CADNexus.com

→ **Dassault Systèmes:** 3DS.com

→ **Nimbix:** Nimbix.net

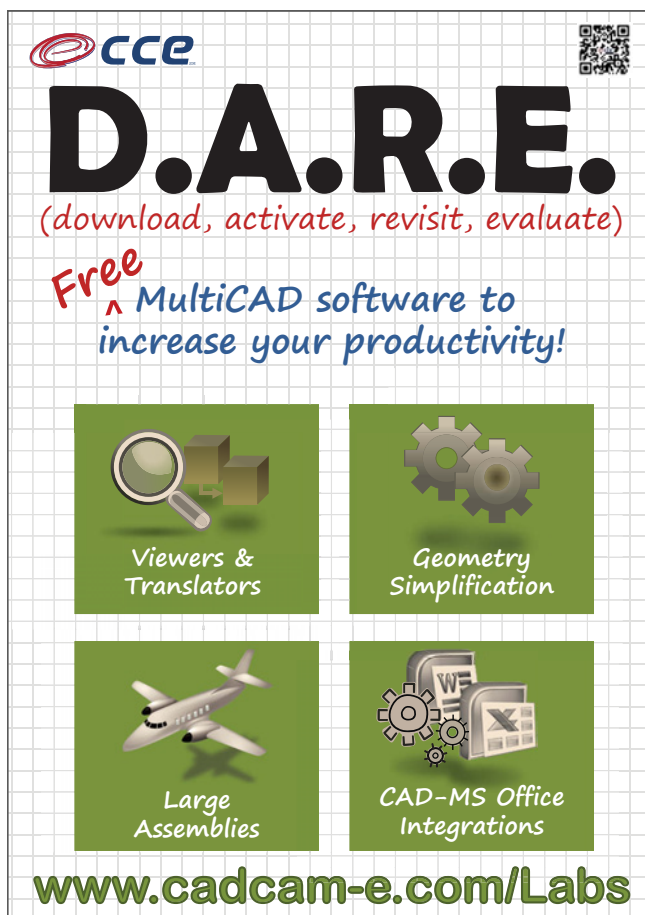
→ **Sabalcore Computing:** Sabalcore.com

→ **Simpson Strong-Tie:** Strong-Tie.com

→ **Trek Bicycle Corp.:** TrekBikes.com

→ **The UberCloud HPC Experiment:** HPCExperiment.com

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



CCE

D.A.R.E.
(download, activate, revisit, evaluate)

Free MultiCAD software to increase your productivity!

Viewers & Translators

Geometry Simplification

Large Assemblies

CAD-MS Office Integrations

www.cadcam-e.com/Labs

Breaking through the Large Assembly Barrier

Increasing product complexity and cross-disciplinary design are adding to the size and scope of CAD models, but new technology advances are facilitating engineers' interaction with large assemblies.

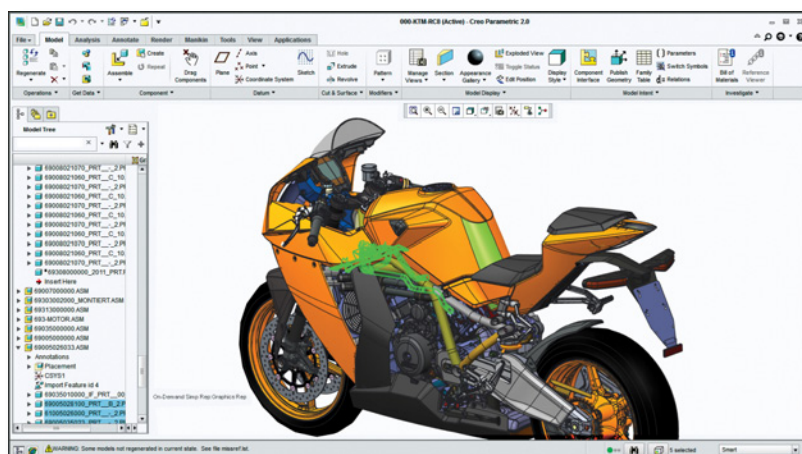
BY BETH STACKPOLE

It wasn't long ago that a complex assembly meant dealing with a 3D CAD model comprised of hundreds, even thousands of parts. Yet what seemed like a heavy load then is now standard fare today, thanks to the ever-increasing sophistication of products—whether it's an aircraft carrier or a state-of-the-art consumer appliance.

In addition to ballooning assembly sizes, engineers are incorporating a richer set of material into the 3D model, such as folding in simulation results or electrical CAD data. It's part of a trend toward creating a complete and realistic digital representation of the product. Besides electrical and embedded software content, the traditional mechanical-only CAD model is also being augmented with metadata, specifications, materials properties, cost estimates and manufacturing options, including the product configuration and variant data required for regional and individual customization of products.

While these highly detailed 3D models go a long way in streamlining the design process and reducing reliance on physical prototypes, they also open up the door to a litany of performance and workflow challenges—among them, the ability to effectively collaborate on concurrent design. According to a survey by market research firm CIMdata, more than 40% of engineers said they experienced excessive delays in loading and interacting with large assemblies, causing great frustration and sapping overall productivity.

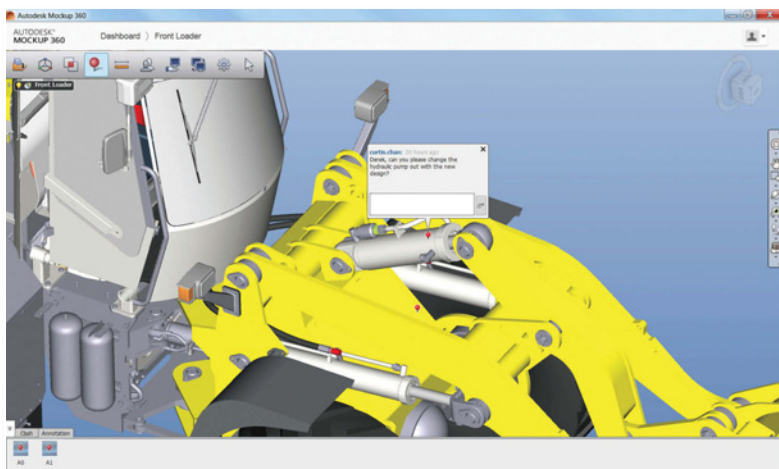
"Ten years ago, companies developing 3D models included just enough information for engineers to communicate to the manufacturing organization how to assemble the product, but we're creating higher-fidelity models today," notes Brian Shepherd, vice president of product management at PTC. "Without the proper tools, engineers can really be



PTC has made its Creo View lightweight design environment available from directly within Creo Parametric to ease the burden of accessing large assemblies. Image courtesy of PTC.

burdened with a lot of waiting time in order to establish the design context that's relevant for them to do their work."

The large assembly problem goes well beyond slowing workstation performance to a crawl: The larger the assembly, the more likely that multiple engineering disciplines are involved in the design process. Thus, managing the structure and concurrency of information is a central part of the challenge, notes Paul Brown, senior marketing director for NX at Siemens PLM Software. Citing the shipbuilding industry as a specific example, Brown explains: "The people doing steel work are different than the people outfitting the piping and HVAC (heating, ventilation and air-conditioning), who in turn, are totally different than the team doing propulsion work. If you end up with a team of 200 design engineers on a vessel, the traditional hierarchical assembly structure is not going to work."



Autodesk Mockup 360 offers a simple way to pull together large-scale project data from many CAD sources in a real-time, cloud-based collaboration and visualization tool. Image courtesy of Autodesk.

At issue, Brown says, is tying up the entire assembly structure simply for one team or a single functional area to gain access to their specific sub-assembly or set of components necessary to do their work. “Traditionally, you’d have to have access to the highest level where the assemblies come together so you can ... get the context of what you need to design,” he explains. “What that means is you have to have access and control over a lot more data, but you risk locking people out and not effectively managing the changes.”

Context Clues

Recognizing that assembly structures are only likely to get bigger and more complex, vendors are actively rolling out solutions aimed at mitigating the pain points around both system performance and concurrent design work. From the hardware side, workstation players like Hewlett-Packard, Dell and BOXX Technologies are pushing the envelope by leveraging, whenever possible, multi-processing and cache technologies — along with over-clocking capabilities to speed up the retrieval, display and navigation of large assembly models.

NVIDIA is also addressing the large assembly problem, deploying technologies like virtualization to help minimize the need to move large 3D CAD files back and forth across the network. NVIDIA is collaborating with partners like Dassault Systèmes’ SolidWorks to optimize CAD software to fully exploit the added horsepower of its graphics processing units (GPUs).

Despite the efforts of hardware and software vendors to speed up large assembly performance, they are starting to reach a point of diminishing returns, reports Ken Versprille, Ph.D., an executive consultant with CIMdata’s Collaborative


Product Development Associates.

“With the pure computer science type of operations, there comes a point where you ask, ‘How much more improvement can you really get?’” Versprille continues. “They’ve already spent 10 years trying to improve their code, and they’ve milked every microsecond they can get already. Now they are starting to look for other types of solutions.”

One proven alternative to boosting large assembly performance is employing lightweight visualization technology wherever appropriate, as opposed to having to load and interact with a full-scale geometric model for every operation. Substituting a lightweight, graphics-only representation of a product



component or subassembly allows engineers to view their portion of a product design in context with related areas — and in some cases, allows for some minor editing.

PTC has embraced this approach with Creo Parametric. The latest release, which incorporates the Creo View lightweight visualization technology, loads the critical graphics and structure of a model to provide the full perspective, only add-




Personal CNC

Shown here is an articulated humanoid robot leg, built by researchers at the Drexel Autonomous System Lab (DASL) with a Tormach PCNC 1100 milling machine. To read more about this project and other owner stories, or to learn about Tormach's affordable CNC mills and accessories, visit www.tormach.com/desktop.

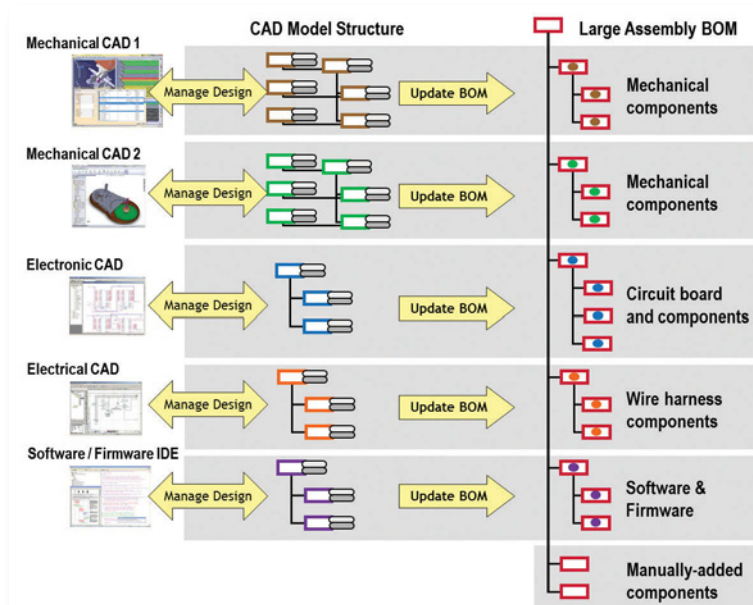
Mills shown here with optional stand, machine arm, LCD monitors, and other accessories.



PCNC 1100 Series 3

PCNC 770 Series 3

[**www.tormach.com/desktop**](http://www.tormach.com/desktop)



Aras' standardized CAD integration platform, which includes unified data and process models, delivers high performance for large assembly modeling for global multi-CAD environments. Image courtesy of Aras.

ing the geometry and feature history when an engineer zooms in on a particular area in the model to get a more detailed view. Siemens PLM Software takes a similar approach, integrating its lightweight JT format into NX's data structure.

While lightweight visualization can address some of the performance issues associated with large assemblies, most CAD providers are moving a step beyond to tackle the broader problem of allowing multiple engineering teams, including those in different disciplines, to collaborate on concurrent design. While each CAD player has a slightly different approach, the concept is pretty similar:

- allow engineers to flexibly define sections of a product assembly that are relevant to their role and work deliverables;
- eliminate the need to load the entire assembly, thus speeding performance; and
- support concurrent design via data and change management workflows.

Siemens PLM Software calls its approach Fourth Generation Design (4GD), a technology making its initial debut in the new Catalyst solution for the shipbuilding industry. As a combination of both NX and Teamcenter technology, 4GD is a new take on the hierarchical assembly structure that doesn't have fixed relationships. Its users can change their view of the structure depending on the job requirement, and without tying up the rest of the assembly data. 4DG, which Siemens officials likens to a usage-based bill of material (BOM), will be integrated in all Teamcenter and NX solutions over time, notes Brown.

"Just like Synchronous Technology freed up the sequential nature of the parametric history tree, 4DG frees up an assembly from the bounds of a hierarchical structure so engineers can see the bits they need to see without tying up data," Brown adds. With 4DG, engineers build up recipes to create context for their areas of interest — defining, for example, that they want to see areas of the assembly related to HVAC design, but not steel work. Based on these parameters, which can be continuously modified, NX and Teamcenter automatically load the parts of the assembly that are relevant instead of pulling in everything and then filtering out unnecessary data, he explains.

Product lifecycle management (PLM) provider Aras is also attacking the concurrent collaboration challenge with large assemblies, albeit from a slightly different angle. Aras has been architected as a standardized CAD inte-

gration platform, allowing disparate functional engineering groups to work inside a common assembly.

"Because we've never built any optimizations around one CAD system, we've been able to define a neutral representation of the CAD structure so you can have Creo and CATIA check in and check out in the same assembly," explains Peter Schroer, Aras president. "By doing so, you allow people in a big design team to work in one assembly and have a consistent structure, workflow and data model."

Aras is also able to leverage its native Web architecture to take advantage of techniques such as compression and multithreading to speed up large assembly design, Schroer says.

Virtualization and the Cloud

Two of the hottest information technologies — virtualization and the cloud — are also being used to tackle the large assembly performance and collaboration problem. NVIDIA and SolidWorks have partnered to certify SolidWorks 2014 to support the new NVIDIA visual computing appliance (VCA), a turnkey GPU virtualization offering that allows workgroups on a LAN to remotely run GPU-accelerated Windows applications on any Mac, Linux, or Windows computer.

"A large assembly means more complex collaboration, and with virtualization and remoting, you can have access to designs [and large assemblies] wherever you are," notes Andrew Cresci, NVIDIA's general manager of manufacturing.

NVIDIA is also working with SolidWorks to optimize SolidWorks 2014, to better leverage the NVIDIA Kepler GPU architecture for large assemblies. In testing, the pair have achieved an average of 2x faster performance with assemblies larger than 500 components compared to SolidWorks 2013, NVIDIA claims.

For Autodesk, the cloud is essential in addressing the large

assembly problem, both for increasing performance and easing cross-disciplinary collaboration. In addition to Inventor capabilities like Express Mode, which speed up assembly load times by leveraging mesh-based data stored in a cache, the company is pursuing several cloud-based options to allow dispersed engineering teams and partners to collaborate on a common assembly without sacrificing the proper change management and tracking.

MockUp 360, for example, is a real-time collaboration and digital mockup tool used to quickly aggregate large-scale 3D CAD data, enabling team members to easily collaborate on large assemblies in the cloud and track changes.

“Mockup 360 lets groups use any CAD system and contribute to a common data model,” explains Peter White, Autodesk’s product manager for 360 Mockup. “Engineers can connect anywhere in the world because it’s cloud-hosted. You simply position your data in the context of the design, and the software translates it and integrates it into the overall product architecture.”

The just-released Autodesk Remote, a cloud service that lets users drive their Autodesk applications on their primary computer from a remote PC or iPad, will help address system performance issues when accessing native design data, he adds.

Even with the breadth of new solutions, it’s safe to say that software and hardware vendors will continue to chip away at the large assembly problem — although it’s unlikely to let up anytime soon.

“It’s a challenge, and it’s something we are continuing to focus on,” says Bill Lewis, director of product marketing for Teamcenter at Siemens PLM Software. “We don’t consider it solved by any stretch.” **DE**

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

INFO → **Aras:** Aras.com

→ **Autodesk:** Autodesk.com

→ **BOXX Technologies:** BOXXTech.com

→ **CIMdata:** CIMdata.com

→ **Collaborative Product Development Associates:** CPD-Associates.com

→ **Dassault Systèmes:** 3DS.com

→ **Dell:** Dell.com

→ **Hewlett-Packard:** HP.com

→ **NVIDIA:** NVIDIA.com

→ **PTC:** PTC.com

→ **Siemens PLM Software:** plm.automation.Siemens.com

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

SPOTLIGHT

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

Forget everything else. Performance IS your reality.

Lowering expectations to meet your workstation's performance?

Spending more time waiting on engineering applications instead of actually using them?

Need a workstation that exceeds your expectations?

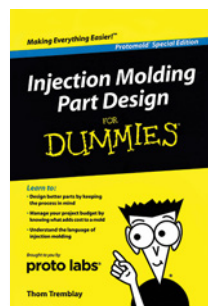
Did you answer **YES** to all three? Get a BOXX.

Build Your BOXX Now!

www.boxxtech.com/Industries/architecture



BOXX



New to the world of plastic parts and injection molding?

Get a FREE copy of our book that will help you make better parts, fast.

proto labs®
Real Parts. Really Fast.™

Visit www.protomold.com/book and enter code DT13DB.

MILLIONS of FREE CAD Models

- Download CAD models
- Supplier Certified models
- All CAD formats supported
- 3D Viewing



COMING SOON!

3DX

TracePartsOnline.net is a Free CAD portal with 100's of 3D supplier catalogs. Millions of CAD models and product data sheets available for designers, purchasing, and manufacturing.

traceparts
Product Content Everywhere™

Check it out yourself. Join our other 1.5 million users at: www.tracepartsonline.net

Metal 3D Printing on the Rise

From ordinary to exotic, AM machines and materials are getting the spotlight.

BY PAMELA J. WATERMAN



Image courtesy
of Renishaw.

Plastics may dominate the media coverage of 3D printing, but metal applications have quietly made tremendous progress in the past decade. We're talking real metal, real parts, that can last longer than traditionally made pieces. In fact, so many systems, service bureaus and material providers now focus on this process that one article can only skim their recent activity.

In witness to change and growth, companies with prototyping in their name now do end-part manufacturing; more sources exist for metal powders, and a major service bureau has been bought up by no less than GE Aviation. Metal additive manufacturing (AM) is truly a mainstream process.

Systems and Processes

Time was you could count metal-capable AM systems on one hand. Terry Wohlers, the AM industry expert whose annual Wohlers Report has been a valued resource since 1996, says 11 established metal AM system manufacturers are now operating worldwide, with three traditionally active in the U.S.

"Among the major players in metals are Arcam, Concept Laser, EOS, ExOne, Phenix Systems (now a part of 3D Systems), ReaLizer, Renishaw and SLM Solutions, with Optomec producing systems for more specialized applications," Wohlers observes.

Metal AM technology has developed impressively in a relatively short time, Wohlers says, adding, "In some ways, it has developed further in 10 years than plastics systems have developed in 25 years."

A key element of this success comes from the wide range of approaches for building layered parts from metal: melting powders with lasers or electron beams, or binding powders with liquids then sintering the products. Similarities and differences

abound with each manufacturer's system, whether in the price, material options, part size and detail accommodated, or required infrastructure (an argon supply, for example). An increasing number of systems already offer options in stainless steel, tool steel, bronze alloys, cobalt-chrome (CoCr), Inconel, aluminum and titanium (See "Additive Manufacturing 101," *DE* November 2010), but developers are not resting on their laurels.

Finding Niches, Filling Needs

German direct metal laser sintering (DMLS) equipment provider EOS has the strongest presence in the States, with systems at both service bureaus and manufacturers. Thanks to the "aging of society," the company sees real growth in the medical business as devices call for more functional integration and miniaturization. To meet this need, EOS has been reaching out to partner with various industrial end-users on joint development projects.

EOS Project Manager Joachim Goebner says many industries are unaware of what can be done with "traditional" AM, let alone his EOS' micro-laser-sintering (MLS) process — a project in development since 2002 that is producing parts with 32-micron wall thicknesses. He points out, "We need to identify needs and projects, and inspire people by showing them the possibilities."

The MLS system can create parts at full density (99%) or with controlled porosities down to 30%. Goebner says almost any metal is a candidate for this approach, even high-temperature tungsten, stainless steel and molybdenum. In July, the

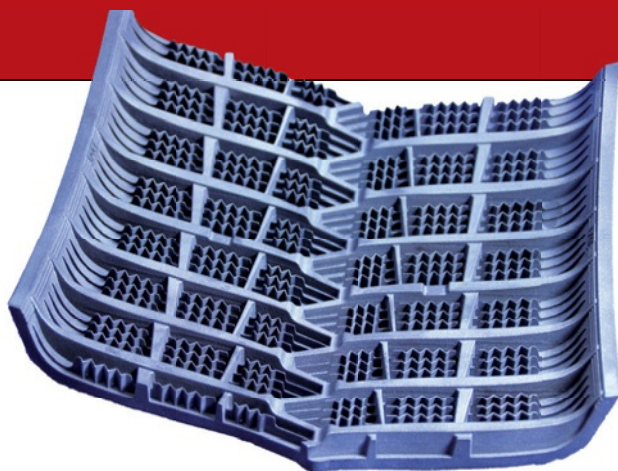
company installed an MLS system at Germany's Technical University of Darmstadt, to continue R&D on such mechanical properties as surface roughness; it plans to offer parts in 2014.

After years of producing more than 30,000 AM custom medical implants for European markets, Arcam has also expanded into the U.S. The new Arcam Q10 replaces the Arcam A1, and has been designed to meet the implant industry's need for high productivity, high resolution, ease-of-use and quality control. The system includes the new Arcam LayerQam, a camera-based monitoring system for inline part-quality verification.

ExOne, the third major U.S. player, takes a very different approach to metal AM, using inkjet printing technology to deposit a liquid binder onto powdered material. From the start, the company has put considerable effort into new material development on both binders and powders, starting with 316 stainless steel, 420 stainless steel, copper alloys and bronze. Based on work done at ExOne's Material Applications Laboratory (ExMAL), the company recently announced the availability of iron (for bronze infiltration) and bonded tungsten — plus continuing efforts with titanium, Inconel and magnesium. The development program includes working out optimal sintering temperature profiles for the new materials, but an attractive feature is that all process parameters are open to users.

Selective laser melting (SLM) technology, as with many of the metal AM technologies, has roots going back to the mid-1990s. This process was developed by Dr. Matthias Fockele (then-owner of F&S GmbH) together with the Institut für Laser Technik, Aachen, Germany, and has evolved both in terms of technology and business.

According to Fockele, ReaLizer SLM systems create parts in alloys of aluminum, CoCr, titanium, steel, gold, platinum and



A section of tread mold created for the tire industry using high-strength maraging steel 1.2709, printed on a 3D Systems Phenix PXL system. Image courtesy of 3D Systems.

Inconel, with more materials in development. Parts can be fully dense, as for dental prostheses, or filled by a lattice structure appropriate for bone growth in implants or the design complexity of jewelry. Though not yet a large force in the U.S., ReaLizer has a solid European customer base of around 200 systems.

Various partnerships and patent agreements have led to the present use of the same term, SLM, to describe the process behind equipment marketed today not only by ReaLizer but also by German-based competitor SLM Solutions. This year, SLM Solutions opened a subsidiary business in Novi, MI, to directly supply its U.S. customers with its four metal-based models.

Another split-off from Fockele's original business opened the door to laser melting systems now sold in the U.S. by U.K.-based Renishaw. Demonstrating the extreme durability of AM parts, a Renishaw system was recently used to build a titanium prototype nose cone for the British Bloodhound Supersonic Car that will attempt to break the 1,000-mph speed barrier in 2015.

In another sign of growth in this arena, Germany's Concept Laser is penetrating the U.S. market. Its laserCUSING systems process a range of metals, including powdered forms of titanium alloy (rematitan CL) and CoCr (remanium star CL), as patented by Dentaaurum Group, and targeted to implantable dental work.

On the domestic front, 3D Systems has completed its acquisition of Phenix Systems, gaining a second direct-metal process for supporting aerospace, automotive and patient-specific medical device applications. Parts made on a Phenix are very dense, and achieve surface finishes as accurate as 5 microns roughness average (Ra) — important for tiny parts. Parts can also be up to 250 x 250 x 300mm. The Phenix systems join 3D Systems' SLS product line of Direct Metal production printers, and, according to the company, represent the capabilities of the next generation.

Two different processes from Optomec offer a different slant on metal options. The New Mexico-based developer employs metal powders in the laser engineered net shaping (LENS) process for part production, repair or enhancement, as well as in Aerosol Jet systems that can deposit metal traces to print electronics. The company also markets the LENS MR-7 system, targeted to materials researchers for rapid production of quality

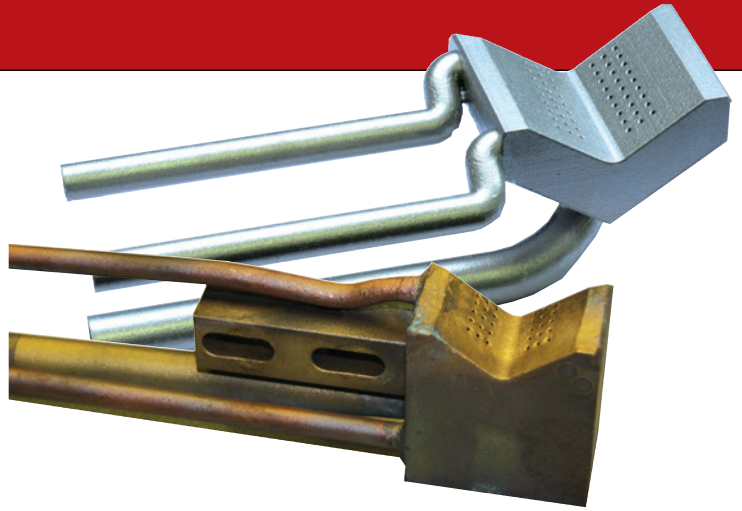
Freedom to Optimize AM Operations

Some of the increased interest and activity in metals AM stems from the widening public exposure to 3D printing in general, but some also comes as more systems let users take charge. On the hardware side, Chuck Alexander, product manager at the Solid Concepts service bureau, says he is glad there are now many players in the AM systems world. He even welcomes the increased openness for powder sourcing.

"In the CNC world," he adds, "they don't tell you what material you can buy or how to program [the machine] for that material. Manufacturers try to lock you into process parameters — for some good reasons, such as keeping you from breaking the machine and trying to keep control over intellectual property. But I want to be able to have the tools to do what I want. When more competition comes out, there'll be more options."

—PJW

Old and new (top) water-cooled gas burners used by customer Havells Sylvania to heat quartz glass. The old unit was welded from 20 pieces and rapidly wore out. The new unit was completely redesigned by LayerWise to take advantage of AM techniques: It is built in one piece, with significant material reduction, and includes conformal cooling channels that are less than 1mm away from the burning gas. The new units tripled the burners' lifetime. *Images courtesy of LayerWise.*



novel materials. In the MR-7, multiple powder feeders can build gradient materials into complex alloys, while sensors collect microstructure property data. Tools for thermal and melt-pool monitoring also support research work and control.

The Nitty Gritty of Materials

Metal AM wouldn't be possible without reliable sources of powdered metals with specified properties and controlled granule sizes. While a decade ago, most system suppliers only sold proprietary materials, customers now have a growing number of sources for buying and co-developing raw powders.

In the U.S., Erasteel supplies metal powders from offices in New Jersey and Illinois. The company, a division of ERAMET, has a 40-year history of metal powder production at its Swedish headquarters. It markets a range of Pearl Micro powders.

Norbert Ludwig, spokesperson for TLS Technik of Germany, lists the many materials his company delivers to both system original equipment manufacturers (OEMs) and end-users worldwide. "Our main powder products are pure titanium, titanium alloys, steel grades, and aluminum, copper and nickel alloys," he says. "The particle size depends on the alloy and machine, and is in the range between 10 to 60 microns."

Another supplier, LPW Technology in the U.K., has had a global business since 2007. It offers custom and particle size distributions, and has introduced LPW PowderSolve — online soft-

ware that gives customers a full-traceability database, important for lifecycle and trend analysis.

Sandvik Materials Technology (also based in the U.K.), a division of Sandvik Osprey, offers standard and customized alloys with sizes starting at 250 microns. It specializes in fine powders (<38 microns) and prides itself on producing spherical granules suited to AM systems for good flow and high packing density.

In the realm of non-traditional suppliers, Rhode Island-based BEGO dental has patented a high-performance chrome-cobalt-molybdenum alloy, Wirobond C+, for building dental implants on EOS systems. The company is also working on a noble powder, a chrome-palladium-trace-element alloy that would provide better stiffness in AM dental restorations; it predicts this will be a "game-changer" for its U.S. market.

Service Bureaus Gain New Opportunities

A description of metal AM, particularly as a true production process, would not be complete without mention of service bureaus. The expertise of Morris Technologies, previously the major U.S. service bureau for metal applications, prompted the company's 2012 acquisition by GE Aviation. Greg Morris, founder and former president, explains his company is no longer a "service provider," although they are selectively working with a few external companies where there is a strategic fit.

Fineline Prototyping, a business with plastics AM experience dating back to 1989, brought on the first two U.S. Concept Laser systems in June 2012 and has just installed a third, the larger M2 machine. Founder and President Rob Connelly says the latter purchase will support increasing requests for larger parts, particularly in titanium, for aerospace applications. Fineline is also qualifying high-temp Inconel 718, which is difficult to machine traditionally. Connelly says he chose the laserCUSING technology for both its high resolution and open-access operation.

Woburn, MA-based Powderpart is a new U.S. business backed by seven years of European AM metal experience. As a service bureau, Powderpart is a branch facility of the Italian company Ci-Esse, a machine shop that launched 35 years ago to support Ferrari. General Manager Luca Colombini founded the Ci-Esse AM metals department, working with Concept Laser systems, and has brought that same technology and expertise to the U.S. market.

"We have a deep relation with Concept Laser," says Colom-

Metal Service Providers

Here's a look at some other service bureaus with strong capabilities in AM metal processes. Check out the resources on their websites:

- 3trpd: 3trpd.co.uk
- Cideas: BuildParts.com
- Directed MFG: DirectedMFG.com
- DM3D Technology: POMGroup.com
- GPI Prototype & Manufacturing Services: GPIPrototype.com
- Incodema: Incodema.com
- LaserSintering.com: LaserSintering.com
- Mydea Technologies: MydeaTechnologies.com
- Solid Concepts: SolidConcepts.com

bini. "These machines are fully equipped with their quality management system [that lets us] constantly control all details running the machine. When the part is finished, we can go through all the data and say it has no issues, or yes, something may have come out that was strange. We have that information."

Colombini says his company's focus is on production parts, often for aerospace and medical markets. He notes they buy materials from various sources as well as develop their own. Powderpart also offers an UltraFine post-processing surface treatment.

Perhaps the best-kept secret in metal AM services is LayerWise, the Belgian company that has developed its own SLM systems that it runs in-house with such materials as stainless steel, titaniums, CoCr, tantalum and tungsten. The company is active in optimized part engineering, helping customers make the best use of AM's design freedom.

"We see that successful manufacturers are looking for a highly specialized metal AM partner like LayerWise. The strength of tight collaboration is the valuable interaction with customers when revisiting the original part or assembly design," says LayerWise Marketing Manager Rob Snoeijs. He reports that activity is increasing in the U.S. in the dental, medical and industrial sectors.

As Morris says, "These are exciting times for additive not only at GE, but in our industry and around the world." **DE**

INFO → 3D Systems: 3DSystems.com

→ Arcam: Arcam.com

→ BEGO: BEGOusa.com

→ Concept Laser: Concept-Laser.de

→ Dentaurem Group: Dentaurem.de

→ EOS: EOS.info

→ Erasteel: Erasteel.com

→ ExOne: ExOne.com

→ Fineline Prototyping: FinelinePrototyping.com

→ GE Aviation/Morris Technologies: MorrisTech.com

→ Institut for Laser Technique: ILT.Fraunhofer.de

→ LayerWise: LayerWise.com

→ LPW Technology: LPWTechnology.com

→ Optomec: Optomec.com

→ Powderpart: Powderpart.com

→ Realizer: Realizer.com

→ Renishaw: Renishaw.com

→ Sandvik Materials Technology: SMT.Sandvik.com

→ SLM Solutions: SLM-Solutions.com

→ TLS Technik: TLS-Technik.de

→ Wohlers Associates: WohlersAssociates.com

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

Powderpart exhaust system component built in Inconel 718 on Concept Laser M2 laserCUSING system (6 x 6 x 4.5 in.), finished with UltraFine postprocessing system (proprietary to Powderpart's parent company Ci-Esse). Image courtesy of Powderpart.



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.

Advertising Index

| | |
|------------------------------------|-------|
| Altair-Product Design..... | 17 |
| ANSYS..... | CV2 |
| Aras Corp..... | 19 |
| Aras Corp..... | CV3 |
| Bluebeam Software, Inc..... | 7 |
| BOXX Technologies..... | 11 |
| BOXX Technologies..... | 45 |
| CADCAM.E.com, Inc..... | 41 |
| CD-adapco..... | 21 |
| Cobham Technical Services..... | 35 |
| COMSOL..... | 5 |
| CPFD Software..... | 9 |
| CST of America, Inc..... | 15 |
| Epson..... | 1 |
| ESTECO SpA..... | 29 |
| HBM-nCode..... | 31 |
| Intel Corp..... | 13 |
| Intel Corp..... | 26-27 |
| MSi Computer Corp..... | 23 |
| MSi Computer Corp..... | 37 |
| National Instruments..... | 3 |
| Okino Computer Graphics, Inc..... | 40 |
| Proto Labs, Inc..... | 45 |
| RTT..... | 39 |
| Solid Concepts..... | 33 |
| Stratasys-Mojo..... | CV4 |
| The Foundry Visionmongers Ltd..... | 25 |
| Tormach LLC..... | 43 |
| Traceparts..... | 45 |

The Simulation-Test Loop

Siemens PLM Software's LMS business segment improves performance and system linkage with Test.Lab Rev 13.

BY RANDY FRANK

A more complete system solution — that's the goal of many technology companies today. And with its acquisition of LMS, Siemens has made significant progress toward this goal.

Saying it's the first product lifecycle management (PLM) software company to provide a closed-loop, systems-driven product development solution with integrated test, Siemens had a strong rationale for its acquisition of LMS, according to Bruno Massa, vice president of the Test Division, LMS, A Siemens Business.

"There is a very strong complementarity between their solutions and our solutions. That's obviously true for the simulation side of things," says Massa. "There is almost no overlap between their tools and our tools."

On the simulation side, Siemens had focused on finite element analysis (FEA) — with NX product design software as one example. In contrast, LMS focused on specific performance attribute domains, such as acoustic simulations — for example, ride and handling simulations and vehicle energy management in automotive applications. As a result, there is a natural fit between these portions of the two businesses.

Test provided another key piece of the technology puzzle, Massa explains: "The test part plays an important role in Siemens' strategy because Siemens wants to evolve in the future toward a tool-centered solution set that not only allows the virtual world type of validation, but also to do that with a closed loop to the real world."

Because test was at the core of the LMS strategy, the combination of the two companies made the fit even better, he adds.

The Testing Approach

Unlike simulation software, test starts with actual hardware — a vehicle, for example. "Typically, we do not start with a CAD model or an FE model," Massa says. "We start from the real thing, a prototype, a basic touchable part."

As seen in Fig. 1, the team instruments the vehicle with "all kinds of sensors," including a microphone, vibration sensors and strain sensors, he adds. The sensors connect to the LMS SCADAS data acquisition system. This hardware provides the first digital treatment of the sensor signals, and then transfers that information to Test.Lab. The software-based testing system performs further treatment of the signals — including time



FIG 1: Testing the acoustic characteristics of an electric vehicle involves traditional road noise and wind noise, as well as high frequencies generated by the inverters used to control the power of AC motors through pulse width modulation.



FIG 2: A dual-channel LMS SCADAS FR4 Flexray module and a four-channel LMS SCADAS DCH4 module for high temperatures are new data acquisition modules in LMS Test.Lab Rev 13.

synchronization and proper visualization of the data for users. It also provides a broad range of analysis tools to understand the root cause and source of noise and/or vibration (see Fig. 2).

“All that chain of activity is quite independent from any CAD environment or simulation environment,” says Massa. “You don’t need simulation or CAD tools to do that job.”

In addition, testing can involve a prototype that is in its final validation/verification stage before sign-off for production. If there are performance aspects that do not meet the targeted requirements, the purpose of the test will be troubleshooting to refine and further correct the prototype, and determine the required design changes.

Another end purpose for test is quite different, and involves simulation. In this case, the testing is performed to capture information from a real vehicle — perhaps a predecessor vehicle or a benchmark vehicle so the information can be fed back into simulation models.

“It’s to take out some uncertainties in your 3D or simplified simulation models, to basically calibrate and validate your models,” Massa explains. “With further refinements of the design using 3D models, you can basically work with validated assumptions from the real world.”

The models can be component, or subsystem levels such as a suspension or vehicle body. It can even be a full vehicle. In applications like automotive, Massa says, there is an increasingly tight link between these two roles.

New Systems Requirements

Today’s vehicles have created new system requirements — as well as the need for more comprehensive testing, especially acoustic testing. The lack of engine noise in electric vehicles, and hybrid vehicles when they operate only on electric power, reveals acoustic noise that was normally masked. Road noise, the connection between the tires and the road, and wind noise, especially at higher vehicle speeds, are two common noise sources.

LMS’ Test.Lab Transfer Path Analysis (TPA) technology allows users to investigate the relationship between the noise level that a driver would experience in a vehicle to determine whether it is. TPA traces the flow of vibro-acoustic energy from a source, through a set of known structure- and air-borne transfer pathways, to a given receiver location.

“Today, the technology that we have developed and also extended in Revision 13 allows us to understand and investigate what the real source of the noise issue can be,” Massa says. He notes it can also assist in determining what to modify to reduce or eliminate any noise problems.

Revision 13 of LMS Test.Lab combines classical frequency-based and time-domain TPA measurements with pass-by noise (PBN) measurements (as defined by ISO 362, with required limits regulated by various governmental bodies) to allow engineers to easily analyze which subsystems contribute most to the total noise level, relative to the vehicle position. With this

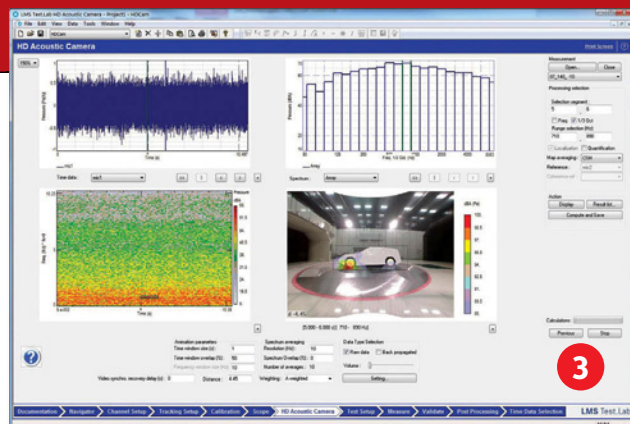


FIG 3: The data from high-definition acoustic cameras allows the precise location of the origin of various sounds.

approach, all interior or exterior PBN measurements can be finalized within just two days (see Fig. 3).

New Systems Capabilities

The Sound Source Localization tool added to Test.Lab Rev 13 provides a 3D visual representation of the sound in the vehicle. During vehicle testing, this allows on-the-spot investigation of where leakages are occurring in the interior — through the heating, ventilation and air-conditioning (HVAC) system, stereo speakers, door and window seals or other openings.

“These are tools that allow users to accelerate the testing — and specifically the conclusion you can take from a test — and bring that at the same exact moment while you are physically performing your test on the vehicle,” Massa explains. “You no longer have to wait a couple of hours, or worst case, a couple of days or weeks to come to an end conclusion of what’s good or wrong with your vehicle.”

Another key aspect of Test.Lab Rev 13 involves data management. According to Massa, the new solution ensures that users don’t “just create and generate large volumes of data, but that they also have a structured solution inside their Test.Lab product to manage the data, and can easily search results, distribute results, create reports and make those reports available.”

Taking Test to a New Level

Indications of closer linkage with Siemens are quite clear, but the how and when are yet to be announced.

“Siemens has a clear direction of not just acquiring a company, but also from the tools perspective, to have integration and further cross-links between the solutions,” says Massa.

For now, the new capabilities in Test.Lab Rev 13 bode well for tighter linkages in future software solutions. **DE**

Randy Frank is a contributor to DE. Send e-mail about this article to DE-Editors@deskeng.com.

INFO → LMS, A Siemens Business: LMSIntl.com

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



Improved Combustion Modeling

Today's on-road vehicles produce more than a third of the carbon monoxide and nitrogen oxides in our atmosphere, and more than 20% of the global warming pollution, according to the Union of Concerned Scientists. Primary forms of transportation, such as cars, trucks, buses, ships and trains, account for 90% of the cancer risk caused by air pollution in the greater Los Angeles area. In response, governmental bodies are promoting new particulate matter regulations, such as the European Union's Euro 6 standard and the California Air Resources Board's Ambient Air Quality Standard, which restrict soot particle size and number limits in addition to total emissions mass limits.

Perhaps the top challenge faced by today's engine manufacturers is designing high performance internal combustion engines that meet regulatory mandates for reduced levels of greenhouse gases and other toxic emissions. Combustion computational fluid dynamics (CFD) simulation offers design engineers the potential of creating lower emission engines while reducing the millions of

bustion). What proved to be good enough for the development of yesterday's engines is insufficient for today's designs.

In addition, the fuel landscape continues to evolve and become more complex as today's engine specifications demand fuel flexibility (not just gasoline, but various fuel blends with ethanol; and not just diesel, but also biodiesel) while achieving lower emissions.

New engine designs, more complex fuel types, and stricter emissions requirements have meant that combustion simulations need to do a lot more computation — in fact, chemistry calculations can account for as much as 90% of CFD simulation time.

In the past, the typical way designers have dealt with time-to-solution issues was to simply buy more CPUs, or run severely reduced data sets to produce a simulation in a reasonable amount of time. However, neither is a good answer: The first leads to increased capital expense, and the latter limits a designer's ability to accurately predict emissions and other combustion effects.

Combustion Modeling

Identifying the correct combustion trends through simulation is critical to making good design decisions. Conventional CFD simulations for spark ignition engines are highly mesh-dependent, involving the generation of a computational grid that approximates the dynamics of fuel combustion. These complex interactions require severe mesh refinements and very small cell sizes to resolve flame location and topology. Conventional CFD solutions are also limited in the amount of chemistry detail they can incorporate within a simulation. Accurately modeling real fuel behavior requires more chemistry than traditional CFD approaches can handle while delivering acceptable time-to-solution.

Reaction Design's FORTÉ CFD Package uses well-established computational techniques, not complex mesh refinements, to create realistic 3D modeling of internal combustion engines. This solution couples proven mathematical techniques and algorithms with detailed chemistry input libraries to simulate spark combustion and predict the effects of operating conditions and fuel variations on engine emissions and particulate matter.

With innovative approaches to relieving the bottleneck in chemistry calculations, predictive engine simulation is now a real tool that can be harnessed for reducing greenhouse gas and toxic pollutants from internal combustion engines. **DE**

Bernie Rosenthal is chief executive officer at Reaction Design. Previously, he was the co-founder and senior vice president in charge of worldwide business operations at Tensilica, and has held various executive and management positions at Synopsis, AMCC and TRW. Contact him via reactiondesign.com or de-editors@deskeng.com.

Predictive engine simulation is now a real tool that can be used.

dollars spent on physical prototyping and experimentation. However, to ensure that engine simulations can accurately predict real-life fuel emissions, both complex mathematical algorithms that describe the physics and thermodynamic behavior of combustion, and a detailed understanding of different types of engines and the chemical makeup of fuels is required.

Virtual prototyping of engines using CFD allows designers to exploit strategies that vary fuel ignition timings and control other combustion parameters and use simulation results to guide design decisions that can impact engine emissions. However, accurately modeling real fuel behavior requires much more detailed chemistry models than conventional CFD packages can handle.

The Model Fuels Consortium (MFC), an industry-led program, has developed the detailed chemical models required to simulate how fuel behaves during combustion. But, commonly used methods to reduce the complexity of this information to the point that it can be incorporated into traditional CFD simulation approaches result in a substantial loss in predictive accuracy for such impacts as emissions.

This issue is exacerbated by the fact that new, high-efficiency, low-emissions engine designs present technical challenges that are highly influenced by chemistry (e.g., dual-fuel engines, staged spray injections, and premixed charge compression ignition com-



DIFFERENT PLM

INNOVATIVE • OPEN • MODERN

Discover Your Next Generation PLM

WWW.ARAS.COM/DIFFERENT



A STRONG CASE FOR COMPLEX DESIGN.

The Mojo 3D Printer knocks out complex parts with ease.

There's only one desktop 3D printer on the market that arrives equipped with everything designers and engineers need to start 3D printing complex geometries. The Mojo is powered by genuine FDM® to build spot-on, functional concept models and rapid prototypes in ABS*plus* thermoplastic. Take it for a spin and we think you'll agree – the competition better get a move on.



Only \$185/month

**See if your printer can make the case.
Visit stratasys.com/case to download a free
STL file of the cell phone case (pictured).**



Stratasys®

Mojo®