Desktop Engineering®

JULY 2009 • VOL. 14, ISSUE 11 • $9.00

TECHNOLOGY FOR DESIGN ENGINEERING

DESKENG.COM

elements of HPC

Wider is Better

Fast, affordable color comes to large-format printers

- HPC Energy Efficiency
- Why Multicore Matters
- Flow Separation Solved
- HP xw8600 with New Quadro
- Scaling with Parallel MATLAB
- Off to the Races with EnSight
- Adobe Photoshop CS4 Extended
- Omnify Empower PLM Examined
- Leveraging Abaqus in a Soft Economy

Accelerate your drive for perfection.

ATI FIREPRO™ GRAPHICS

Give your designs an edge with ATI FirePro™ 3D graphic accelerators.

It's the power and speed to take your project as far as your imagination can go. Never settling, constantly revising, testing, improving, and tuning, all the while knowing that you're working faster than you ever thought possible. Relishing the fact that your complex design is that much better, that much more efficient, and that much more advanced as a result. After all, ATI FirePro™ 3D graphic accelerators detect, optimize, and tune applications for superior performance. Plus, they give you rock-solid reliability for extreme productivity.

What does that mean for your final design? It's up to your standards. And for you? Standards that just got a whole lot higher.

The time has come. Put ATI FirePro 3D graphic accelerators to work for you.

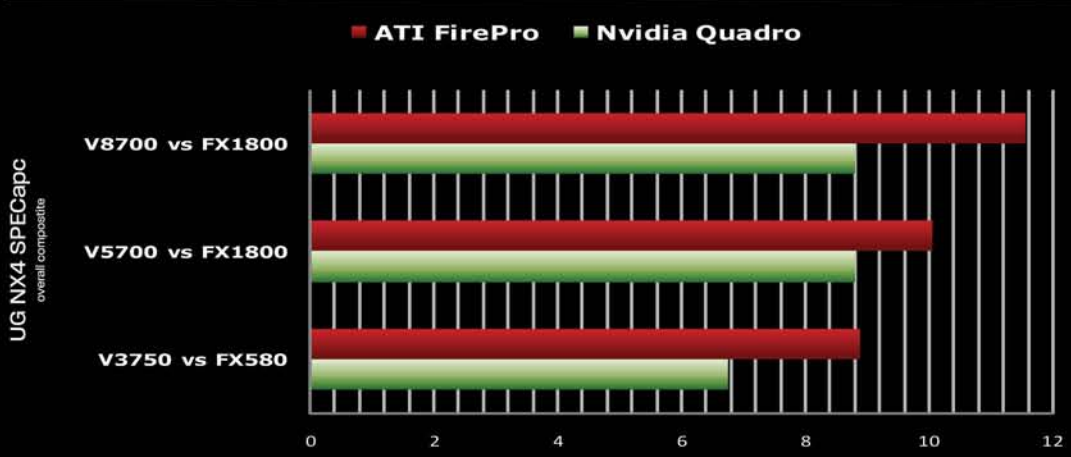

Learn more at ati.amd.com/FirePro

Where Do I Go for Data Acquisition Products?

omega.com, of Course!

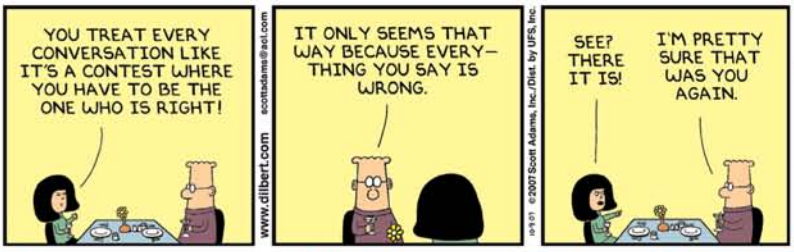

Need More DILBERT? Visit omega.com to order BLUECAT® New Horizons® in Data Acquisition Systems Version No. 23 — featuring 120 DILBERT CARTOONS!

"Visit us at Semicon West 2009, Booth No. 5356"

FREE!

Go to: www.omega.com/dilbert for your daily dose of DILBERT!

CTXL-TRH-W *PATENTED
Microprocessor Based Portable Universal Circular Chart SUPERECORDER™
Starts at $795
MADE IN USA

iBTHX iServer MicroServer™
Barometric Pressure, Temperature and Humidity Transmitter
*PATENT PENDING
Starts at $295
MADE IN USA

8-Channel Thermocouple Input USB Data Acquisition Module
Starts at $325

USB-4718

OM-CP-RHTEMP2000
Starts at $399

MADE IN USA
Humidity/Temperature Data Logger with LCD Display, NOMAD®

iDRN Series
Starts at $250
MADE IN USA

Analog and Digital Output Signal Conditioners/Transmitters

Portable Handheld Data Logger
Starts at $995
OM-DAQPRO-5300
Includes software, 110 Vac adaptor, USB cable, manual

OM-CP-TEMP1000
Starts at $199

Submersible Temperature Data Logger NOMAD®
MADE IN USA

iTHX-W
iTHX iServer MicroServer™
Temperature and Humidity
Starts at $295
MADE IN USA

Miniature *WIRELESS* Thermocouple Connectors
The Smart Connector™

Starts at $79

MWTC Series *PATENTED
MADE IN USA
Manufactured By Omega

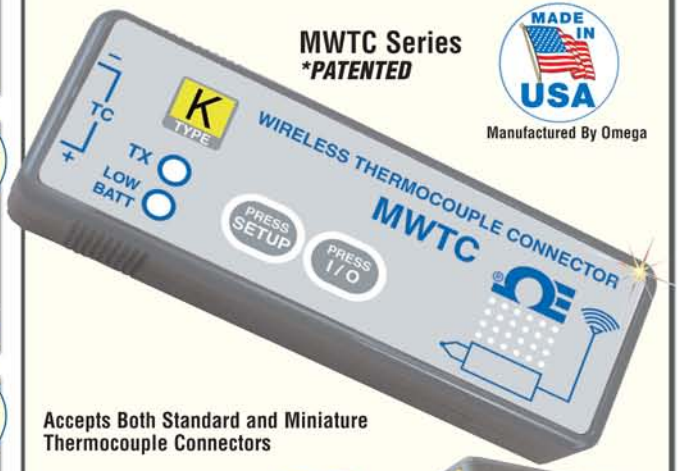

Accepts Both Standard and Miniature Thermocouple Connectors

- Available in Types J, K, T, E, R, S, B, C or N Thermocouple Calibrations
- Built-In Cold Junction Compensation and Linearization
- Model MWTC-REC1 Receiver Works with Multiple *WIRELESS* Remote Connectors
- Low Power Operation and Sleep Mode For Extended Battery Life

Your single source for process measurement and control products!

For Sales and Service, Call TOLL FREE
1-800-327-4333SM
1-800-DAS-TEEE

Direct Link: omega.com/709DK3

*PATENTED NOTICE
Covered by U.S. and International patents and/or patents pending

Shop Online at
omega.com®
OMEGA®

Dilbert © United Feature Syndicate, Inc.

Getting the Most from Parallel Systems

Peter Varhol
pvarhol@deskeng.com

The next generation of high-performance computing is upon us, and we must figure out how to make use of it. This generation involves slower processors, but these processors have multiple cores, or self-contained execution engines. My dual-core laptop is now three years old and I anticipate that my next purchase will have either four or eight cores.

The problem is that actually using these cores in the acceleration of one or more applications is highly problematic. Most engineering applications are single-threaded, in that they follow only one execution path at a time. By their nature, they can execute on only a single core while the rest stand more or less idle.

A few applications are multithreaded and each thread can theoretically be scheduled to run on different cores. This is somewhat of an improvement as long as the operating system can schedule individual threads and the threads accomplish a lot of work independently of one another.

> MatLab is an easy way to take advantage of multicore and parallel systems.

The MathWorks is trying to change that. I had the opportunity to visit the sprawling campus of this engineering software company in suburban Boston and speak with high-performance computing manager Silvina Grad-Freilich. MATLAB is already multithreaded, she explained, meaning that it can execute parts of the code that it knows can be executed independently in parallel.

One capability of MATLAB is that its Parallel Computing Toolbox can automatically parcel out work based on the availability of units of computation whether they are cores, clusters, or grids. Further, the Parallel Computing Toolbox enables users to run MATLAB applications in up to eight cores locally, taking advantage of the latest trends

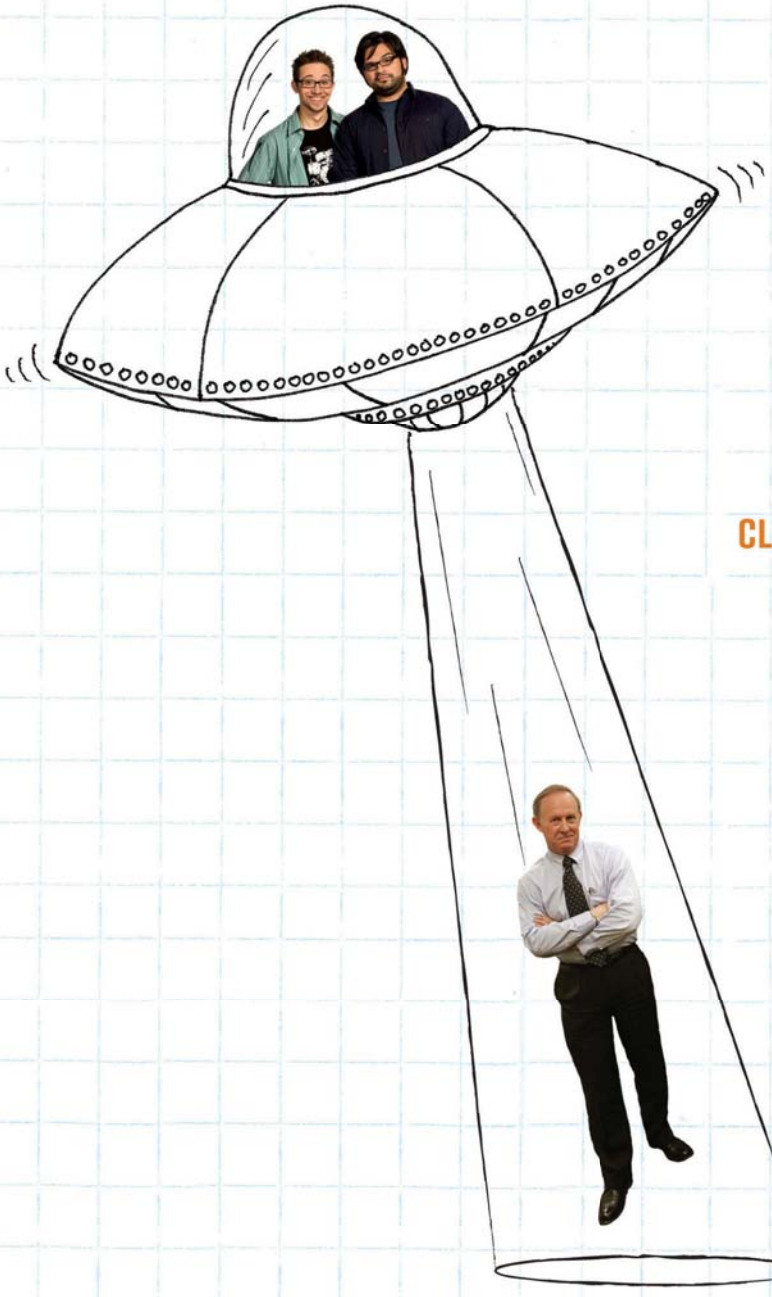

GET READY FOR MORE
CLOSE ENCOUNTERS OF THE CAD KIND.

A NEW EPISODE HAS INVADED.

3DUDES GONE 3D

SolidWorks® software presents a new series about three PRODUCT DESIGNERS forced to work in an overflow trailer. Starring a USER COMMUNITY OVER ONE MILLION STRONG, a friendly NEIGHBORHOOD SOLIDWORKS RESELLER, and more than 150 USER GROUPS. And featuring ACTIVE ONLINE DISCUSSION FORUMS, a set of CUSTOMIZED TRAINING CLASSES, and one CROWD-PLEASING RETURN ON INVESTMENT.

NR
NOT RATED

WATCH NOW AT 3DUDESGONE3D.COM

SolidWorks

in desktop system processors.

For those of you who write your own MATLAB code, you can add some code to help it even more. Using parallel commands as a part of, for example, FOR loops (assuming that the repeating computations are independent), you can explicitly tell those loops to execute simultaneously on whatever resources are available. Rather than FOR, the MATLAB instruction is PARFOR. And any code will execute on all defined resources in the project.

There is more to analysis than code; you can also have extremely large datasets. By breaking up problems on different computers in a cluster or grid, you can keep a smaller set of data in memory during parallel computation and still collect the data and analyze it later. If you can't run on a desktop system because of memory limitations, maybe now you can.

The end result is you can take existing MATLAB code and run it in parallel on multiple cores, often with few or no changes. Depending on where you are executing this code, you can get some pretty significant performance improvements—and best of all, largely using your existing MATLAB routines.

While not everyone uses MATLAB, it provides engineers with the opportunity to make full use of the power of their desktop computers and other inexpensive clusters and grids, including the Amazon EC2 cloud. Just be prepared to say goodbye to renting time on your favorite supercomputer. ■

Peter Varhol has been involved with software development and systems management for many years. Send comments about this column to DE-Editors@deskeng.com.

CONTENTS

JULY 2009 • VOLUME 14 • ISSUE 11

Desktop Engineering®

DESKENG.COM

"Computer science is no more about computers than astronomy is about telescopes."

> E.W. Dijkstra

COVER STORY

18

LARGE-FORMAT PRINTERS

18 Bigger, Bolder Printers Fill the Market > David Essex

Fast, affordable color comes to large-format printers tailor-made for technical documents.

22

PLM

22 ENOVIA SmarTeam: Instant PLM for SMBs > Jos Voskuil

Engineering Express helps manage bills of materials within PLM, offering added efficiency to mid-market companies.

25

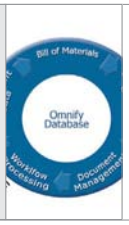

PLM

25 Omnify Brings Savings and Clarity to Crystals > Alaine Portnoy

Manufacturer of oxide single crystals cuts engineering change cycles by 16 percent.

30

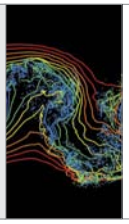

VISUALIZATION

30 Tecplot: Calculating Flow Separation > Vince Adams

Researchers relieve the pressure in calculating flow separation using Tecplot 360.

35

CFD & VISUALIZATION

35 Getting More from Less with Fluent and EnSight > Kara L. Gray

Missouri Science and Technology racing students use CFD and visualization to gain a leg up on the competition.

39 

CAE

39 Leveraging Abaqus CAE During a Downturn > Frank Popielas

CAE and analysis-based design represent technology that can help you weather the recessionary economy.

42 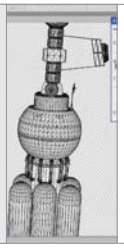

RENDERING

42 Photoshop CS4 Extended Improves Flow > Mark Clarkson

A strong upgrade with new features for transforming your rendering.

48 

WORKSTATION

48 HP xw8600 Makes Perfection Even Better > David Cohn

HP updates its top-of-the-line workstation with the latest Intel Xeon quad-core CPU and NVIDIA Quadro FX graphics board.

Elements of HPC

54 

54 Scaling Beyond the Desktop with MATLAB >Chakravarti, et al

Parallel MATLAB applications make programming simple and portable.

58 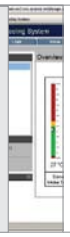

58 More than a Question of Energy > Tom Kevan

Performance, power consumption, and cooling make up a complex stew, leaving the search for energy efficiency with no simple recipe for success.

64 

64 Multicore Matters > Peter Varhol

Multicore processors will change the way we think about engineering computing, but only when software catches up.

71 

PRODUCT OF THE MONTH

71 New Autodesk Upgrades for Building Information Modeling

Readers picked Revit Structure 2009, Revit MEP 2009, and NavisWorks 2009

75 

Commentary

75 Successful PDM Implementation for SMBs > Dave Chadwick, Siemens PLM Software

EDITORIAL

Steve Robbins — Executive Editor
Jonathan Gourlay — Senior Content Manager
Anthony J. Lockwood — Editor at Large
Margaret S. Gurney — Copy Editor

CONTRIBUTING EDITORS

Mark Clarkson • Al Dean
David S. Cohn • Mike Hudspeth
Tom Kevan • Susan Smith
Peter Varhol • Pamela J. Waterman
Kenneth Wong

PUBLISHER

Brian Vaillancourt (x263)

ADVERTISING SALES

603-563-1631 • Fax 603-563-8192
Brian Vaillancourt — Publisher (x263)
Gail Robinson — Sales Manager (x252)
Jeanne DuVal — Account Manager (x274)

ART DEPARTMENT

Darlene Sweeney — Art & Production Director (x257)

A LEVEL 5 COMMUNICATIONS PUBLICATION

Steve Robbins — Chief Executive Officer
Thomas Conlon — President

ADVERTISING, BUSINESS, AND 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

BPA WORLDWIDE

COMSOL CONFERENCE WORLDWIDE

Milan · Boston

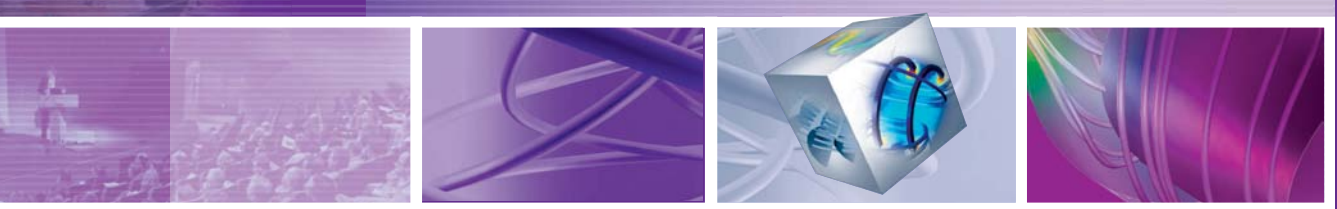

Join the Leaders in Multiphysics Simulation at the COMSOL Conference 2009

BOSTON, OCT 8-10 / MILAN, OCT 14-16

Hands-on Minicourses
Get hands-on training in more than a dozen topics

Keynote Talks and User Presentations
Tap into users´ applications and research

Exhibition & Poster Session
Showcase of simulation products and services

Demo Stations
Test drive new products and meet with support and application engineers

Tutorial Presentations
Demonstration of specific modeling and simulation techniques in COMSOL Multiphysics®

Awards Banquet
Meet and interact with your colleagues.
Prizes for best papers and posters

GOLD SPONSORS

MatWeb
ALTASIM TECHNOLOGIES
Microsoft

BRONZE SPONSORS

The MathWorks
simpleware
Jones and Bartlett Publishers
World Scientific
SolidWorks

MEDIA SPONSORS

Tech Briefs
Design World
DE
Microwave Journal
Chemical Engineering Progress

www.comsol.com/conference2009

COMSOL

editorial index

>Analysis 10, 30, 35, 39

>CAE 35, 39, 64

>Commentary 2, 75

>Computing Systems 2, 54, 58, 64

>CFD 30, 35

>Design & MCAD 25, 30, 39, 42

>FEA 39

>HPC 2, 53, 54, 58, 64

>Innovation 2, 14, 30, 54, 58, 64

>Manufacturing 10, 14, 25

>New Products 14, 18, 42, 48, 70

>PLM 22, 25, 75

>Large-Format Printers 18

>Reader's Choice Award 71

>Rendering 42

>Visualization 30, 35, 42

>Workstation 48

departments

DOF 2

BRIEFINGS 10

EDITOR'S PICKS 15

AD INDEX 72

DE Desktop Engineering®

online @ deskeng.com

>New Blogs
Read Tony Lockwood @ deskeng.com/burning_questions and Kenneth Wong @ deskeng.com/virtual_desktop.

>DE Marketplace see page 47
>DE Product Showcase
see page 74
Hardware, software, and publications. Live links connect to suppliers.

>SpecSearch
Search for products and services from more than 11,800 supplier catalogs.

>Newsletter Registration
Editor's Pick of the Week, Elements of: Analysis, MCAD, Engineering IT & Computing, Rapid Technologies, Newslink

ON THE COVER> Large-format printers that are fast are commonplace in engineering. What isn't so common, however, is color, but recent advances are changing color's role in technical communication. Here's a look at a selection of the latest in fast large-format printers tailor-made for technical documents. *To read David Essex's article, turn to page 18.*

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

Desktop Engineering® (ISSN 1085-0422) is published monthly by Level 5 Communications, Inc., 1283D Main Street, P.O. Box 1039, Dublin, NH 03444, 603-563-1631. Periodicals postage paid at Dublin, NH, and at additional mailing offices. Desktop Engineering® is distributed free to qualified U.S. subscribers. Subscription rates for non-qualified: U.S. $108 one year; Canada and Mexico $126 one year; all other countries $195 one year.
LIST RENTALS: For information on list rentals, contact Statlistics, Danbury, CT: 203-778-8700.

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

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

Revolutionize how you bring products to market!

MapleSim 2

High-Performance Multi-Domain Modeling and Simulation

Key Advantages of MapleSim:

- Mix physical components with signal-flow blocks
- Multi-domain models are assembled from pre-built components
- Model diagrams map onto the real system
- System equations are automatically generated and optimized
- Integrated with powerful analysis and documentation tools
- 3-D animation delivers immediate insight

Now Available

MapleSim™ is a high-performance tool for multi-domain physical modeling and control systems development. The world's most advanced symbolic computing engine comes together with traditional numeric solvers to supercharge the simulation and modeling process.

To request a free evaluation copy of MapleSim 2, go to www.maplesoft.com/de

What's New in MapleSim 2:

- 3-D Visualization and Animations: Transform your multibody models into realistic animations
- Result Management Tools: Easily manage the results from multiple simulations

Request your free evaluation copy of MapleSim 2!
Go to www.maplesoft.com/de

Maplesoft
Mathematics • Modeling • Simulation

www.maplesoft.com | info@maplesoft.com
Toll-free: (US & Canada) 1-800-267-6583
Direct:1-519-747-2373

ADDITIVE MANUFACTURING SHOWS GROWTH OPPORTUNITY

Wohlers Associates, Inc. has announced the publication of *Wohlers Report 2009*, a 250-page analysis of the newest developments and trends in additive manufacturing (AM) worldwide.

AM systems build parts layer by layer. They are used for a range of applications, including design review and validation, fit and function prototyping, pattern making, and tooling. Increasingly, AM is used for custom and replacement parts, limited and short-run production, and sometimes full production.

The global market for AM products and services grew to an estimated $1.183 billion in 2008, according to *Wohlers Report 2009*. The services segment of the industry was a bright spot, with growth of 7.9 percent. The AM industry is expected to more than double in size by 2015, according to the report.

The use of additive manufacturing for direct part production has grown to become the second most popular application of AM technology. Companies representing thousands of users and customers of AM technology from around the world responded to a survey on the subject. Respondents believe that part production from AM will represent 35.9 percent of their business in five years and more than half in 10 years.

"Methods of additive manufacturing are creating new markets for limited-edition and one-of-a-kind products," said Terry Wohlers, principal author of the report and president of Wohlers Associates. "Additive manufacturing is causing some companies and individuals to rethink the way products are developed and brought to market."

The study was created with support from 54 co-authors, 65 service providers, 29 system manufacturers, and others. The report includes 31 charts and graphs, 45 tables, and 154 photos and illustrations. The report sells for $475 in the U.S. and $495 in other countries.

INFO

Wohlers Associates, Inc.
Fort Collins, CO
wohlersassociates.com

CIMDATA REPORTS NC MARKET GREW BY 5.2 PERCENT IN 2008

CIMdata estimates that, based on end-user payments, the worldwide NC software and related services market grew by 5.2 percent in 2008 to reach a level of $1.47 billion. CIMdata projects that in 2009, the growth rate of these payments will slow to 1.3 percent as the market approaches the $1.5 billion level. The expected reduction in the market growth rate in 2009 is a result of the downturn in the global economy—specifically the manufacturing industry.Since a similar downturn in the 2000 to 2002 period, NC software has shown modest but steady growth as global economies improved. There was worldwide growth in the sale of machine tools and manufacturing output, greater emphasis was placed on the efficient operation of machine tools as manufacturing firms strengthen their competitive position, and the overall product lifecycle management market, of which CAM is a component, continued on a strong growth path.

INFO

CIMdata
AnnArbor, MI
cimdata.com

SME'S RAPID 2009 & 3D IMAGING CONFERENCE AND EXPOSITION INSPIRES ATTENDEES

More than 1,300 manufacturers and engineers attended the conference sessions and worked the exhibit hall at the recent RAPID 2009 and 3D IMAGING show May 12-14 with energy and enthusiasm. The conference, produced by the Society of Manufacturing Engineers (SME), took place at the Renaissance Schaumburg Hotel & Convention Center in Schaumburg, IL.

"Even in the midst of a depressed economy, you could feel the buzz in the air fueled by passion and innovation," said Todd Grimm, chair of SME's Rapid Technologies and Additive Manufacturing Technical Community. "You could sense that attendees were there to make changes to the way they do business, which can help them come out of this recession strong."

Described by one attendee as one of the best keynote speakers in recent memory, Mike North, formerly host of Discovery Channel's "Prototype This," set the tone for the event with his vibrant and dynamic "Get 'Er Done" presentation on day one. North discussed the challenges and resources needed to create innovative machines at breakneck speeds, and talked about the implications of the home inventor having access to a wide range of one-off manufacturing tools, causing a paradigm shift from a Do-It-Yourself to an Outsource-It-Yourself mentality.

Sharing their wisdom, vision, and inspiration were four of the founders of rapid prototyping and 3D imaging technologies—Marc Soucy of InnovMetric Software; Ping Fu, Ph.D., of GEOMagic; Chuck Hull of 3D Systems Corporation; and Carl Deckard, Ph.D., a past chairperson of SME's Rapid Technologies and Additive Manufacturing community and president of Met-L-Flo, Inc.

New conference tracks that featured rapid implant manufacturing (coordinated with Materialise), metal casting (with the support of the American Foundry Society), and architecture, brought in attendees with different views and fresh ideas beyond the typical design and manufacturing engineering discussions.

DE's Steve Robbins at the BJB Enterprises booth wondering where the gin went.

Two forums hosted by Materialise (Schaumburg, IL, materialise.com), Medical Implant Scanning and Manufacturing as well as the Rapid Implant Manufacturing Forum, were deemed successes by the company. More than 30 people attended the Medical Implant Scanning and Manufacturing Workshop to learn about the process of producing custom implants from medical scans using rapid technology. The workshop was led by Koen Engelborghs of Materialise, Andy Christensen of Medical Modeling, Ola Harrysson of N.C. State University, and Dr. Stephen Rouse of Walter Reed Medical Center. The Rapid Implant Manufacturing Forum was a one-day conference track attended by more than 70 people. It addressed the latest trends in rapid manufacturing and custom implants. Participants learned about the diverse uses of image processing and manufacturing from various experts in the industry.

INFO

Society of Manufacturing Engineers
Dearborn, MI
sme.org

ANSYS SOFTWARE DESIGN KEEPS FLOODWATERS AT BAY

ANSYS, Inc.'s software has been successfully used to build pumping station platforms in storm-prone New Orleans. The platforms were designed to withstand hurricane-force winds by a group of engineers charged with improving water drainage pumping capacity who were faced with a number of design challenges. Because the platforms are partially submerged, all design problems had to be identified and addressed before any construction began. In addition, the project needed to be completed prior to the onset of hurricane season. Work was contracted after one of the costliest and deadliest hurricanes in the U.S., Katrina, devastated the Gulf Coast.

When Category 2 hurricane Gustav hit in 2008, the platforms and pumps successfully kept floodwaters in check.When a hurricane or other major weather event hits, high-power pumping units must work at full capacity to drain excess water out of sub-sea level New Orleans. The heavy equipment—turbine pumps, diesel engines, gearboxes, and associated piping—produce vibrations and other stresses that can cause the massive platforms supporting the equipment to fail. Design and analysis consultant Mechanical Solutions, Inc. (MSI), working for the platform design-build contractor Weston Solutions, Inc., used software from ANSYS to simulate these stresses in the pumping platform along the

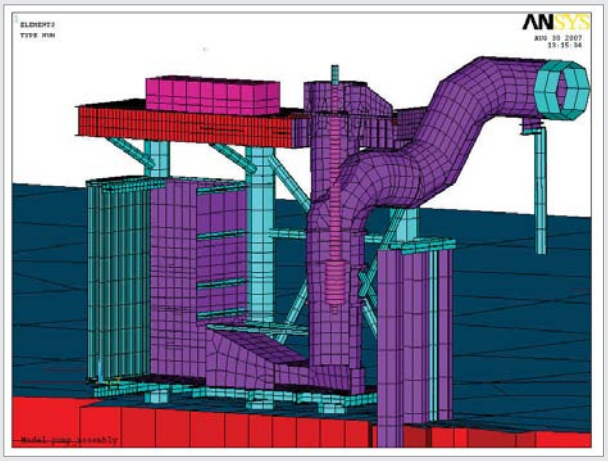

17th Street Canal, ensuring that the system would perform optimally during storms.

"Our approach with every project is to bring value to the table by identifying and mitigating risk early on to reduce the potential for costly modifications after construction has been completed," said William J. Kelly, principal engineer, Mechanical Solutions, Inc. "By using software from ANSYS to simulate the vibrational response of the platform, we were confident that the pumping system would perform optimally, and it did. In addition, engineering simulation helped us meet a rigid deadline by identifying issues early in the design process."

For more information, visit ANSYS.

INFO
ANSYS
Southpointe, PA
ansys.com

CORRECTIONS

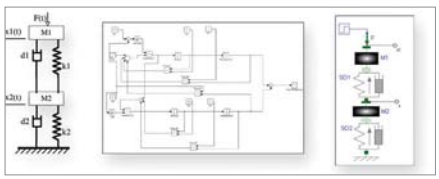

We published the wrong image in an article on ***MapleSim*** in the May issue on page 37. The image that appears on that page is not the MapleSim graphical working environment. This is the image that should have been printed with this caption: Physical model diagrams generated by MapleSim (on the right), unlike those created by traditional tools (center), map directly to the original system (left).

In June, a caption on page 34 mistakenly cited ***Abaqus*** 6.5. Abaqus is currently in version 6.9 and includes failure capabilities to make assessments straightforward and efficient.

RIGHT HEMISPHERE UNVEILS DEEP EXPLORATION 6

Client Software from Right Hemisphere features new user interface, rendering engine, and integration enhancements. Right Hemisphere announced a new release of its Deep Exploration client software. Deep Exploration 6 has evolved to become integrated into Right Hemisphere's visual solutions and supporting technology infrastructure. New features include a configurable workflow-oriented user interface (UI); improved ability to view, select, and navigate 3D models; a new XML file format; and a new rendering engine.

"Unifying visual and text-based product information is critical to enabling the Visual Enterprise. What Deep Exploration does for the Visual Enterprise is deliver navigation, authoring, and repurposing of complex product information across a wide variety of business workflows," said Mark Thomas, founder and CTO at Right Hemisphere.

Deep Exploration 6 has a redesigned UI to support a variety of different end-user tasks, workflows, and deliverables. The new UI allows users to optimize their screen real estate for higher productivity. The heart of the new UI is the Layout Manager that manages the visibility and order of the application's tool tabs. The tool tabs present only the capabilities and panels required for the specific task at hand and guide the user through the necessary steps to create an end product or deliverable such as a 2D technical illustration, a set of manufacturing work instructions, or a photorealistic 3D image.

Deep Exploration 6 also includes a new rendering engine that employs high dynamic range (HDR) imaging technology. HDR imaging techniques support a wider range of luminance in a 3D scene—from direct sunlight to subtle shadows. With this HDR rendering engine, Deep Exploration 6 users can drag and drop HDR materials, environment maps, and background images also known as "back plates" into a 3D scene.The product also contains new

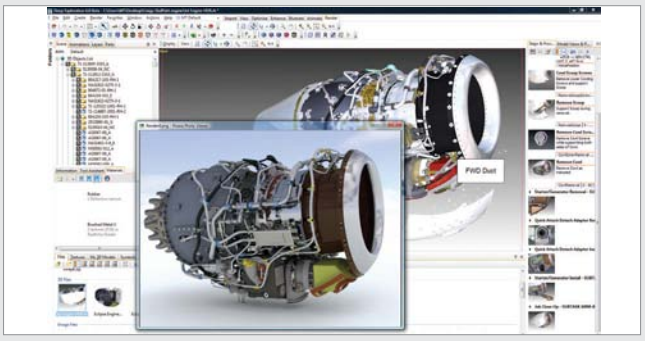

tools and features to enable users to navigate, select, and review large, complex 3D models. The navigation of any model is enhanced with the new Volume Select tool, which lets users choose an area of the model by defining variables of the "geometric envelope"—the area inside a square, sphere, or rectangle.

In other news, Right Hemisphere has teamed up with EMC Corporation to offer a 3D product image management solution. The solution is based on Right Hemisphere's Deep Server enterprise software and the EMC Documentum Enterprise Content Management (ECM) Suite. The combined offering will allow companies to take 3D CAD data from engineering and directly integrate it with EMC Documentum. This 3D product content can then be viewed, searched, manipulated, shared, published, and managed inside EMC Documentum. The result is an end-to-end 3D file and content management system for non-engineering users responsible for service procedures, technical publications, and marketing deliverables.

Deep Server enterprise software complements EMC Documentum with its ability to read more than 80 2D and 3D file formats. This includes support for the commonly used CAD applications such as CATIA, ProENGINEER, SolidWorks, and AutoCAD. For more information, visit EMC and Right Hemisphere.

INFO

EMC
Hopkinton, MA
emc.com

Right Hemisphere
Pleasanton, CA
righthemisphere.com

BOOTHROYD DEWHURST DESIGN FOR MANUFACTURE & ASSEMBLY FORUM CENTERS ON U.S. MANUFACTURING

More than 75 engineers and analysts charged with cost reduction and manufacturing attended the International Forum on Design for Manufacture and Assembly (DFMA) in Providence, RI, last month hosted by Boothroyd Dewhurst. Attendees spent two days learning about the value of cost modeling and how DFMA software has benefited users.

To begin the proceedings, Richard McCormack of *Manufacturing & Technology News* painted a bleak picture of the decline of U.S. manufacturing and pointed at both the Bush and Obama administrations as sharing a good portion of blame. McCormack's session was drawn largely from testimony he gave before the U.S.-China Economic Security Review Commission in March.

Mike Shipulski of Hypertherm followed up with a markedly more optimistic look at manufacturing and made the case that a systematic DFMA deployment will reduce parts in assemblies, cut production costs, and improve product quality. His own use of the software, he says, leads him to believe a similar rigorous process can guide companies toward increased profitablility. "Why aren't you guys doing this?" he asked, pointing to a host of metrics that supported his presentation.

A day-two panel with McCormack, Shipulski, Ned O'Donovan of Rensselaer Polytechnic, and Chris Tsai of Global Productivity, Inc. turned into an animated discussion of insourcing, and offshoring. While the offshoring trend still continues, the panel seemed to agree that U.S.-based design might help stem the flow of job losses before bringing manufacturing back via targeted cost reductions. Intellectual property issues and the Chinese culture of engineering were also discussed.

INFO

Boothroyd Dewhurst, Inc.
Wakefield, RI
dfma.com

MSC.SOFTWARE HOSTS WEBINAR SERIES

With more than 15,000 medical device manufacturers worldwide competing to bring new and increasingly complex device concepts to market first, more and more are turning to simulation-driven development to bring new product concepts to life faster and at less cost. In a webinar series from MSC. Software, attendees will learn from industry experts how simulation-driven development is used by medical device manufacturers to gain a competitive edge.

Remaining webinars take place September 8: Applying Simulation-Driven Design Methods to Optimize Dental Restorative Procedures; and October 13: Blood Flow Analysis for Cardiovascular Applications Increases Understanding of Device Performance. For more information or to register, visit MSC.Software.

INFO

MSC.Software
Santa Ana, CA
mscsoftware.com/events

Would you trust this guy? Well that question has already been answered by thousands of readers who have indicated they already do, implicitly. So here are Lockwood's most recent musings about the products that have really grabbed his attention, and deserve yours.

Reverse Engineering Tool Integrates into SolidWorks

Scan, capture, measure, inspect, and compare model to part directly.

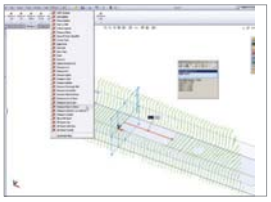

The folks at Reverse-Engineering.com have been busy lately, and their latest product is called ReSpect. Out of r.e.s.p.e.c.t to them an Aretha Franklin, I'll skip a song and dance filled with puns.

ReSpect is a plug-in reverse engineering toolset for SolidWorks that's all about respecting your time as a designer tasked with getting a physical object scanned and prepped for the next step. "Next step" can be MRO (maintenance, repair, and overhaul), first or periodic part inspection, prototype to CAD comparison, alignment checking, or legacy part reconstruction. Doesn't matter. ReSpect honors your time by enabling you to scan an object into SolidWorks directly where you can begin working on it at the same time.

And by "working on it" I mean that you can measure features, inspect points of interest, and compare CAD-to-parts—in real time. ReSpect aligns the scan data to your model as you scan it. You do your checking and update your SolidWorks design simultaneously. Since ReSpect integrates with SolidWorks and with your portable CMM natively, you recoup additional time by skipping the scan data translation step and the seek-and-repair data translation noise processes. .

Read my message in full at DE Online:
http://www.deskeng.com/articles/aaapxt.htm

ANSYS 12.0 Launched

Every major physics discipline in engineering simulation environment suite enhanced.

It all comes down to competition in what's turning out to be a completely new ballgame in these challenging times. Yesterday, new rules governing fuel economy and automobile emissions for all new trucks and cars sold in the US joined the push for alternative energy sources, smart grids, plug-in cars, maglev trains, and so on. That's a lot of design and analysis of complex systems, and it needs to be done quickly—2016 for the new emissions and fuel standards announced yesterday. Not only are all of these engineering challenges, they are all multiple-physics challenges. That is why it's important news that ANSYS 12.0 has come out.

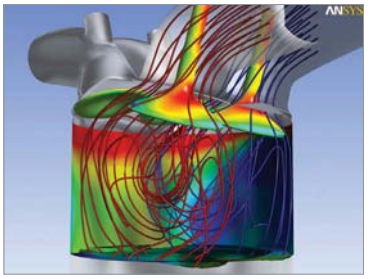

ANSYS Inc. has been at the forefront developing integrated high-level engineering software for years because, simply, we live in a multiple-physics world. Over time, this vision naturally led the company to its ANSYS Workbench platform that integrates multiple-physics disciplines so that you can create complex products and know how they will perform in the real world when subjected to multiple-physics forces.

As the ANSYS CAE center of gravity, Workbench is where design, physics, and data come together so that you can study, prototype, and optimize designs. In version 12.0, ANSYS has enhanced Workbench extensively. The most notable enhancement is the new project schematic functionality in your project page. If you're familiar with NI's LabVIEW, you have the basic idea. You use the Workbench schematics to connect systems and, as you do, you build complex analyses of multiple-physics phenomena. But it's more than that. It's like a dashboard: It tells you data dependencies and the state of a cell—e.g. the cell is up to date, meshed, needs to be solved, etc.—and you can launch applications to do whatever it is you need to do.

Read my message in full at DE Online:
http://www.deskeng.com/articles/aaardc.htm

Nastran FEA Toolset Comes with MCAD

NEi Fusion has a SolidWorks engine, making it a full analysis and simulation system.

I get spooked beginning something using words like "get started,""combination of technologies," and the like because readers can interpret that to mean "stunningly dumbed down" or recall some useless multi-tool they once got on their birthday. So, you are forewarned: This is not that sort of thing. So here goes.

NEi Fusion Designer brings a combination of technologies—3D parametric MCAD powered by SolidWorks and Nastran FEA—together so that you can get started using FEA early and often then throughout the product development process. Think of "get started" in two ways. One, NEi Fusion is engineered to get the non-FEA savvy designer working productively with FEA right away. Two, it is also a tool—a set of tools,really—that enables small- and mid-sized outfits, the recently laid-off who are picking up some jobs as contract "consultants," and so

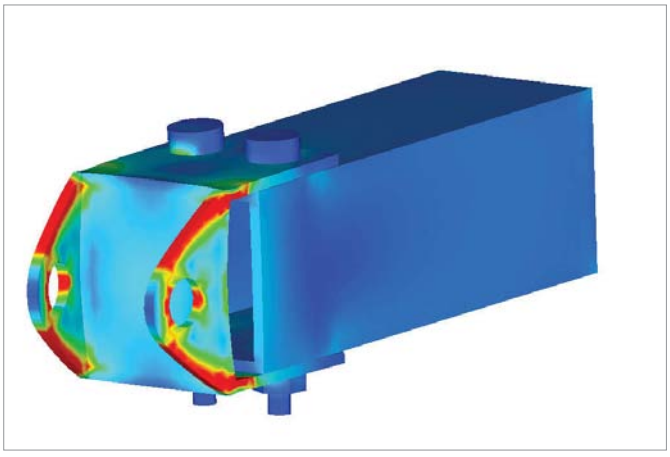

on to bring robust, professional-level Nastran-based FEA into their quiver without getting all a-quivering over the money they shell out for it.

NEi Fusion Designer provides the most used and needed analysis capabilities—linear statics, steady state heat transfer, modal, buckling, and prestress.

Read my message in full at DE Online:
http://www.deskeng.com/articles/aaaran.htm

Parametric Exchange Speeds Scan to CAD Process

Geomagic Studio 11 enhances reverse engineering, product design, and rapid prototyping workflows.

A couple of weeks ago, Geomagic announced a new version of its Geomagic Studio, its software for creating 3D digital models and surfaces out of scanned data. Geomagic has been pushing the envelope on what you can do with scanned data for quite some time now, and version 11 is yet more proof of that.

Geomagic Studio 11 has a lot of enhancements

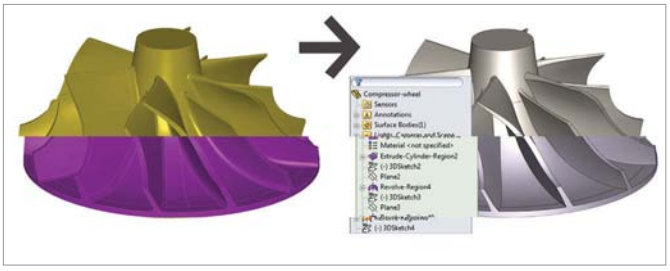

intended to compress design cycles—things like a new tangent hole capability that helps you fill in complex, missing data in your scan data. Also new are interactive polygon sculpting tools for editing and reshaping models in early design stages, as well as curvature-sensitive smoothing, which is said to increase polygon mesh quality by polishing noisy data in flatter areas while preserving highly curved areas and details. Version 11 also introduces you to the Mesh Doctor, which automatically detects and corrects errors in a polygon mesh. I suspect that we're just getting to know the Doctor, so keep your eye on it in future releases.

But what you really should eyeball in Geomagic Studio 11 is its new parametric exchange functionality. Let's stop here a second. The first thing to keep in mind is that Geomagic Studio lets you extrapolate design intent as well as create optimized CAD surfaces from your scanned data. Now, after you're done using Geomagic Studio to, for example, optimize a swept or freeform surface, fix meshes, and fill in holes, you then use its parametric exchange capability to convert and export your Geomagic model to native Inventor, Pro/ENGINEER Wildfire, or SolidWorks CAD geometry.

Talk about compressing the process. Lots of time savings available.

Read my message in full at DE Online:
http://www.deskeng.com/articles/aaardz.htm

What if it were possible...

to get custom parts ***so fast***
that you ***actually*** could

Save a life!

Well now you can.

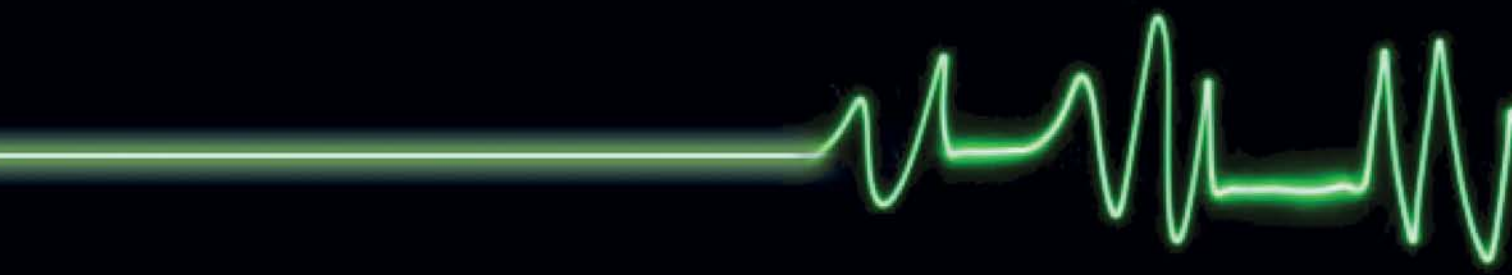

We Save You Time,
So You Can Save Lives.

Rapid Prototypes in **1 Day**
Machined Parts in **5 Days**
Cast Urethanes in **6 Days**
Injection Molded Parts in **10 Days**

QUICKPARTS

www.quickparts.com
770.901.3200

Bigger, Bolder Printers Fill the Market

> Fast, affordable color comes to large-format printers tailor-made for technical documents.

BY DAVID ESSEX

Large-format printers that are fast are commonplace in engineering. What isn't so common, however, is color, but recent advances are changing color's role in technical communication.

Their price premiums shrinking, color printers can reproduce black and white and gray scales as well as their monochrome counterparts, and their cost-per-page is comparable if you're mostly printing line drawings with little or no color.

The move to color is usually a move to inkjet, and HP dominates the market, although Epson, Océ, and especially Canon are players. Toner-based, or LED, printers, are the monochrome mainstay, with KIP America, Kyocera Mita, Océ, Ricoh, Seiko Infotech, and Xerox sharing the segment.

"Océ's really the only crossover company that does both," says Tim Greene, director of wide-format and jetting technologies at InfoTrends. Greene calls the KIP 3100 multifunction unit "a category leader" in monochrome LEDs.

But color LEDs are so expensive that they most often appeal to reprographic and other high-volume, quality-conscious customers. They are usually 3-4 times faster, largely because they lack the drying step that is inkjet's Achilles' heel, and thus don't require more expensive coated paper. Inkjet vendors have greatly improved drying, however, with new paper and ink technologies, Greene says, while boosting speed and resolution with wider, more precise nozzles. Canon and HP have economy modes that put less ink on paper, saving drying time and ink costs.

For high-volume engineering and design uses, HP has released the DesignJet 4020 (pictured) and 4050 printers with onboard processing boosted by hard drives and memory 2x that of the T1120.

BUYING ADVICE

Large format starts at 24 in. wide, though 36 in. is the minimum in most technical environments. Many vendors sell 24- and 44-in. options, and

a few go wider. Printers come in two configurations: multifunction (MFPs) or stand-alone. MFPs can replace analog copy machines by combining copier-like scanners and printers in a more compact unit. Higher-volume printers have multiple paper rolls that switch in automatically when one runs out.

Beware mismatches between the color abilities of scanner and printer. Some spec sheets tout color scanning, but the printer is actually monochrome.

Mechanical speed ratings can be misleading. What matters is "click to plunk"—the time between the mouse click and the plunk of paper in the tray. Printers vary in the amount of file processing they do onboard instead of on the computer. The best way to gauge throughput is to bring your files to the showroom and run print jobs yourself. Greene suggests quantifying monthly volume and extrapolating it to page-per-minute claims to avoid over- or underestimating your need for speed.

The baseline resolution for wide-format monochrome is 600dpi, with some Epson printers reaching 2,880dpi, according to Greene. By comparison, the sweet spot for color is 1,200 to 2,400dpi, says Marc Jongen, a worldwide marketing manager in Hewlett-Packard's technical printing division, and adding photorealistic rendering isn't possible at 600dpi. Higher resolutions reduce performance, and the tradeoff isn't usually worth the negligible increase in image quality. Xerox, for example, sells printers that go up to 1,440dpi, but most users are happy with 720, says Karen Serrano, Xerox's worldwide marketing and training manager for wide-format products.

CANON

Already HP's biggest challenger in color inkjets, according to Greene, Canon has upgraded its imagePROGRAF line substantially in the past year, targeting both the reprographic and technical markets. Introduced last fall, the 44-in. iPF810 and 820 can print two pages per minute at 1,200dpi and come bundled with RATIO's CADSTATION PS+ workflow software for managing a variety of graphical and technical file formats.

The compact Xerox 6279 is a 600dpi, monochrome printer and can print up to nine pages a minute while scanning at 6ips with its optional monochrome scanner, which is upgradeable to color.

HEWLETT-PACKARD

The longtime leader in office color inkjets, HP has been pushing further into specialized, wide-format markets. It introduced three stand-alone and three multifunction models in March 2009 and upgraded its inkjet control software to LED-like functions, Greene says.

"The new printers have major improvements in their productivity," says Jongen. He notes that while mechanical speeds haven't changed from earlier models, onboard file processing has been improved. This helps minimize the time waiting in line, and smaller file sizes free your

workstation more quickly for other work. Web-based control software makes it easier to track consumables remotely, he says, and service contracts are rarely needed, unlike with LEDs.

Of the three new MFPs, the one targeting engineers is the DesignJet T1120 SD MFP ($16,495), a 1,200dpi color unit that contains a contact image sensor (CIS) scanner, a more affordable technology than the charged-coupled device (CCD) technology in many scanners.

"It's great for line drawings for CAD applications," Jongen says. The SD unit sells for roughly $5,000 less than the previous comparable model and its small-footprint, all-in-one design is perhaps best suited for workgroups that have limited space possibilities.

A $21,495 HD version uses CCD for sharper scan resolutions of up to 400dpi, twice that of the SD. Both can print line drawings at slightly fewer than two D-size pages per minute, and use a six-color system that adds gray and matte black to the usual CMYK. Both come in 24- or 44-in.-wide versions. You can get the printer by itself for as little as $3,595.

For high-volume engineering and design uses, HP also released the DesignJet 4020 and 4050, printers with onboard processing boosted by hard drives and memory with more than twice the capacity of the T1120.

The company also introduced large-format "cockle-free" coated papers (they stay flat after printing, even with lots of ink) and recycled bond papers that improve print quality and are geared to technical applications, according to Jeanette Volk, HP's sales development manager for large-format media.

The Océ ColorWave 600 is a mid-volume LED printer with a six-roll capacity. Some see it as a best-of-both-worlds replacement for separate monochrome LED copiers and color inkjet machines.

HP also sells entry-level printers for less than $2,000 for individual users and small workgroups, such as the 24-in. DesignJet 110plus and 36- and 42-in. 510 model.

OCÉ

Noting a trend away from centralized print rooms and toward distributed printing, Océ has designed its software interface and printer controls to make it easier for walk-up users to scan, copy, or print, says Penny Holland, director of business development at Océ. The printers have dual hard drives and other features that allow them to scan, process, and print multiple files simultaneously and avoid bottlenecks. On monochrome models, which are all LEDs, Radiant Fusing technology

replaces heated rollers to provide instant-on printing and better reliability and energy efficiency, according to Océ.

LED printer and multifunction models cover the gamut of capacities, from the entry-level, TDS320, which scans, copies, and prints up to 9 linear feet per second, to the 600x1200, mid-to-high-volume TDS700 and high-end TDS800 printers.

For color, Océ sells the TCS 300 and 500 inkjets and some LEDs. Holland says engineers should consider the ColorWave 600, a mid-volume LED printer with a six-roll capacity, as a best-of-both-worlds replacement for separate monochrome LED copiers and color inkjet machines.

The newest offering, introduced in late April, is the PlotWave 300. The printer-only version costs $17,000, while multifunction systems cost $22,000 to $23,700 depending on the paper rolls. Holland says the top-mounted scanner in the multifunction version makes it well suited for walk-up, departmental use.

XEROX

Xerox touts its ease of use, thanks in large part to a simple graphical user interface, along with a modular approach that makes it easy to add features and capacity as needs change.

Its lead product in the low-volume arena is the Xerox 6204, a monochrome LED-based MFP that outputs 4 to 5 pages per minute, Serrano says. It starts at $18,325.

The Xerox 6279, a 600dpi, monochrome unit, can print up to nine pages a minute while scanning at 6ips with its optional monochrome scanner, which is upgradeable to color on site. It starts at $23,000 and is considered a mid-range volume unit, and will replace the 6050A and 510, Serrano says. "Over time, the 6279 will replace all products worldwide," because, she says, the 6279 is more compact than previous, and competing, multifunction systems. "In the past, they might not have been able to get that multifunctionality in a single footprint."

Serrano says Xerox's high-volume machine is the 721, which spits out a whopping 22 pages per minute. Priced at $75,200 for its most popular configuration, it appeals mostly to reprographic businesses, but is sold into some high-volume engineering shops, she says.

Whether you need to print lots of color or high-quality documents, there's almost certainly a large-format machine to fit your needs. ■

David Essex *has covered information technology for 23 years. He was a* BYTE *editor and has written for* Computerworld, PC World*, and numerous other publications. Send him an e-mail about this article to DE-Editors@deskeng.com.*

INFO

Canon USA
Lake Success, NY
usa.canon.com

Epson America
Long Beach, CA
epson.com

Hewlett-Packard
Palo Alto, CA
hp.com

InfoTrends
Weymouth, MA
infotrends.com

KIP America
Novi, MI
kipamerica.com

Kyocera Mita America
Fairfield, NJ
usa.kyoceramita.com

Océ
Trumbull, CT
oce.com

Ricoh Americas
West Caldwell, NJ
ricoh.com

Xerox
Norwalk, CT
xerox.com

For more information, visit deskeng.com.

ENOVIA SmarTeam: Instant PLM for SMBs

> Engineering Express helps manage bills of materials within PLM, offering added efficiency to mid-market companies.

BY JOS VOSKUIL

Every few decades, a major change occurs that affects how we work inside our organization. When e-mail and the Internet became standard applications for employees, a different type of communication was born. Initially heralded by early adopters, these revolutionary new elements of work life soon matured and were considered safe and commoditized, spurring more companies to adopt these channels of communications.

When visionaries push ardently for change and there is a valid technology solution, entire industries are transformed. One recent game-changing technology, geared for the manufacturing industry, is product lifecycle management (PLM).

Today, a growing number of manufacturing companies are enhancing their PLM implementations and working on shared data and processes among their development and production organizations. As development and manufacturing increasingly becomes dispersed around the globe, it is even more necessary to implement a change that will bring more efficiency, quality, and speed to an organization.

DIFFERENT THINGS TO DIFFERENT PEOPLE

First of all, we need to realize that the term "bill of materials" (BOM) is perceived differently by different people. For most people in a company, the BOM is a list of materials needed to manufacture a product—this is often called the MBOM (manufacturing BOM). For engineers, usually a smaller group in the organization, the BOM is a list of components needed to define a product—this BOM is often referred to as an EBOM (engineering BOM).

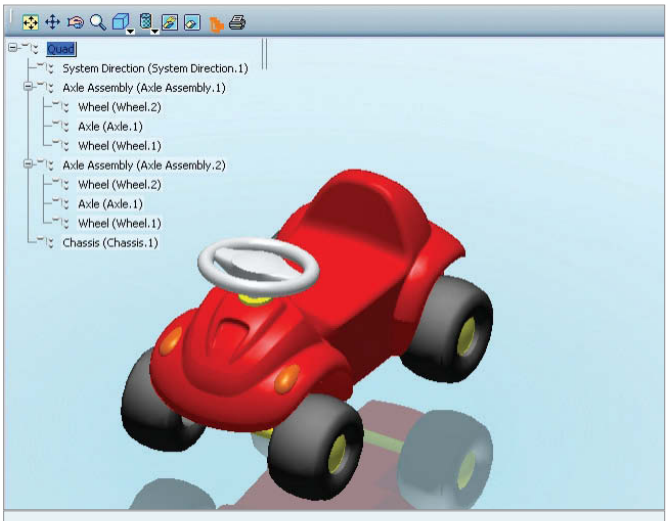

Developing a BOM directly from the solid model is automatic in a PLM system developed for small and medium-sized businesses and will save time in correcting introduced errors.

ERP systems were introduced years before PLM systems and were the first systems to manage a BOM—the MBOM, naturally. The MBOM defined the components needed to be available in order to manufacture and

deliver a product. This meant that manufacturing, drawings, assembly instructions, operations, and manuals were all, by some means, stored and referenced in this BOM.

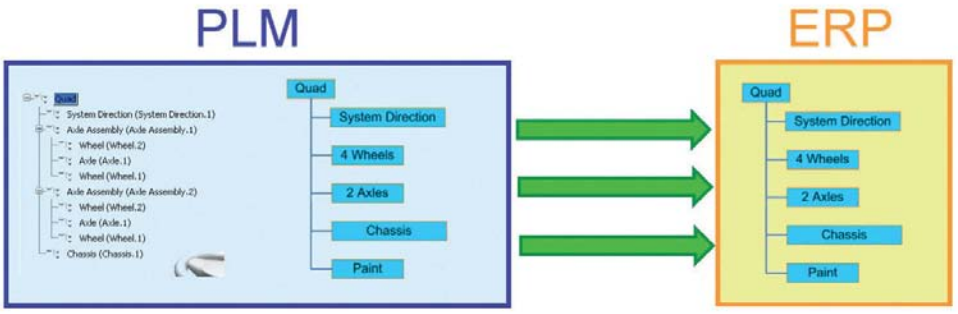

ENOVIA SmarTeam Engineering Express from Dassault Systèmes is an out-of-the-box PLM system with quick-start features that provides the infrastructure, the functionality, and the methodology to manage BOMs inside PLM and to automatically connect to ERP systems.

Before 3D design was popularized in solutions such as Dassault Systèmes' CATIA, the above situation was achieved through a somewhat unstructured process starting with engineering and ending in production. At first, a designer delivered the EBOM to a production engineer, in the format of paper or Excel spreadsheet. The next step was for the production engineer to complete or fix the data provided by the designer and define the BOM for manufacturing in the ERP system. This was often done by manually entering most of the information defined by the designer. Historically, this created an error-prone and labor-intensive situation, as people had to enter or search for the same data twice.

When 3D design was introduced, it became apparent that the structure built in the 3D CAD system resembled strongly or was even identical to the structure of the EBOM. Meanwhile, it also led to the need for product data management (PDM) systems, as all the relations between the products and parts had to be managed in to establish a correct version of the product. In the 2D world, there was no management system in place, as production engineers had to manually track the status of all the parts in an assembly.

Because 3D CAD offered improved quality and control, many companies developed the idea of connecting the 3D CAD system directly to the ERP system. The PDM system was used to store and extract the right version from the BOM.

ISSUES WITH ERP

Unfortunately, this approach remains plagued by several problems. First, since the EBOM and the BOM in ERP (the MBOM) are not identical, some components defined in the EBOM do not exist in the MBOM because they are not needed for manufacturing purposes.

For example, a subassembly of two parts is logically grouped together—the subassembly does not exist as a half-fabricated or production order in the ERP system. Another problem is that several MBOM items are not in the EBOM, as they are not part of the design. For example, paint or grease is needed to manufacture a product, though they are not required as part of the EBOM definition.

Sending the EBOM to ERP and manipulating the data manually does not really solve these problems nor does it bring many benefits. Although errors may be eliminated, the need to manually manipulate data in order to

get the right MBOM with all related information remains. In the ERP system there is still no visibility regarding the full product definition. What about the 3D models, the related drawings, and other documentation?

One of the most common tools for describing a BOM outside an ERP system is still the Excel spreadsheet. From the perspective of users, the Excel spreadsheet is easy to use and flexible, as users can write or define anything in this file. However, there is no control over the content or version of the Excel spreadsheet, and there is no control over changes to the spreadsheet. Again, this leads to a lot of manual overhead related to confirming, fixing, and adjusting data between engineering and production.

To avoid these problems, it makes sense to manage BOMs from design to production inside the PLM system. It's proven and lean. PLM systems have become the core system for all designers and engineers working on product data as they provided full visibility on the actual status of any BOM and its related information within a single system.

The bigger companies who have already implemented PLM have proven this method. Case studies have demonstrated that quality, time to implement, and time to react to changes have dramatically improved with PLM, involving fewer human resources overall, and eliminating the need for manual transformations and searching.

The challenge now lies for smaller companies to shift their mindset as well. Implementing PLM means letting go of Excel as the intermediate BOM format. It also means giving up the ERP system as the only place where actual production information is defined and referenced.

THE SMB CHANGE

Although mid-market companies cannot compare to big enterprises in terms of how they manage EBOM, MBOM, and other BOMs, they can choose a solution specifically built for their needs.

A targeted solution for mid-market companies, ENOVIA SmarTeam Engineering Express from Dassault Systèmes is an out-of-the-box PLM quick-start that provides the infrastructure, the functionality, and the methodology to manage BOMs inside PLM and to connect to ERP. It can also be used to learn and implement the real change for managing BOMs inside PLM and further expand on it with a proven ROI.

Like the change from the traditional post to e-mail, the PLM change may be fully understood in a few years. However, to gain a competitive advantage in early adoption, consider making the change now. If you were reading this article in the last century and wanted to consult a PLM expert, you might have had to buy a stamp. ■

***Jos Voskuil** is a management consultant for the ENOVIA SmarTeam division of Dassault Systèmes. You can send comments about this article to DE-Editors@deskeng.com.*

INFO

Dassault Systèmes
Paris, France
3ds.com

For more information on PLM, please visit deskeng.com.

Omnify Brings Savings and Clarity to Crystal Technology

BY ALAINE PORTNOY

> Manufacturer of oxide single crystals cuts engineering change cycles by 16 percent.

Since 1967, Crystal Technology, Inc. (CTI), an EPCOS company, has manufactured oxide single crystals and selected optical components based on these crystals. Located in Palo Alto, CA, CTI products are sold globally for a wide variety of applications in electronics, optics, and acoustics.

With product quality and value as top priorities, a staff of 150 people works to deliver a number of highly customized product lines involving various product development teams. As a volume manufacturer (the company produces 60 tons of single-crystalline lithium niobate annually), Crystal Technology operates 24 hours a day, seven days a week. This environment, coupled with increasing pressure from customers to deliver products faster, required Crystal to evaluate its existing product development processes and find ways to improve efficiencies.

"First and foremost, we needed to eliminate

Streamline Product Development and Change Management Processes

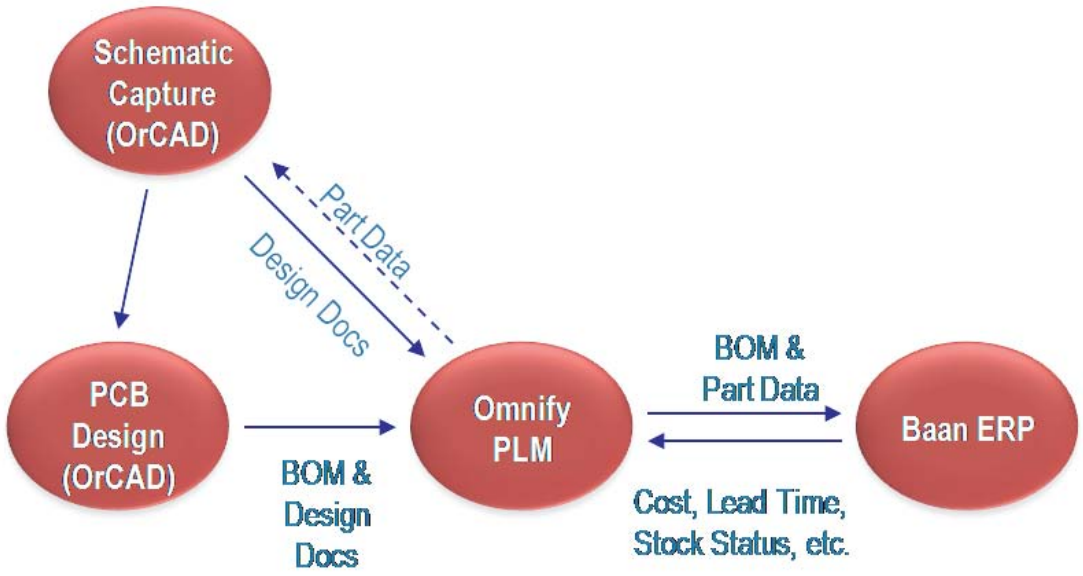

Engineering change orders (ECOs) are implemented and controlled within Omnify and automatically pushed to Infor ERP Baan. It enables catching any mistakes early on.

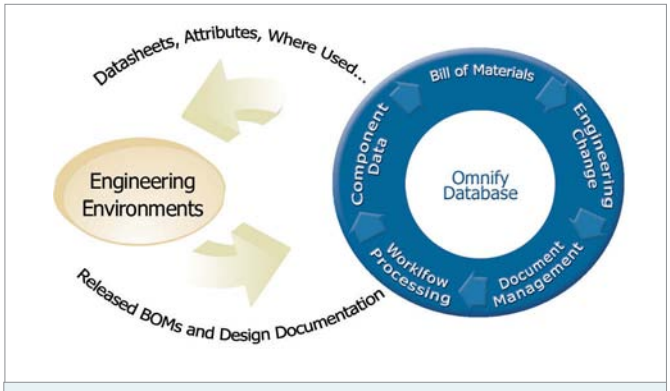

Empower is now the master database for all Crystal's parts and products . The system was easy to configure and required little customization.

waste associated with filing and searching for information," says Fred Garderes, director of supply chain management for Crystal. "We could no longer search through file cabinets to find information where one hour could be spent looking for one piece of information."

After decades of using manual, paper-based processes, Crystal needed to automate the exchange of internal and external information so it was available throughout the organization for faster reviews and evaluation. "We needed our data to be consistent and available across all of our systems," says Garderes.

CENTRALIZED PRODUCT DATA

Crystal wanted a central repository that would be the "bearer of the truth" for all data. With many product lines, and mostly customized products (each part is a custom application), it required a solution that could manage a large amount of part information such as specifications, package size, and specific output. Crystal needed to make sure new product designs could be easily reused to cut costs and eliminate re-work.

Crystal was looking to automate processes to streamline product development and change management between its Infor ERP Baan system and Cadence OrCAD design environment. "We were looking for continuity of information among these systems with some consideration for migrating information to our outsource partners, whether on the manufacturing or design side," says Garderes.

In addition, some of Crystal's electronic assembly design is outsourced and it wanted a system that would simplify the packaging of information sent to its contract manufacturing (CM) partners including bill of materials (BOMs), part information, assembly instructions, and test specifications.

OMNIFY THE SOLUTION

Crystal soon realized that a product lifecycle management (PLM) system was what it needed to manage all product information. According to Garderes, Baan did not have the features for managing detailed design information or the necessary level of access control.

Crystal evaluated three key PLM vendors, including Omnify Software. After assessing each product, the company picked Omnify's PLM solution because it addressed all Crystal's requirements at a reasonable cost, easily integrates with its existing design and manufacturing systems, and Om-

nify's team was willing to support company needs. And while Crystal engineers were comfortable with their paper processes and were hesitant about using a new automated system, Omnify's intuitive interface helped overcome that reluctance.

"Omnify Empower is now the master database for all of our parts and products from cradle to grave," says Garderes. "The system was easy to configure and required little customization to reach our key needs, including the interface to Baan. In order to maintain full integrity between Omnify and Baan, we decided to mirror a number of tables from Baan to support the parts creation process in Omnify. Today, 140 fields are passed between Baan and Omnify." All of Crystal's product documentation is solely vaulted in Omnify. "We plan to eliminate all of the file cabinets we have filled with paper and deposit that information in Omnify as well."

EFFICIENT DESIGN PROCESSES

Omnify helps to create more efficient design processes by providing engineers with instant access to information from their desktops via a direct interface to OrCAD. All of the detailed part, BOM, and document information stored and managed in Omnify is at the fingertips of Crystal's design engineers. "Omnify is used as a sandbox for our engineers to play with designs by adding items to BOMs, scratching off items from BOMs, creating parts, and changing drawings," says Garderes.

Crystal can now easily package required information from Omnify to send to its CMs and avoid redundancies. Engineering

change orders (ECOs) are implemented and controlled within Omnify and automatically pushed to Infor ERP Baan. It enables catching any mistakes early on. Crystal currently has 28 workflows in place for a specific product team where changes are reviewed and approved in Omnify and then automatically released from document control to Baan, eliminating any previous inconsistencies due to manual data entry.

Omnify Empower: Key Benefits

Centralized Product Data
- Streamline communication
- One version of the truth for all data
- Easy access to information
- Ensured accuracy of information across internal and external teams

Automated Processes
- Efficient design processes
- Eliminate manual entry of data/errors
- Direct data sharing between engineering and manufacturing
- Meet ISO and RoHS compliance requirements

Time and Cost Savings
- Reduce ECO cycle times
- Estimated overall savings of $250K over five years
- Support part re-use, reduce re-design

Business-ready Solution
- Easy to use/intuitive interface
- Fast user adoption
- Open technology platform for integration with existing engineering and ERP systems
- Cost-effective

"We wanted Approved Vendor Lists, costs, suppliers on the AVL, and potential switches from supplier to supplier available to engineers to better meet customer requirements and save on costs," says Garderes. "Having all of this data available in one view in Omnify lets engineers know the inventory on hand and on order, who the supplier is, the availability, etc., so they can make better design choices."

As an ISO 9001:2000 and ISO 140001 registered company, Crystal is required to document certain procedures and maintain adequate records to conform to both regulations. By providing an automated process for documentation, Omnify helps Crystal pass audits and maintain accreditation. Omnify is also used to meet Restriction of Hazardous Substances (RoHS) requirements for identifying, tracking, and reporting compliance information.

CUSTOMER SUCCESS/ RETURN ON INVESTMENT

Implementing the Omnify PLM solution has helped Crystal meet its goal to improve product development processes and effectively meet customer demands while maintaining its focus on quality and cost leadership. "We are very satisfied with the results of our Omnify implementation," says Garderes. "We have been able to reduce our ECO cycle by 16 percent, from 12 to 10 days, and this is trending towards even further reduction."

Crystal is also realizing significant time savings in finding information, making it easy to

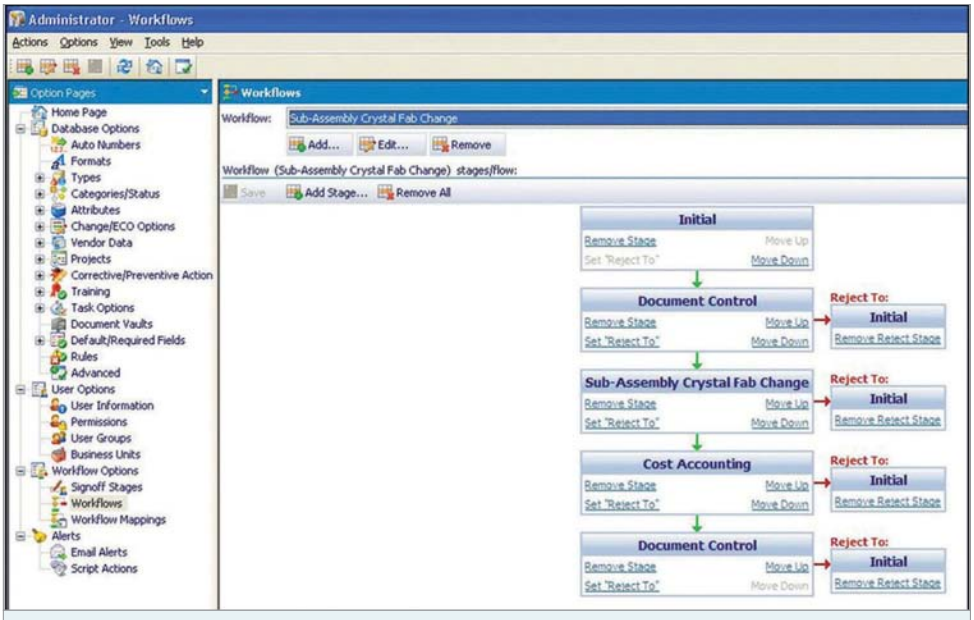

Crystal currently has 28 workflows in place for a specific product team where changes are reviewed and approved in Omnify and then automatically released from document control to Baan, eliminating any previous inconsistencies due to manual data entry.

retrieve historical as well as current information. No more diving into file cabinets.

Finally, the company is getting a return rate of approximately 2.5 times its investment over the first five years, with a projected ROI rate of 6.25 thereafter. "We have estimated a substantial cost savings of about $250,000 over a five-year period due to the enhancements made to our product development processes."

As Crystal expands its use of Omnify Empower with new Quality Management and Project Management modules, it will eliminate all paper processes and home-grown systems. Further, Crystal will have all document control and quality information in one central location, giving engineers added knowledge about what issues accompanied a part during its lifetime, and it will leverage the Project Management module as a platform for its new product introduction (NPI) process to help shrink those NPI cycles. ■

Alaine Portnoy *is the vice president of marketing at Omnify Software. She has experience in the manufacturing and PLM industry in the mechanical, defense, medical device, and electronics markets. Send comments about this article to DE-Editors@deskeng.com.*

INFO

Cadence Design Systems
San Jose, CA
cadence.com

Crystal Technology, Inc.
Palo Alto, CA
crystaltechnology.com

Infor Global Solutions
Alpharetta, GA
infor.com

Omnify Software
Andover, MA
omnifysoft.com

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

Tecplot: Relieving the Pressure in Calculating Flow Separation

> Researchers relieve the pressure in calculating flow separation using Tecplot 360.

BY VINCE ADAMS

In any simulation, obtaining accurate results is only half the battle. Whether the domain is material stress or fluid flow, any data is only as good as the engineer's ability to interpret and apply it.

One of the more important yet difficult calculations to understand is the precise behavior of the breakaway of fluid from walls. However, after five years of research, a team of scientists from MIT, San Diego State University, and United Technologies has developed a breakthrough methodology that allows engineers to precisely determine the location and angle of flow separation in both 2D and 3D as well as steady and unsteady flow conditions.

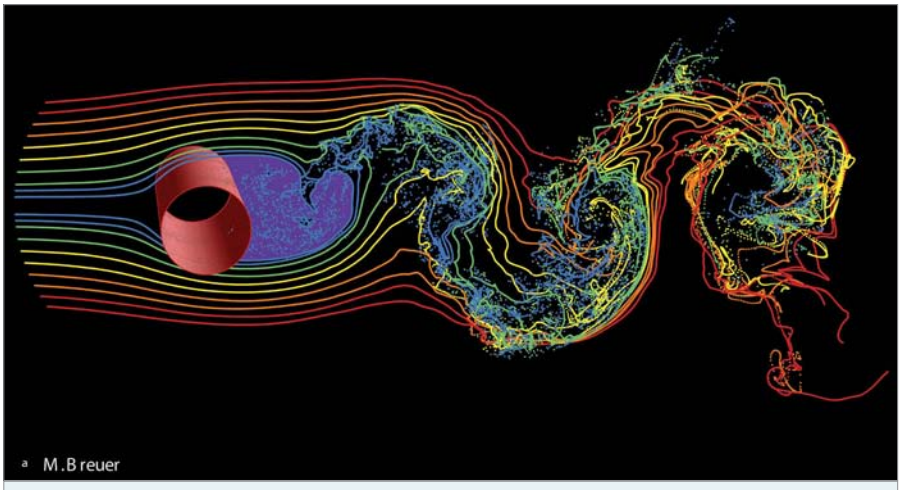

Figure 1: CFD Results of Flow Separation Behind a Cylinder. *Image Courtesy of Tecplot Inc. and M. Breuer, Department of Fluid Mechanics, Helmut-Schmidt University Hamburg.*

"Visualization of steady flow vs. unsteady flow can be compared to looking at a static image vs. a movie," notes Mike Peery, president of TecPlot, Inc. in Bellevue, WA. "The work they did isolates the steady parameters from the craziness of unsteady flow."

The ramifications of this development are far reaching, from minimizing drag in automobiles to improving cooling of turbine blades (*see "Turbine Benefit"*), as industries strive for greater energy efficiency and fuel economy.

THE PROBLEM

Engineers and scientists have been aware of flow separation for more than 100 years. It started with Ludwig Prandtl's identification of the phenomenon. Prior to Prandtl's recognition of the boundary layer in fluid flow around an immersed body, scientists were unclear about the impact of viscosity and fric-

tion at the solid surface. Most flow analysis neglected this phenomenon out of necessity, thus adding a large degree of uncertainty to the calculations that would lead to predicting drag. Figure 2 shows the flow about a cylinder under these assumptions.

Note that the flow stream, including velocity, is symmetric about the center of the cylinder. Bernoulli's Principle suggests that with equal velocities, the pressure distribution at the front of the cylinder must be equal and opposite to the pressure at the rear of the cylinder, thus canceling each other out in the flow direction. Theoretically, there should be no drag on the cylinder in an inviscid flow! Jean le Rond d'Alembert identified the discrepancy between this and experimental results in the mid 18th century but failed to explain it.

Prandtl's theoretical work on boundary layers, published in 1904, clarified that flow around any solid body consisted of two distinct regions, the boundary layer and free flow. At the surface of the solid, velocity is zero and increases over a small distance to the velocity of the free flow. This transition region is what we now know as the boundary layer. Since the boundary layer is thin, the velocity gradient is high. Newton's shear stress law then predicts that a large amount of shear stress must be present, creating what is now known as skin friction drag force. This drag force alone does not

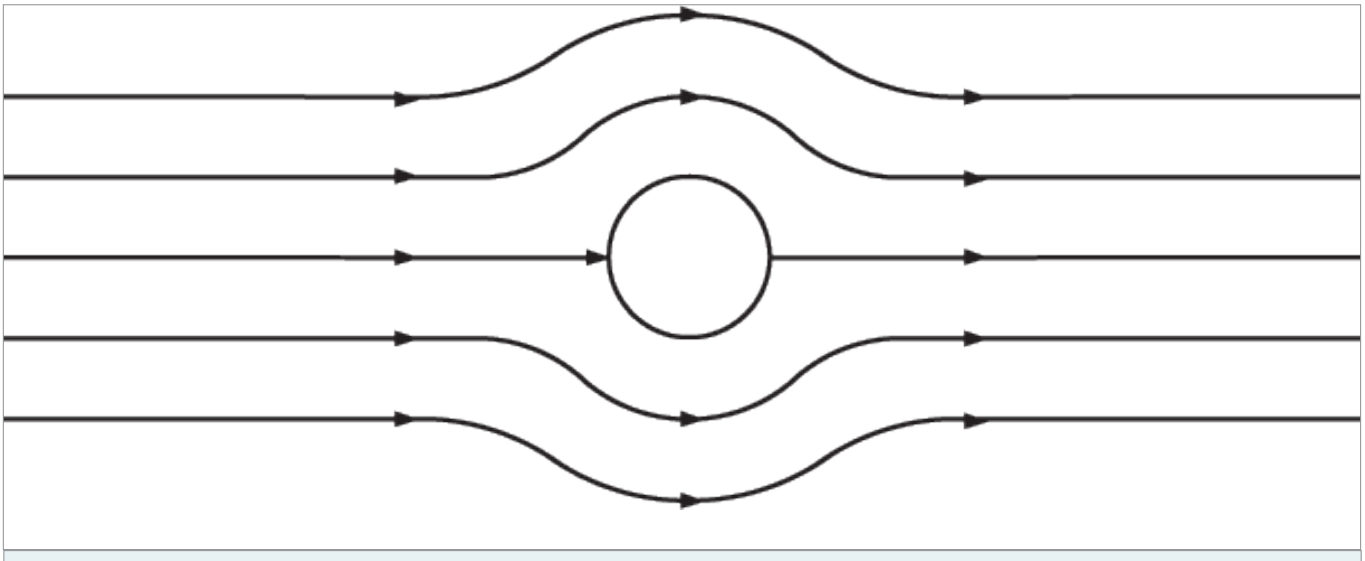

Figure 2: Theoretical Flow Around a Cylinder.

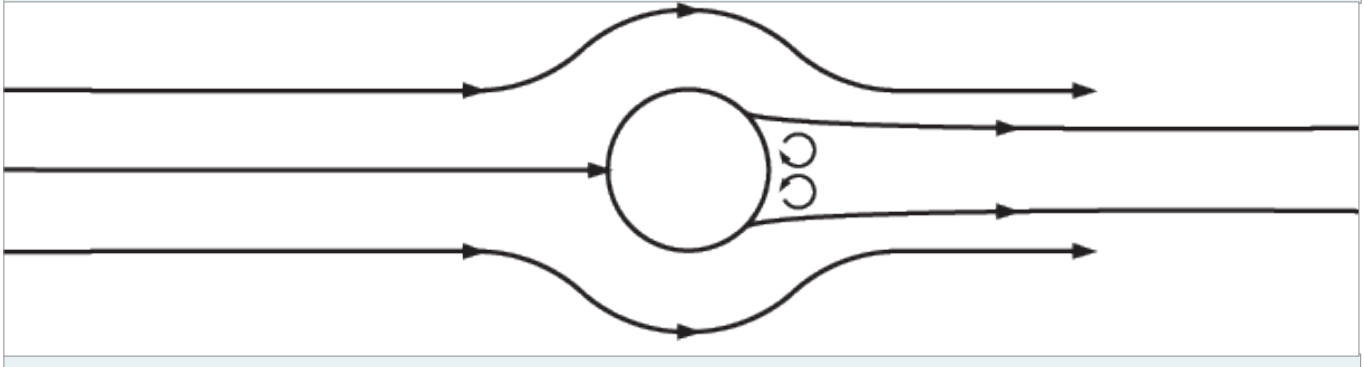

Figure 3: Actual Flow Around a Cylinder with Separation.

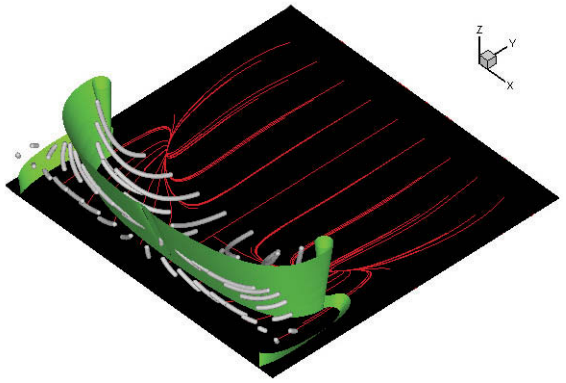

Figure 4: It is now possible to calculate with reasonable precision the exact location of flow separation and to visualize the geometry of its introduction into free flow as a 3D isosurface.

compensate for the discrepancy identified by d'Alembert.

Prandtl also identified that flow separation was related to the boundary layer behavior. The large shear stresses in the boundary layer, in the presence of increasing pressure, cause the boundary layer to separate from the surface and trail off into the free flow *(see Figure 3*). When the flow separates from the body, the fragile assumption of balanced pressure is violated. The large pressure region on the leading face of a body in a flow stream is no longer canceled out by the negative pressure at the trailing face and form drag is introduced.

Form drag is usually much larger than the skin friction drag and increases with increasing fluid speed. Consequently, minimizing the pressure differential by reducing or eliminating the area on the trailing face beyond the point of boundary layer separation can have a dramatic effect on the efficiency of a body that travels through a fluid. Applications critically dependent upon this include airplanes, automobiles, boats, and even golf balls. (Hence the dimples!)

THE SOLUTION

Aerodynamicists, or any engineer looking to improve the drag on a product, traditionally relied on a wind tunnel to examine the interaction between the body and the fluid. The location at which the flow separates from the body is typically characterized by a transition from laminar flow to turbulent flow. A smoke stream will highlight this transition with swirling or vortex shedding. While the presence of vortex shedding is readily apparent to even the untrained eye, quantifying the location on the body where the flow actually separates is extremely difficult, even for experts.

In conditions of steady flow, where the input stream is time-independent or cyclic, shear stress on the body's surface can indicate the location of separation as well. Special paints on a body in a wind tunnel can identify changes in shear stress, thus indicating flow separation but they provide no feedback on where the fluid goes when it separates. This is also important for the development of cooling and

mixing applications and doesn't address the more general case of unsteady flow.

For about 20 years, practically speaking, engineers have been able to use computers to predict fluid flow using computational fluid dynamics (CFD). This provides many advantages over wind tunnel testing in that variables can be controlled and results can be visualized in a number of graphical means. Moreover, these results can be animated, saved, replayed, and scrutinized at an engineer's leisure. This flexibility in visualization makes the determination of separation location easier but, as can be seen in Figure 1, which replicates the smoke stream results, it is still imprecise.

"You can see those vortices forming and flipping around and it's all very pretty but how do you get engineering understanding out of that?" asks Peery.

Turbine Benefit

Gas turbine engines actually run at temperatures well above the melting temperature of the blades for increased efficiency. Engineers use a technique called film cooling so that the blades can operate in this environment. This requires small holes in the blade where compressed air creates a cool film near the surface that keeps it from melting. The flow separation of the boundary layer caused by these jets is one of the limiting factors in how hot a gas turbine can run. If the engineers could truly understand the flow separation in the region of these small jets, the location and pattern of the holes on the blades or the cooling airflow rates can be optimized, greatly improving the efficiency of the engines.

THE BREAKTHROUGH

It is now possible to calculate with reasonable precision the exact location of flow separation in certain conditions and to visualize the geometry of its introduction into free flow as a 3D isosurface. Professors George Haller (MIT), Gustaaf Jacobs (San Diego State), Dr. Amit Surana (United Technologies) and their team were able to illustrate this phenomenon with a 3D isosurface to represent the actual computed separation result with the help of Tecplot 360 (*see Figure 4*).

"This is a real change in the way we look at fluids and will affect many areas of flow engineering," says Jacobs. "While drag is certainly an important application, the improvement of mixing processes may turn out to be as important due to the transport mechanism being modeled."

According to Peery, "Tecplot 360 helps engineers interpret and understand the computed unsteady details in ways that actually are totally profound, different than what anyone has thought before, at least in common knowledge. With the growing computing power available to engineers, the calculation of detailed unsteady results is more commonplace. This creates a need to interpret that data in a way that makes sense."

Craig Jensen, an industry analyst with ASI USA, believes this helps to address a growing problem in the automotive industry: "The industry needs to move more aerodynamic study into the realm of simulation but lengthy solution times and the inability of engineers to communicate trends has encouraged continued reliance on physical testing. This visualization tool can certainly help to break down that barrier." Larry Rinek, a business consultant

with Frost and Sullivan and a published aviation historian, concurs. "It is important to keep flow attached as far back as possible. A few inches may provide a measurable improvement in fuel economy and CO2 emissions." Larry also points out that flow separation is a common cause of instability in aircraft.

The fact that these calculations support both steady and unsteady flows should not be overlooked. Prior theory only supported the estimation of flow separation in steady, or predictable and repeating, flow conditions. The determination of a mean isosurface representing flow separation in an unsteady flow condition is unparalleled. Jacobs is quick to point out that the work that his team recently completed is built on significant development of many others in the fluid dynamics field. "Our work rests on the shoulders of hundreds of scientists who have been working on this problem for more than a century," says Jacobs.

Jacobs also credited the power and flexibility of Tecplot 360 in the development of this breakthrough. He has been using Tecplot software for more than 10 years and believes that its greatest strength is its user friendliness and customization. "The layout files and macro features are particularly useful." This breakthrough required a visualization component that didn't exist in commercial software to date. The ability to work within an existing postprocessing, or visualization, framework freed his team to focus on the science, not the user interface.

Peery is quick to note that while he and his company appreciate the opportunity to provide visualization technologies for this study, the focus must lie with Haller, Jacobs, Surana, and the research. "Our understanding of flow separation has just taken a big leap forward."

The methodology uses high-order Navier-Stokes equations, the backbone of all modern CFD calculations. The solution expressly addresses behavior in the boundary layer.

According to Jacobs, "The lower-order methods employed by most commercial CFD codes are efficient for general flow results but may have trouble getting accurate results in the boundary layer." More precision was required to fully understand separation. Jacobs suggests it may be a few years before commercial applications incorporate this research. However, demand from the engineering community will ultimately drive leading CFD providers to this and other ways to improve understanding of increasingly complex data. ■

INFO

ASI USA
Bingham Farms, MI
asiusa.com

Frost & Sullivan
frost.com

Tecplot, Inc.
Bellevue, WA
tecplot.com

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

Vince Adams *has been working, writing, and teaching in the field of simulation for 20 years. He is the co-author of 3 books and numerous articles on FEA and design analysis. He can be reached for questions or comments at vince_adams@sbcglobal.net. Comments should be sent to DE-Editors@deskeng.com.*

Getting More from Less with Fluent and EnSight

> Missouri Science and Technology racing students use CFD and visualization to gain a leg up on the competition.

BY KARA L. GRAY

What happens when you marry high-performance formula-style racing with high performance computing (HPC)? A group of students from Missouri University of Science and Technology (Missouri S&T) have discovered that this winning combination yields both powerful results and outstanding opportunities.

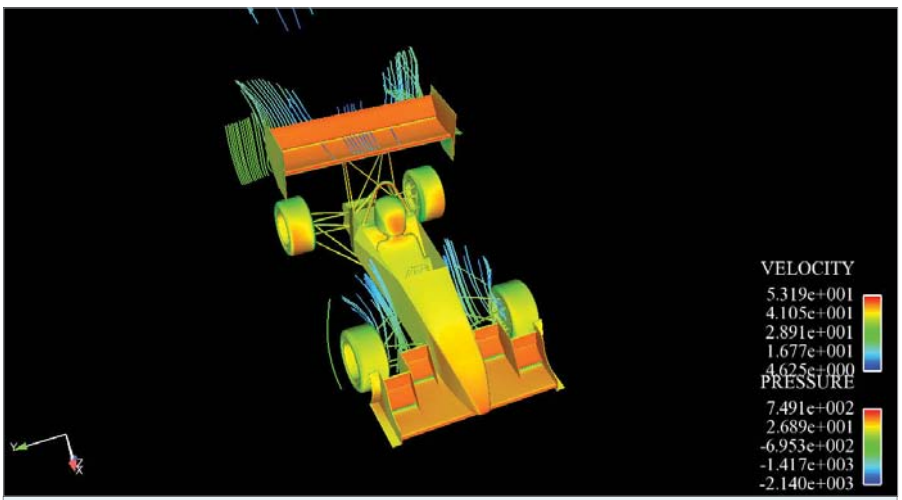

Missouri S&T students simulated airflow around the #92 car to gain optimum downforce in specific areas without inflicting detrimental effects on other parts of the car. This image is an isometric view showing static pressure contours on the surface with pathlines colored by velocity.

For the past 21 years, Missouri S&T's Formula SAE racing team, part of the university's Student Design and Experiential Learning Center (SDELC), has set its sights on a first-place finish in the Formula Society of Automotive Engineers (SAE) collegiate competition that challenges engineering students to build and compete with a formula-style race car against teams from around the world. The event is designed to simulate a real-life scenario where a "corporation" commissions a group of engineers to design and build a prototype car. Annual competitions are held worldwide, including U.S. races at Michigan International Speedway, Virginia International Raceway, and California Speedway.

Each year, more than 300 teams take part in the two-phase competition that includes a series of static events, including cost analysis, sales presentation, and engineering design, as well as dynamic events like acceleration, skid pad testing, and both autocross and endurance races. Static events

are judged by marketing and product development specialists at the Big 3 automotive companies: Ford, General Motors, and Daimler-Chrysler.

For students like senior Andrew D'Hooge, a double major in mechanical and aerospace engineering, being part of the Missouri S&T racing team is more than an extracurricular activity. It's an integral part of the curriculum, an opportunity to reinforce classroom lessons with hands-on design and build tasks in an environment that fosters teamwork, leadership, and communication—essential skills in the engineering field.

AERODYNAMICS: A NOVEL APPROACH

As the aerodynamics group leader, D'Hooge is responsible for designing and testing all aero components on the Missouri S&T #92 car, with an eye toward the ultimate goal of achieving maximum downforce with minimal mass.

"We're one of only a handful of teams that consistently run an aerodynamics package," D'Hooge says. In a sport where the average course speed is around 35 mph and speeds in excess of 60 mph are a rarity, reducing drag isn't as much of a concern as is optimizing downforce. "Since there aren't a lot of straight-aways in Formula SAE, cornering is a major component. The more downforce you can generate, the better handling and lateral maneuvering capability you can achieve."

But generating this valuable downforce must be accomplished with precision and without added weight. Generating downforce is one thing, but generating it where it can do the most good is vital—too much in one area of the car can impact maneuverability. And, the weight of added components can overwhelm the aero benefits.

ACHIEVING OPTIMUM DESIGN THROUGH CFD

To achieve the optimum design and analyze the aero impact on the overall performance of the car, D'Hooge and his teammates are using cutting-edge computational fluid dynamics (CFD) and extreme visualization software to perfect the design conceptually, before any manufacturing ever takes place. Because the group must essentially start from scratch to build a new car each year, the CFD approach significantly reduces build time and eliminates physical trial and error.

"The goal with implementing CFD in the design process is to reduce manufacturing time to allow for more testing time," D'Hooge

INFO

ANSYS, Inc./ Fluent
Canonsburg, PA
ansys.com

CEI, Inc.
Apex, NC
ensight.com

Siemens PLM Software
Plano, TX
siemens.com/plm

Missouri S&T Formula SAE Racing
fsae.mst.edu

Formula SAE Collegiate Design Series
students.sae.org/competitions/formulaseries/

For more information on visualization and CFD, please visit deskeng.com

says. “Once we’ve arrived at a final design with the CFD, we can then validate the results in a wind tunnel. Ensuring reliability of the components is equally important. No matter how they impact the performance of the car, we have to be sure they will stand up to the rigors of the competition.”

D’Hooge and his team begin the process the same as any other team, with concept generation, brainstorming, and discussion of implementation plans. Once the concept for an individual component is established, like the front wing for example, a CAD image is rendered using Siemens PLM Software’s NX 5 CAD/CAM/CAE software. From this image, team members then turn to Fluent’s GAMBIT 2.2 to generate the million-plus component finite mesh. The mesh is then imported into Fluent 6.3 for flow modeling.

EXTREME VISUALIZATION/ QUICK RESULTS

To visualize the model, the team uses EnSight extreme visualization software by CEI, Inc. of Apex, NC. EnSight allows the team to quickly and accurately preview the results of possible design iterations before any changes are made on the shop floor.

“We had been doing our postprocessing work in Fluent, but EnSight has proven to be far easier to use and offers far more expanded capabilities,” D’Hooge says. “Thanks to the CFD and the ability to visualize the components before we make them, we can now turn around parts in a day, from the initial concept through drawing, meshing, and postprocessing.”

The ability to visualize the results of design changes is a critical time-saver for the Missouri S&T team. With no on-site access to a full-scale wind tunnel to test their components, the team must be well-prepared for its annual Spring Break visit to the Ford Motor Company wind tunnel in Detroit. This, the final step in the build process, is an opportunity for the team to validate its results by correlating CFD data with real data gleaned from the wind tunnel tests.

“Once we’ve validated the CFD results in the wind tunnel, we can then use this data as a broader design tool,” D’Hooge says. Even though at this point in the process there is no time left to redesign components before the first competition, the data still has incredible value for the team. While they do have to start the build from scratch next year, they can use the same design elements and CFD models in an evolutionary design strategy, again saving the team precious time in development.

With EnSight on their team, the Missouri S&T students have definitely found a competitive advantage both on and off the track. EnSight’s flexibility and ease of use have made design alterations quick and effective.

“EnSight’s interface is really easy to use and it’s very intuitive,” D’Hooge says. “We can jump right into using it without having to sit down and read a manual. It gives us the ability to iterate individual components through numerous designs, and even incorporate these analyses into a full-car CFD, to see how each component affects the whole car before we even start to build.”

SELLING THE SCIENCE

In addition to providing the pre-build analysis that is critical to the car’s performance in the dynamic events, the images and anima-

tions generated in EnSight are used as part of the marketing presentation portion of the competition and for promotional materials to gather sponsors for the team.

Performance benefits aside, D'Hooge says one of the greatest advantages of using cutting-edge tools like EnSight is the opportunity to gain valuable experience with industry-standard software and design procedures, just like it's done in the real world.

"A lot of former members have had some excellent job opportunities as a result of being part of the team, with OEMs, the Big 3, and we even have several alumni in the motorsports industry," he says. "It's a valuable experience to have the opportunity to learn how to use these tools, but the team aspect is also a major benefit. These are skills you can't teach in a classroom. But here we get hands-on experience using the same software and manufacturing processes as the Formula One teams are using."

The 2009 Missouri S&T team is looking to build upon its previous three top 10 finishes in the past six seasons and is set to compete in three domestic competitions this season at the Virginia International Raceway, at Michigan International Speedway, and at California Speedway. ■

Kara L. Gray is a professional writing and public relations consultant specializing in technology, manufacturing, and non-profit communications. You can send an e-mail about this article to DE-Editors@deskeng.com.

Leveraging Abaqus CAE During a Downturn

> CAE and analysis-based design represent a technology to help you weather the economy.

BY FRANK POPIELAS

Engineering companies, like everyone else affected by the current economy, are looking to control expenses. And many seeking to cut resources would be remiss in cutting computer-aided engineering (CAE) resources, and especially CAE engineers who can be seen as the keys to business survival today. In fact, CAE engineers can be considered the best insurance for emerging competitively when the climate improves.

Analysis-led design will keep you focused on what's really important now: producing quality products cost-effectively. A well-developed CAE system lets you design your product properly upfront and test it virtually before you start prototyping. This helps manage expenses without jeopardizing the business.

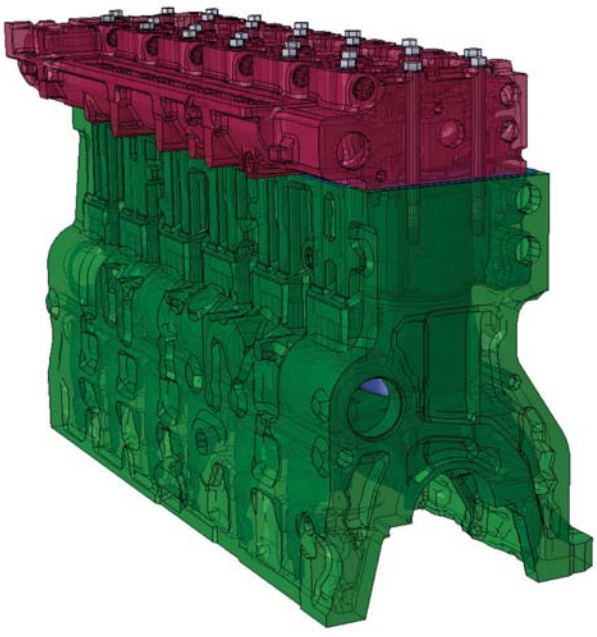

Abaqus FEA image of a powertrain model as used by Dana Holding Corporation. The CAE department of the engineering company performed two HPC processing-time benchmarks using similar models and found that software upgrades dramatically reduced the number of cores needed for the analysis.

FIND YOUR "SWEET SPOT" FOR CAE SYNERGY

In a down economy it is essential to identify the "sweet spot" for your HPC (high-performance computing) system as it runs today's large, complex models. HPC, used in tandem with the right software, will give you the proper scaling and performance your CAE system demands. The combination of these two elements defines your sweet spot, that synergy of hardware, software, and engineering costs that gives you the lowest expense-per-unit for your product.

A critical way to maximize existing HPC investments is to upgrade with the latest software releases that make best use of hardware per-

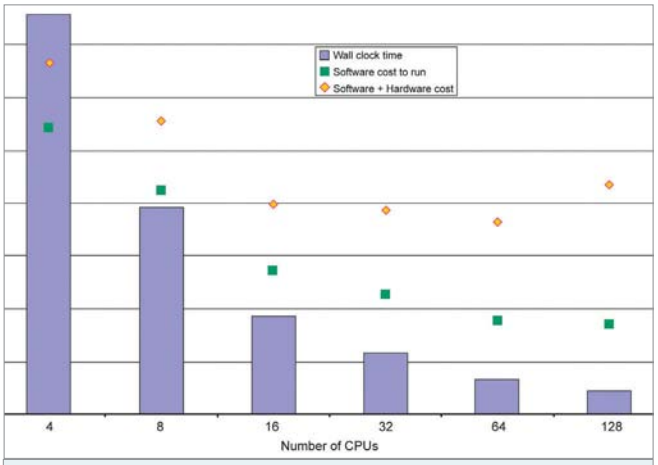

This graph shows wall clock time required for FEA analysis (in this case Abaqus V6.7) of a powertrain model and the cost of software and hardware needed to run the job. Note that with this earlier version of the software the cost minimum (known as the "sweet spot") did not occur until 64 cores were used.

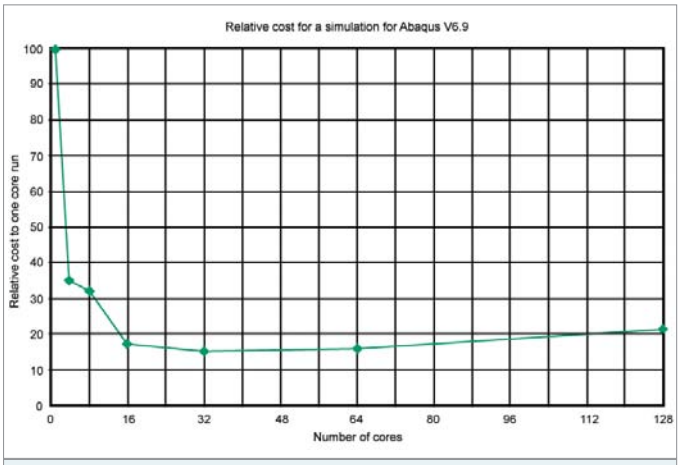

This graph shows a similar cost analysis to that in the graph at the left, but this time the job was run using Abaqus V6.9. The sweet spot shifted to 32 cores, instead of 64 as seen in the image at left (benchmark study performed two years earlier). This dramatic improvement is largely a result of software upgrades.

formance, minimizing the cost per unit of both hardware and software. Recent developments in software are boosting performance even further, to a point where hardware costs are only a minor factor in simulation costs overall.

Two years ago Dana Holding Corporation of Toledo, OH, performed a software/hardware benchmark study using Abaqus FEA from SIMULIA on a basic powertrain model. At that time, we identified a sweet spot of around 64 cores. Earlier this year we completed beta testing of the latest software release from Abaqus (V6.9) and found that on essentially the same model, we only needed about 32 cores—half what we used two years ago—to hit the sweet spot for a similar FEA analysis. HPC provided some degree of better handling on the hardware side, but the main reason for the cost reduction was improvements made to the software.

Previously, once we passed our CAE sweet spot, costs would increase again. But with the enhanced software scalability of the latest Abaqus release, combined with improved hardware management in place, we find that costs remain almost flat after that. With this system we can use additional cores to speed up our simulation when we are under the gun to meet a deadline, without increasing our expenses.

Boosting CAE resources in this way is certainly easier for large OEMs than for smaller companies. But it's not enough just to have the resources: you need to be well organized to use them efficiently. At Dana we make sure that whenever there is a new software release, our global CAE team upgrades. We began investing in CAE in a more focused way a decade ago and are now at the point where our hardware

and software work in close tandem, with the right system in place to enable all team members to easily communicate with each other.

HOW LEVERAGING CAE SAVED A BUSINESS

An example of how this kind of long-term CAE development strategy pays off can be seen in a large customer of ours. In the past they wasted time and money getting a product to market that was no longer profitable or competitive in their field.

In order to survive, they had to completely rethink the way they did business—and they did it by leveraging the power of CAE. They could have gone the traditional way of slashing costs by just cutting staff, and they did reduce their headcount to some extent. But they also reexamined the way they allocated their existing resources, changing their mindset from an outdated process flow of build-test-error, build-test-error, to a new focus on upfront design. They began investing in CAE technology, using analysis-led design to optimize their product on the computer before cutting prototypes. This helped them keep costs under control while they built up a new technology foundation.

CAE was the key to significantly reducing operating expenses at this company and, more importantly, to long-term success as well as building better-quality products. After several years, the company had doubled its market share.

INFO

Dana Holding Corporation
Toledo, OH
dana.com

SIMULIA
Providence, RI
simulia.com

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

INVESTING IN CAE PROMOTES RECOVERY

How can you start investing in CAE during an economic downturn? You need to take a step back and look at the big picture: Even if you have just a small CAE team in place, you should support them now because they are your starting point for the future. But first make sure you've established the basic simulation techniques and standards that fit your product. Only then should you look for an HPC system that will support those needs.

While you may not see immediate results in the current downturn, any delay will only expose your business further. Companies that have established their focus on CAE earlier will lead the recovery in their respective areas and benefit from their focus on analysis-led design.

Hardware and software costs are coming down. What is your company's sweet spot for a product? What will be the paybacks for your investment? While it's hard to put an exact figure on returns because they are dependent on how, and how well, you use your CAE capabilities, a return on investment within two years seems very likely. So hang on to your CAE team members, support them at the level you can, and stay close to the leading edge, so you'll be ready when the economy picks up again. ■

Frank Popeilas *is manager of advanced engineering for the Sealing Products Group at Dana Holding Corporation. He has an M.S. in engineering and has been with Dana for 15 years. He is responsible for development of advanced materials, testing, and simulation techniques. Send comments about this article to DE-Editors@deskeng.com.*

Adobe Photoshop CS4 Extended Improves Flow

BY MARK CLARKSON

> The latest version is a strong upgrade with plenty of new features for transforming your rendering and navigating your scene.

In the nearly two decades since its release, Adobe's Photoshop has become the definitive image-editing program—an application so ubiquitous, it's become a verb. Does Adobe have anything left to add? I found out.

A NEW LOOK

There are significant changes to the Photoshop UI, but nothing disorienting. The improved Screen Mode button has been moved to the top of the interface, for example, next to the new Rotate View control and

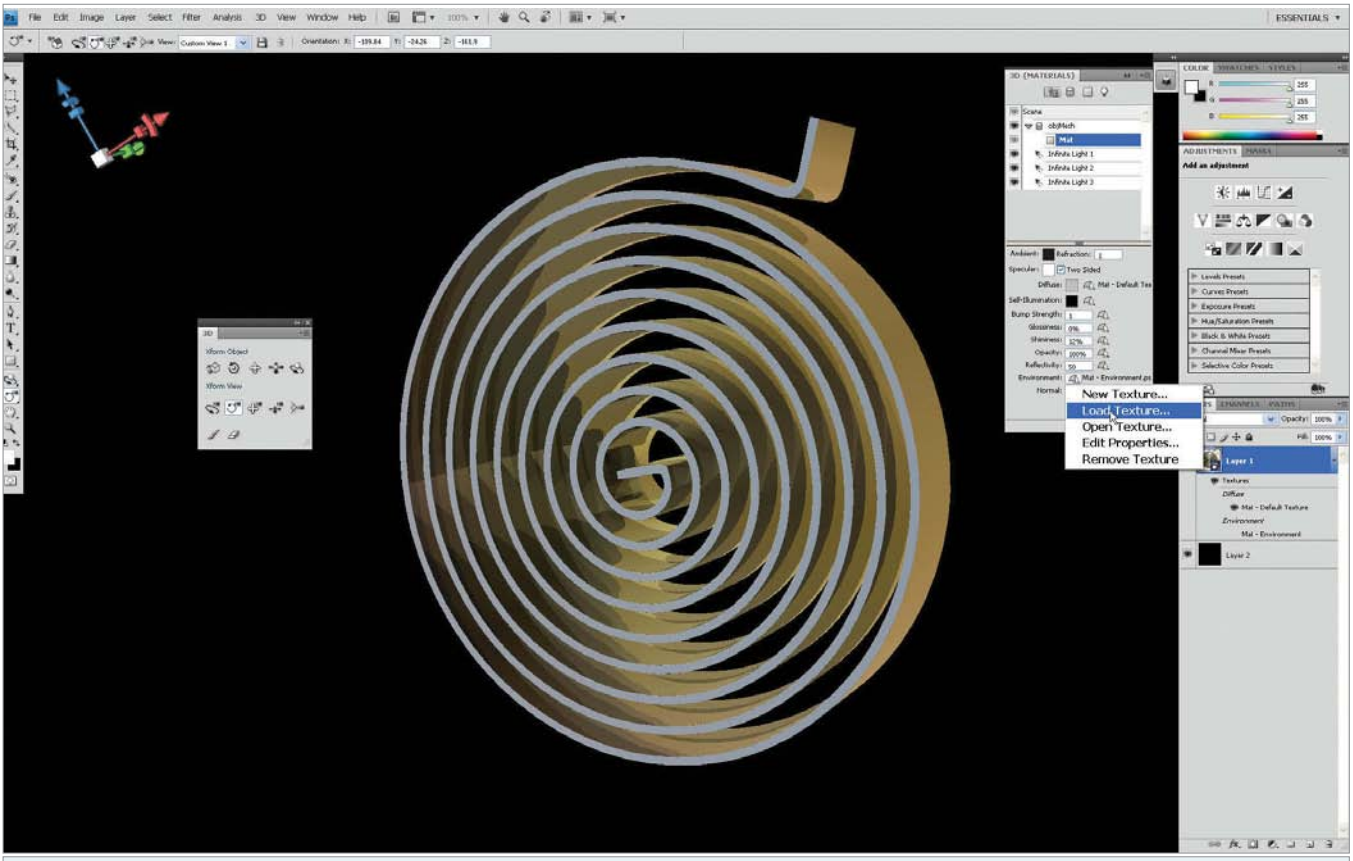

Photoshop CS4 provides transparency, volume, and ray-traced reflections as well as shadows on any 3D object you can design.

a much-needed Arrange Documents tool.

Layer masks now have interactive, non-destructive controls for density and feathering. And adjustment layers are easier than ever to use (and if you don't use adjustment layers, I'm frowning at you right now!)

Photoshop CS4 sports de rigueur tabbed documents that greatly facilitate flipping through a dozen photos. You can tear a tab off and float the document free from the rest of the Photoshop interface, even moving it to a different monitor.

OPENGL ACCELERATION

In a trend we're seeing everywhere, Photoshop CS4 is making use of GPU hardware acceleration, although it's not offloading computational tasks such as filters just yet. For now, CS4's OpenGL acceleration is largely limited to 3D (about which there will be more later) and moving around in large images—specifically zooming and panning, both of which have been subtly yet radically reformed.

With GPU acceleration, zooming is continuous. Click with the Zoom tool and you can zoom smoothly in or out, instead of jumping between discrete zoom levels as was the case.

You can "flick pan" documents: give them a little shove and they continue sliding smoothly across the screen before slowing to a stop.

Best of all is the new Birds-eye View tool, which allows you to flip back and forth between zoomed-in and zoomed-out views and easily navigate around the largest images. Like many past Photoshop features, these seem gratuitous until you've worked with them for a while; then you find you just can't live without them.

CONTENT-AWARE SCALING

One of the coolest new technologies in

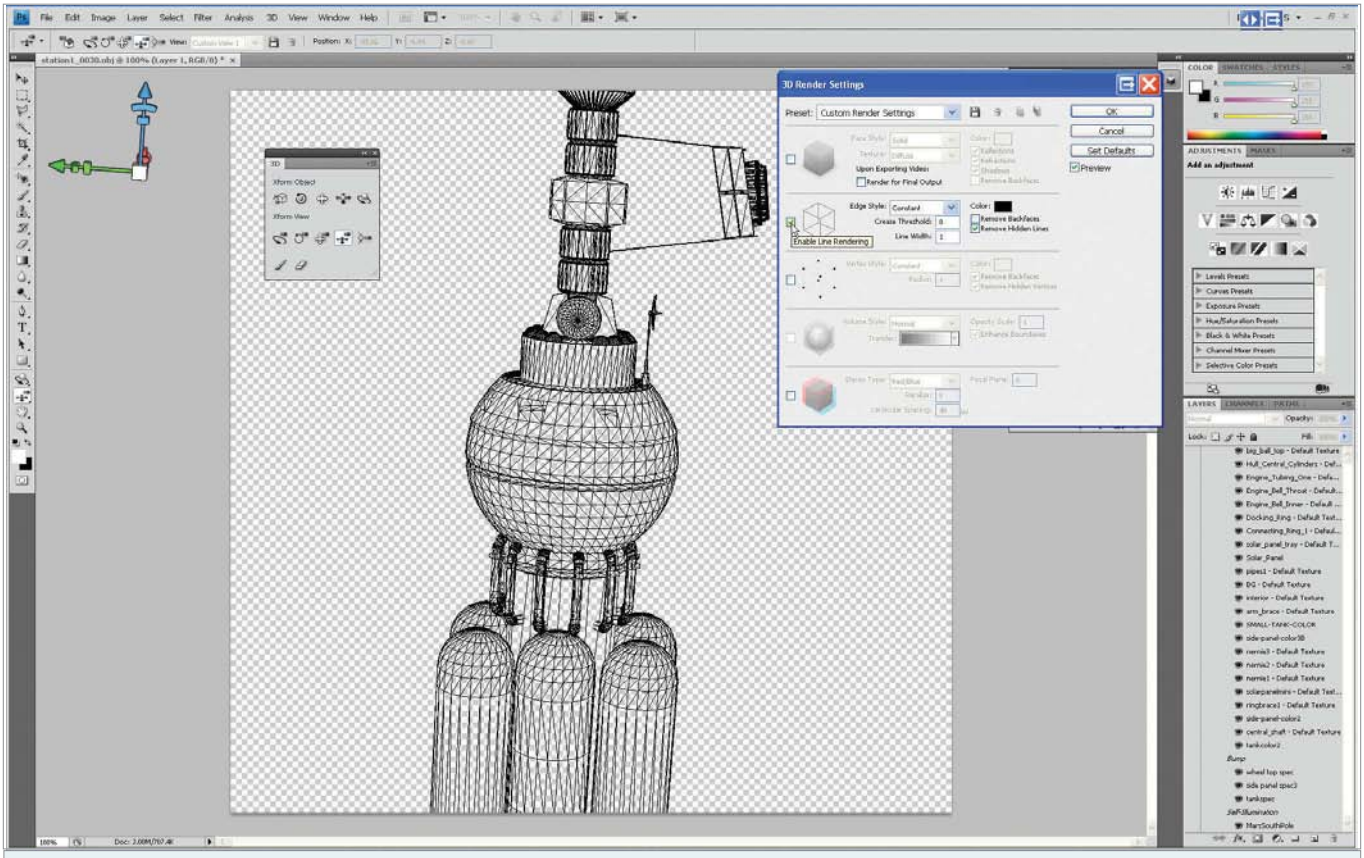

Photoshop CS4 imports 3D objects, complete with textures. Objects can be displayed as solids, lines, or points.

CS4 is content-aware scaling, which intelligently scales your images, stretching and squashing less important pixels (say, trees and sky in the background) while preserving more important pixels, such as a group of people in the foreground. You also can manually protect certain areas of the image to help the process out.

3D, NOW MORE THAN EVER

CS4, like CS3 before it, allows you to import 3D objects (3DS, OBJ, U3D, KMZ, and COLLADA) as well as 2D images. 3D objects appear with their full complement of maps, editable just like any other Photoshop documents. With CS4, you can now paint directly on models, creating and modifying texture maps with Photoshop's standard tools: the brush, pencil, clone stamp, and so forth.

Photoshop isn't the best place to work on your 100-million-poly airplane, but for small and medium models it works pretty well. I've pulled in OBJ files in excess of a million polygons and, while things got a bit chunky, it definitely works.

CS4 now lets you create a few basic 3D objects within Photoshop itself, including cylinders, boxes, pyramids, soda cans, and hats. You can also elevate 2D layers to 3D 'postcards' similar to the 3D layers found in

Adobe After Effects. You've always been able to distort 2D layers to provide the illusion of 3D perspective, but 3D postcards are actual 3D layers. They can reflect the environment, cast and receive shadows, and so forth, as well as spin, scale, and move in 3D space.

Photoshop will render 3D objects, imported and created, with ray-traced reflections, refractions, and shadows. You can apply texture maps (colors, images) and bump maps. You can change materials' brightness, reflectivity, shininess, and add and subtract lights.

A 3D-axis tool in one corner provides you with colored, visual cues as to your object's orientation. You can easily scale, rotate, and move your objects with this one intuitive widget. Photoshop also gives you individual tools for transforming your object and navigating your scene.

Still, I find Adobe's 3D navigation a bit awkward, probably only because I'm used to working in "real" 3D apps. I doubt Photoshop's 3D workflow will ever be as optimized as, say, 3ds Max or SolidWorks.

CONFIGURATOR

While it is not actually a part of Photoshop, Configurator is a fascinating little application available via free download from Adobe Labs. Configurator allows you to easily create your own panels and palettes for Photoshop CS4. You can drag and drop tools,

NVIDIA Quadro CX

My computer is a relatively tame dual-core machine with 4GB of RAM. To give it a leg up on Photoshop CS4, I installed a new Quadro CX video card from NVIDIA (MSRP $1999).

The CX is a dual-width PCIe card, with dual DisplayPorts and a single DVI connector. The entire card sits in its own enclosure pierced by a fan on one side to control airflow. It's really a three-slot card, as any card placed directly above it will be about 1/8 in. from the exhaust fan, and will soon become the temperature of molten lead (the card draws 150W). If you've got six slots free, the CX supports SLI allowing you to link two of the cards together.

The CX is promoted by NVIDIA as the CS4 Suite accelerator and, with 192 full-precision 64-bit floating-point processors and 1.5GB of fast GDDR3 RAM, it accelerates the heck out of things. The difference between Photoshop CS4 with and without GPU acceleration is like the difference between hopping in the street with your ankles bound and rolling down the avenue in a red convertible. Huge photos (400 megapixels+) were smooth as glass.

This speed isn't exclusive to the CX—any GPU-accelerated card will take advantage of Photoshop's new features; they're implemented via OpenGL, after all. A lot of the CX's value comes from the rest of the CS4 suite, especially its hardware acceleration of H.264 video encoding in Adobe Premiere and After Effects. If you don't work with video, consider the Quadro FX 4800 instead; essentially the same card, it sells for a few hundred dollars less. *— MC*

menu items, Photoshop actions, and text. What's more, you can drop in images, audio, video, and even SWF and FLV movies and animations. And all of this goes beyond the usual customizable tool palettes many of us have used.

You could, for example, create a panel that guides users through a standard process—such as preparing photographs for printing—complete with carefully organized tools, actions, text hints, and even animated tutorials or videos. All the images, animations, and video are referenced via URL; you're not embedding a 2MB tutorial video within the actual panel.

HELP ME!

Like that of most big applications, Photoshop's help has gone online. Selecting Help from the menu takes you to the Web where you're hooked into help, community forums, and the Adobe knowledge base.

This works pretty well, but sometimes the trees seem to get lost in the forest, and you'll find yourself searching through a list of discussion threads rather than the help entry you'd hoped for.

The new 3D features deserve more space than they're given. Reflective surfaces reflect the environment map, but where do you set the environment map? Help never says.

THE VERDICT

If you regularly work with very large images and you have a newer video card, you'll definitely find this upgrade worth it. Likewise it's worth it if you'd like to integrate Photoshop (MSRP $999; upgrades from $349) more tightly with your 3D workflow, or create your own panels.

Otherwise, this is still a strong upgrade with an improved workflow and plenty of new features I haven't touched on, though older machines—especially those with older video cards—may balk a bit.

Go ahead and give it a try. You can download a trial version of Photoshop and all Adobe's applications from adobe.com/downloads/ ■

*Contributing Editor **Mark Clarkson**, a.k.a. "the Wichita By-Lineman," has been writing about all manner of computer stuff for years. An expert in 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 c/o DE-Editors@deskeng.com.*

INFO

Adobe Photoshop CS4
Adobe Systems, Inc.
San Jose, CA
adobe.com

NVIDIA Corporation
Santa Clara, CA
nvidia.com

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

DE Marketplace

Making Perfection Even Better on theHP xw8600 Workstation

> HP updates its top-of-the-line workstation with the latest Intel Xeon quad-core CPU and NVIDIA Quadro FX graphics board.

BY DAVID COHN

If the HP xw8600 workstation that arrived at our office recently looked familiar, that's because we had indeed reviewed this system before (see *DE* June 2008) and because it also came wrapped in a colorful graphic skin similar to the HP xw4600 we also reviewed earlier this year (see *DE* March 2008). But while the system was virtually identical in appearance, that's where the similarities ended, because inside this new iteration of HP's top-of-the-line Intel-based workstation were a pair of 3.4GHz Intel Xeon X5492 processors and a new NVIDIA Quadro FX 4800 graphics accelerator.

The Xeon 5400-series, formerly code

We take another look at the HP xw8600 workstation, this time equipped with a pair of 3.4GHz Xeon quad-core CPUs and the new NVIDIA Quadro FX 4800 graphics accelerator.

Benchmarking the HP xw8600 Workstation

		HP xw8600 workstation (two 3.4GHz Intel Xeon X5492 quad core CPUs, NVIDIA Quadro FX 4800, 4GB RAM)		**Lenovo Thinkstation S10** workstation (2.66GHz Intel Core 2 Q6700 quad core CPU, NVIDIA Quadro FX 4600, 2 GB RAM)	**Alienware Area-51 ALX Crossfire** workstation (Intel Core 2 Extreme 9650 quad core 3.0GHz CPU overclocked to 4.0 GHz, two ATI Radeon HD 3870, 4GB RAM	**HP xw8600** workstation (two 3.16GHz Intel quad core CPUs, NVIDIA Quadro FX 4600, 4GB RAM)	**HP xw6600** workstation (two 3.0GHz Intel quad core CPUs, NVIDIA Quadro FX 1700, 2GB RAM)	**HP xw4600** workstation (3.0GHz Intel dual-core CPU, NVIDIA Quadro FX 1700, 2GB RAM)	**Appro Xtreme WH 5548** workstation (four 1.9GHz AMD Opteron quad-core CPUs, NVIDIA Quadro FX 5600, 32GB RAM)
Price as tested		$9,307.00		$2,589.00	$6,163.00	$6,915.00	$4,611.00	$2,319.00	$9,217.00
Date tested		12/22/08		6/30/08	3/24/08	12/24/07	12/21/07	12/20/07	12/27/07
Operating System		XP	Vista	Windows XP	Windows Vista	Windows XP	Windows XP	Windows XP	Windows XP
SPECviewperf	higher								
3dsmax-04		52.24	**54.61**	37.88	19.61	35.26	33.16	37.37	**19.23**
catia-02		**63.17**	62.48	48.25	**17.06**	46.11	43.01	46.98	25.30
ensight-03		**54.44**	50.82	43.33	**24.88**	41.89	31.03	30.01	38.74
maya-02		**234.50**	193.15	191.10	**32.16**	175.60	111.20	111.40	128.50
proe-04		52.73	**57.15**	48.86	**13.04**	40.46	38.27	43.98	21.82
SW-01		109.91	**119.29**	90.90	**28.64**	74.28	55.44	62.12	40.32
tcvis-01		**29.84**	27.58	24.46	**6.26**	23.57	15.34	15.06	24.24
ugnx-01		**34.17**	31.14	27.04	**12.75**	25.25	14.90	14.29	31.80
SPECapc SolidWorks	lower								
Score	seconds	**164.71**	n/a	188.01	n/a	174.62	184.36	167.24	**282.78**
Graphics	seconds	**54.18**	n/a	60.13	n/a	57.97	61.52	57.46	**89.03**
CPU	seconds	44.36	n/a	41.48	n/a	50.52	50.03	**40.40**	**69.12**
I/O	seconds	**69.96**	n/a	90.18	n/a	**69.96**	77.04	73.43	**125.34**
SPECapc SolidWorks	higher								
Score	ratio	**4.84**	n/a	4.56	n/a	4.46	4.28	4.82	**2.98**
Graphics	ratio	**3.55**	n/a	3.15	n/a	3.33	3.11	3.29	**2.15**
CPU	ratio	7.27	n/a	7.72	n/a	6.39	6.45	**7.99**	**4.67**
I/O	ratio	**4.52**	n/a	3.51	n/a	**4.52**	4.11	4.31	**2.53**

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

named "Harpertown," consists of dual-die quad-core CPUs manufactured on a 45nm process. The X5492 is one of four halogen-free Intel Xeon processors recently introduced. Intel also announced that all previously launched versions of the Xeon 5200 and 5400 series will also now be halogen free. The CPUs use a Hafnium-based high-k metal gate formula and the X5492 is drop-in compatible with existing systems based on the Intel 5400 chipset, which explains how HP was able to quickly roll out the new CPU in its existing workstation platform. Equally appealing is the fact the X5492 uses just 150 watts, not bad considering its speed. But with the increased power requirements, HP included a 1050-watt 80 PLUS efficient power supply instead of the 800-watt unit in the previous system.

Like the earlier xw8600 we reviewed, the system measures 8.3 in. x 20.7 in. x 17.9 in. (WxDxH) and weighs approximately 40 pounds. The front panel hosts a pair of USB 2.0 connectors, headphone and microphone jacks, and a FireWire connector. There are three external 5.25-in. drive bays, one of which again contained an HP 16X DVD+/-RW dual-layer optical drive with HP LightScribe technology and another housing a 3.5-in. floppy drive.

HP offers a wide range of hard drive options, including SATA drives up to 1000GB and SAS drives up to 300GB. With five internal 3.5-inch drive bays, the xw8600 can accommodate up to 5 terabytes. Our evaluation unit came with the same 250GB 7200rpm Seagate Barracuda drive as in our previous evaluation unit.

The rear panel adds five more USB connectors, PS/2 keyboard and mouse ports, a second FireWire connector, audio-in, audio-out, and microphone jacks, a 9-pin serial port, and two RJ45 LAN connectors for the integrated Broadcom 5755 NetXtreme Gigabit PCIe LAN. There's an additional USB connector inside the case so you can hide a USB-based dongle where it can't be tampered with.

Inside the tool-less chassis were those two Xeon CPUs, each with a big cooling tower and 3-inch fan. Our evaluation unit came with 4GB of RAM installed as two 2GB DDR2 memory modules; 800MHz memory is required with the X5492 processor. The eight available memory slots can accommodate up to 64GB of memory using 8GB DIMMs, or up to 128GB using an optional memory riser.

The motherboard has a total of seven slots, six of them full-length: two PCI-Express x16 graphics slots, three PCI-Express x8 slots (two x4 electrically and one switchable as x1 or x8), one PCI-X 133MHz slot, and one legacy PCI slot. One of the graphics slots was filled with a just-released NVIDIA Quadro FX 4800 graphics accelerator with 1.5MB of GDDR3 memory.

In addition to providing double the memory of its predecessor, the Quadro FX 4800

is based on a second-generation NVIDIA GPU unified-architecture with 192 processor cores (compared to 96 cores in the FX 4600). While this all leads to greater graphics performance, it also increases the board's power requirements to 146 watts. Like other high-end boards, the FX 4800 requires an auxiliary connection to the computer system's power supply and the board's large cooling fan and plastic cowl block access to the adjacent motherboard expansion slot. We'll provide a detailed look at the NVIDIA Quadro FX 4800 in a separate article.

To say that the new xw8600 was fast would be an understatement. This system was incredibly fast, surpassing every other system we've ever tested by a considerable margin. On the SPECapc Viewperf graphics benchmark, the HP xw8600 equipped with the NVIDIA Quadro FX 4800 turned in the fastest results we've ever recorded on all eight datasets.

BREAKING RECORDS AGAIN

Since HP graciously provided a second identical Seagate hard drive with Windows Vista loaded, we were able to perform our benchmark tests under both Windows XP and Windows Vista. Amazingly, thanks in part to the latest video drivers from NVIDIA, the results under Vista actually surpassed those under Windows XP on several of the datasets; and our Vista results were obtained with the full Aero interface enabled. (For a detailed comparison of Vista and XP, see "Vista vs. XP: Windows On the Mat," April 2009 DE. The FX 4800-equipped xw8600 was one of two systems referenced.)

INFO

HP
Palo Alto, CA
hp.com

HP Workstation xw8600

>**Price:** $9,307 as tested ($1,361 base price)
>**Size:** 8.3"x20.7"x17.9" (WxDxH) tower
>**Weight:** 40 pounds
>**CPU:** two Intel Xeon X5492 3.4GHz quad-core w/ 12MB L2 cache and 1600MHz front-side bus
>**Memory:** 4GB (128GB max) DDR2 800MHz
>**Graphics:** NVIDIA Quadro FX 4800
>**Hard Disk:** Seagate 250GB 7,200 rpm SATA
>**Floppy:** 3-1/2" floppy
>**Optical:** DVD+/-RW Dual-Layer Lightscribe
>**Audio:** integrated Realtek audio w/ microphone, line-in, headphone, line-out jacks and jack retasking
>**Network:** dual integrated Broadcom 5755 NetXtreme Gigabit LAN
Modem: none
>**Other:** seven external and one internal USB 2.0, PS/2 keyboard, PS/2 mouse, two IEEE1394 FireWire, and one 9-pin serial
>**Keyboard:** 104-key HP keyboard
>**Pointing device:** two-button HP scroll mouse

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

On the SPECapc SolidWorks test, which is more of a real-world test (and breaks out graphics, CPU, and I/O performance separately from the overall score), the HP xw8600 reclaimed its spot as the fastest CAD workstation in spite of the overhead of managing two separate CPU sockets. We were also pleased to see that the I/O performance exactly matched that of the earlier xw8600 we reviewed. As we noted in our recent review of the HP EliteBook mobile workstation, we are no longer running our SolidWorks test under Vista. We will update this test for future reviews.

The AutoCAD rendering results were even more impressive. Those eight CPU cores enabled the HP xw8600 to complete the test image in just over a minute. Only the Appro Xtreme WH 5548 equipped with a total of 16 processor cores has completed this test faster.

As usual, HP rounds out the xw8600 with its excellent 104-key keyboard and a two-button optical mouse. Users have a choice of 32- or 64-bit versions of either Windows Vista Business or Windows XP Professional. HP backs the system with a limited 3-year warranty on parts and labor and offers more extensive coverage for an additional charge.

Of course, all of this power is going to cost more. Although the starting price for an HP xw8600 has come down at bit (to $1,361), that's for a system with a single 2.0GHz quad-core CPU, 1GB of RAM, an 80GB hard drive, and an entry-level graphics card. At the time of our review, we couldn't price our configuration online, because HP was not yet offering the NVIDIA Quadro FX 4800 as an option. But based on a similarly configured system plus the new NVIDIA graphics accelerator, we estimated the price as tested at $9,307. While that's a lot—and well beyond the needs of many midrange CAD users—it should be very attractive to those looking for the ultimate in performance. ■

*Contributing Editor **David Cohn** is a computer consultant and technical writer based in Bellingham, WA, and has been benchmarking PCs since 1984. He's the former editor-in-chief of* Engineering Automation Report *and* CADCAMNet*, and the author of more than a dozen books. Please send comments about this article to DE-Editors@deskeng.com. You can also contact David at david@dscohn.com.*

JULY 2009 • VOL. 14, ISSUE 11 • A SPECIAL SUPPLEMENT

DESKENG.COM

KNOWLEDGE TO SUCCEED
NOW AND TOMORROW

elements of HPC®

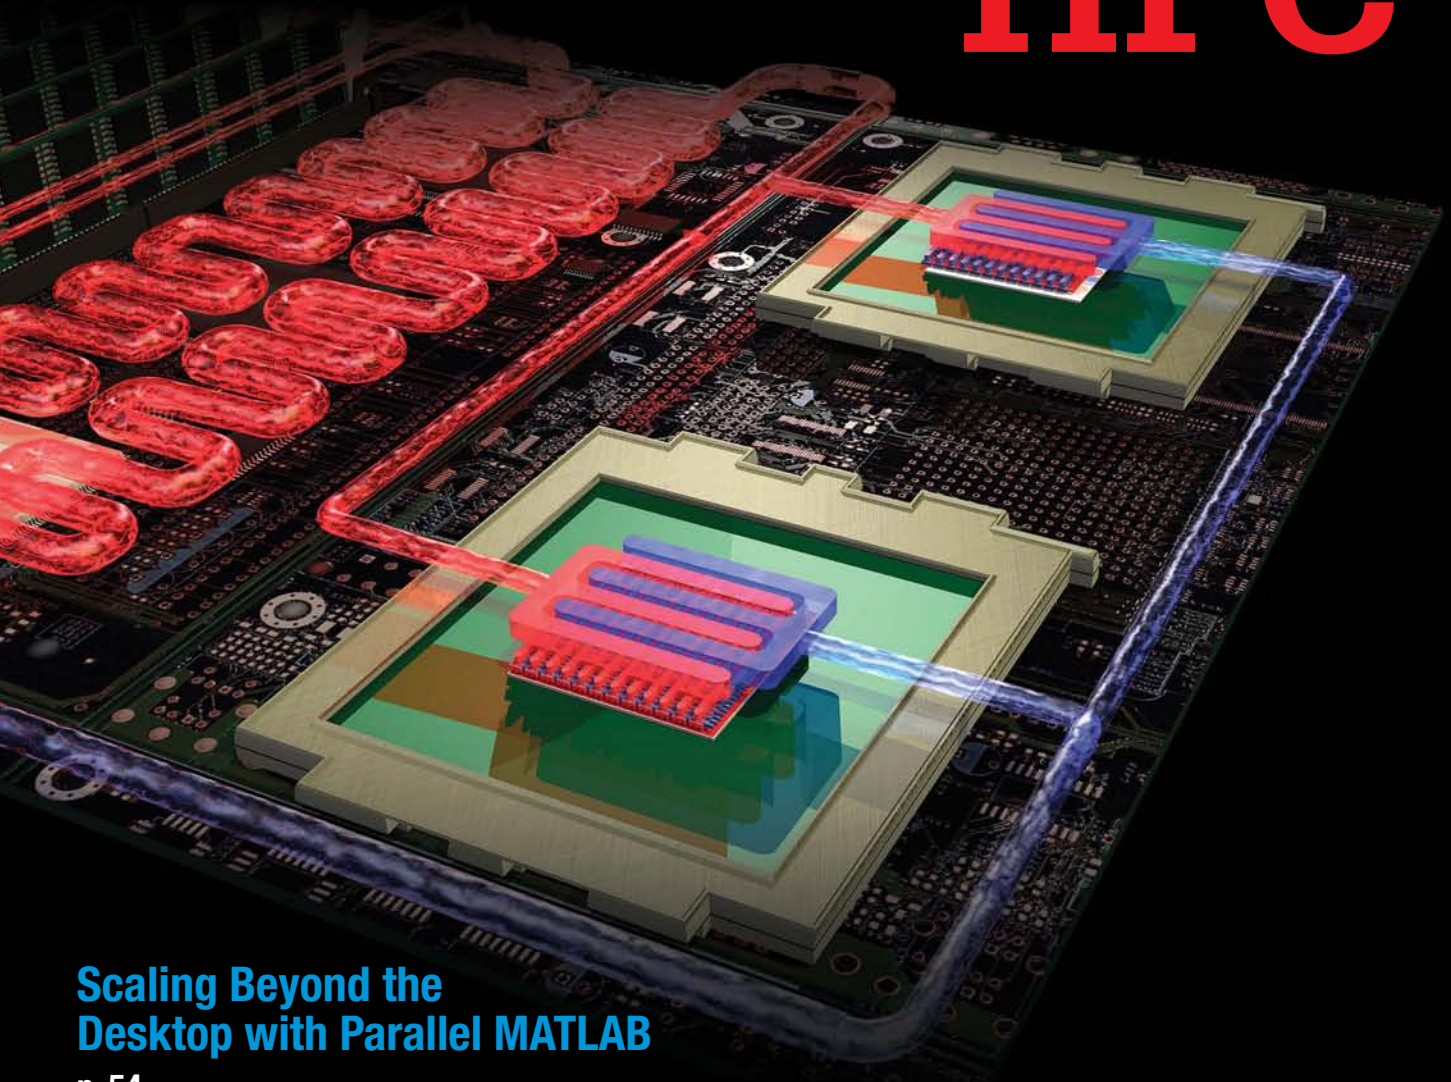

Scaling Beyond the Desktop with Parallel MATLAB

p. 54

Energy Efficiency in HPC Systems is No Simple Matter

p. 58

Multicore Will Change the Way We Think

p. 64

This image is a simulation of a future HPC cooling technology being examined by scientists at IBM's Zurich, Switzerland, research lab. The goal is to bring the water coolant to the hottest location—directly on the chip itself. See page 58 for the full story.

Scaling Beyond the Desktop with MATLAB

Parallel MATLAB apps make programming simple and portable.

BY A. CHAKRAVARTI, P.S. IVANOV AND J. RORISON

Parallel computing technologies offer engineers the means to accelerate solutions of their computational problems by using multiple hardware resources. The ability to solve very large problems by scaling computer programs to run on multicore workstations, clusters, grids, and clouds can help engineers gain significant research and competitive advantages.

Engineers who need their computer programs to run faster or process larger data sets would prefer that every program automatically use parallelism. However, automatic parallelism is still a subject of basic computer science research. Thus, the responsibility of using parallelism to run programs across multiple cores or computers is shared by the designers of programming

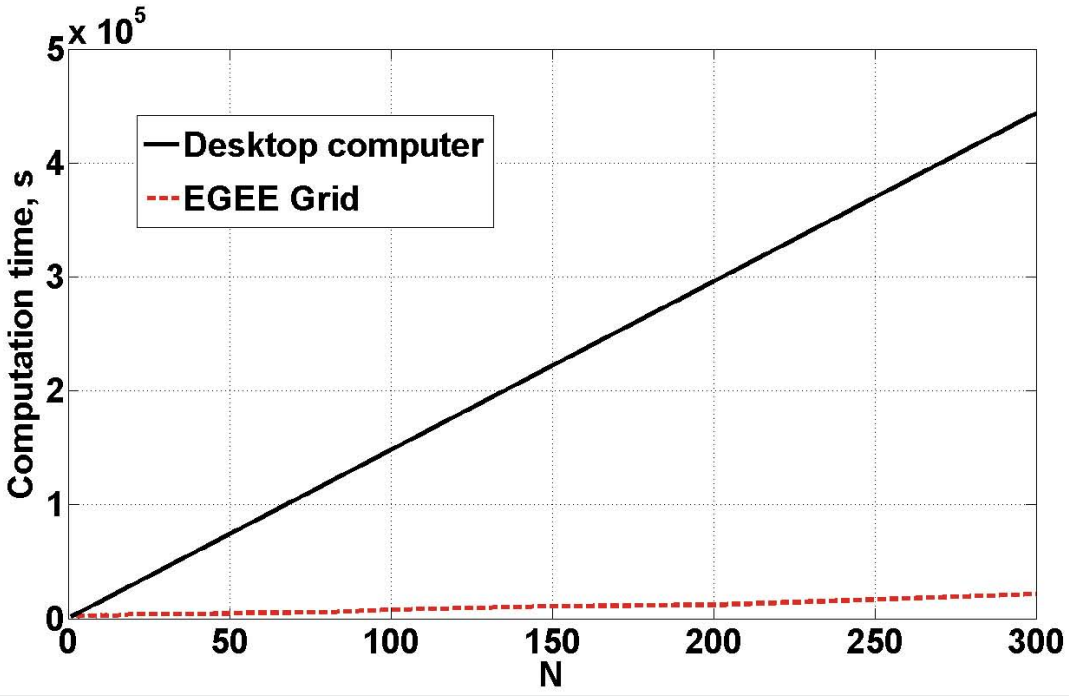

This is a comparison of computation time on one workstation core to time on the EGEE grid.

languages/environments and their users.

This article explores the expectations that engineers have when using scalable parallel computing technologies as well as how the designers of MATLAB have attempted to meet those expectations.

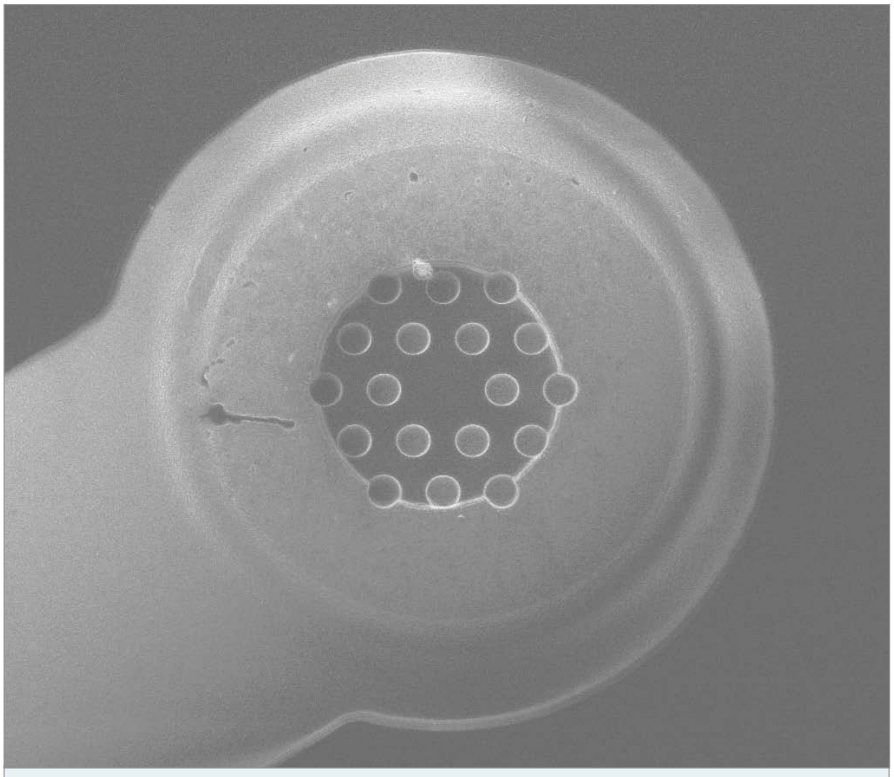

Top view of a PC-VCSEL. The photonic crystal has holes with 2mm diameter and 4mm center-to-center spacing.

Required Capabilities

Engineers are primarily concerned with solving complex problems within their technical domains. While some may be experienced programmers, most would prefer to be shielded from the finer points of parallel programming, such as multithreading, synchronization, and data management across clusters. For this reason, a parallel computing environment needs to make it as easy as possible for engineers to write, use, and maintain parallel programs.

Engineers use a variety of operating systems and hardware. They do not want to change code when migrating applications from one operating system to another, or from a multicore desktop computer to a large cluster. Their colleagues and customers should also never have to modify code. Scalability and portability are key requirements for a parallel computing environment because most engineers would just want their parallel applications to seamlessly use whatever resources are available.

A major roadblock for any engineer who wants to use remote cluster hardware is the need to have specific knowledge about the cluster. Most engineers would prefer that the cluster administrator write system-specific scripts, set environment variables, and manage job queues. The separation of user and administrator tasks is an important requirement.

Scaling Beyond the Desktop

There are a number of parallel computing technologies available to an engineer. Some, such as Intel TBB and Cilk, enable programmers to write parallel programs that use multicore computers. However, the same programs cannot scale up to use remote resources like clusters. Programs often need to be rewritten to use other technologies such as MPI, which are complex and require specialized knowledge. This workflow

violates the important requirement that the same parallel program scale from workstations to clusters without any recoding.

Specialized technologies like MPI have the additional drawback of requiring the user of a parallel program to have some knowledge of the system on which it will be run. This reduces the portability of code and the number of people who can use it.

The MATLAB Approach

The designers of MATLAB attempt to fulfill the parallel computing requirements of ease-of-programming, scalability and portability, and separation of user and administrator roles by offering different levels of control to programmers. Some programs require no recoding, while others require the use of low-level programming methods. The most commonly used programming techniques involve adding annotations to code. For example, a FOR loop with independent iterations can be annotated as a PARFOR loop. At runtime, the computing environment will attempt to run the loop iterations in parallel across multiple MATLAB workers (execution engines that run in the background on a workstation or cluster).

The MATLAB programming language is separated from the execution environment. This means that the same parallel program can run on multicore workstations as well as on clusters, grids, and clouds. An engineer who writes a parallel MATLAB program does not need to know anything about where the program will eventually run, lending portability across different hardware and operating systems. For example, an engineer who develops a program on a Windows workstation can run the same program on a Linux cluster, or share it with a colleague who uses a Mac laptop.

Scaling up a parallel MATLAB program from workstation to cluster does not require the user to have any knowledge of the cluster because MATLAB allows for the roles of user and cluster administrator to be independent of each other. The administrator stores information about the cluster in a configuration file, e.g., how to submit jobs and transfer data, and sends it to all cluster users. In fact, a user could receive several configurations, one for each remote resource. The user imports the configurations into the MATLAB user interface and selects one of them as the resource on which to run the parallel MATLAB program.

Here's the typical workflow of an engineer who wishes to solve a large technical problem in MATLAB: The user writes a serial program and then parallelizes it by using constructs such as PARFOR. The user

INFO

Cilk Arts, Inc.
Lexington, MA
cilk.com

Intel
Santa Clara, CA
intel.com

For more info, please visit deskeng.com.

tests and debugs the program with small inputs on the workstation. Finally, the user increases the size of inputs to the program, imports a configuration for a remote cluster, and reruns the program on that cluster.

MATLAB in the Real World

The Optics group at University of Bristol performs research on Semiconductor Vertical-Cavity Surface-Emitting Lasers (VCSELs), which are widely used in fiber-optical telecommunication networks. The group develops new generations of VCSELs with photonic crystals (PC-VCSELs). In order to perform numerical simulations using models of PC-VCSELs, MATLAB solvers were used for 2D partial differential scalar Helmholtz equations and ordinary differential laser rate equations. The approximated solution time of equations of models varied from 10-700 minutes for some models and 4-60 hours for others. Since these models used many input parameters, the computation of PC-VCSEL characteristics and their optimization required hundreds of solutions of equations. The problem was clearly too large for a desktop system; performing the computations on a lab workstation would have taken several days.

Researchers parallelized the MATLAB program by structuring it as a job that computed parameters of optical modes of PC-VCSELs N times. Therefore, there were N tasks, each of which computed parameters of optical modes. Researchers first tested and debugged the program by using multiple MATLAB workers on a workstation. Once the correctness of the parallel MATLAB program had been established, it was decided to run it on a grid system provided by EGEE (Enabling Grids for E-sciencE), the consortium that provides more than 70,000 processor cores to users worldwide. By using a portion of this infrastructure, the time for computation of 300 tasks was reduced from more than five days to just six hours—an improvement of 21 times.

Parallel programming environments that allow engineers to solve technical problems on and beyond desktop computers are important enablers of technical activity. MATLAB provides engineers with the ability to parallelize programs and to then scale beyond the desktop without modifying the program or writing any additional code. ■

A. Chakravarti *works for The MathWorks, Inc., and P.S. Ivanov and J. Rorison are with the University of Bristol. Send comments about this article to DE-Editors@deskeng.com.*

More than a Question of Energy

Performance, power consumption, and cooling make up a complex stew, leaving the search for energy efficiency in HPC systems with no simple recipe for success.

BY TOM KEVAN

If you're looking for simple, straightforward guidelines on the use of technologies, designs, and best practices to achieve energy efficiency in high-performance computing (HPC) systems, forget it. Answers to this question rapidly become a confusing swirl of what often appear to be conflicting goals and attributes. Performance, parallelization, power consumption, and cooling are critical qualities screaming for balance. Unfortunately, the right balance varies with each application.

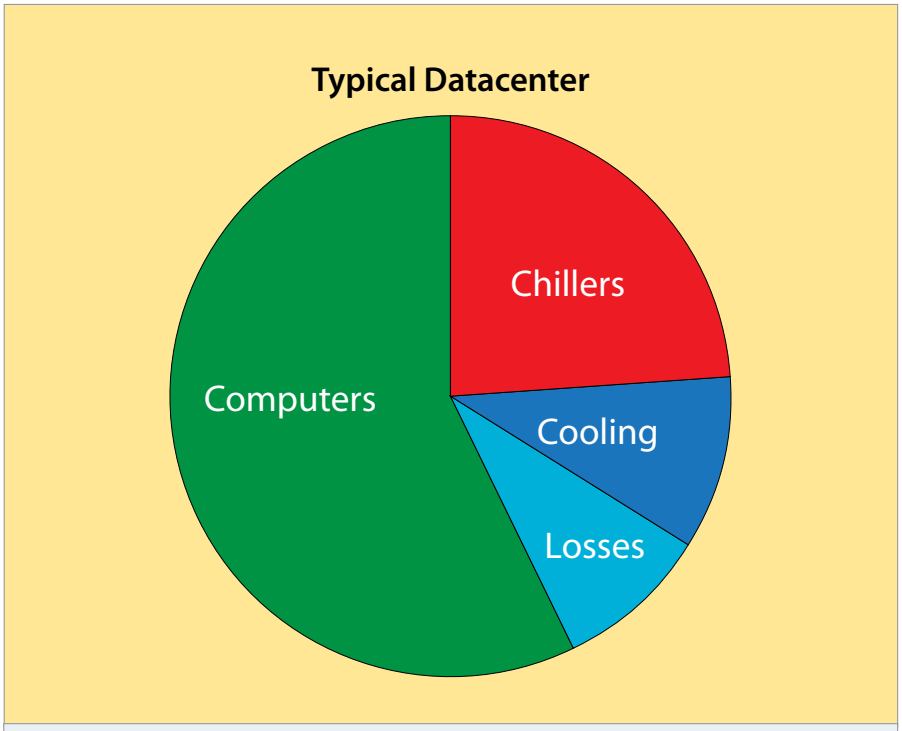

Figure 1: In a typical datacenter, just 50% to 60% of the power entering the datacenter goes to computers and processors while 30% to 35% goes to chillers, computer-room AC units (typically only 60% to 80% efficient), and cooling motors, fans, and pumps. Electrical losses due to inefficiency and AC to DC conversion account for 10% to 15% of power. Chart courtesy of Cray, Inc.

A Question of Balance

Begin by recognizing that performance is the holy grail of HPC. The ultimate goal is to attain the ability to do the computing work—to solve often large, complex problems—as quickly and accurately as possible. Unfortunately, there isn't one design approach that's right for all situations.

Now add to the mix the search for energy efficiency. One approach to reducing power consumption has been to reduce the voltage and frequency of the processors.

Four or five years ago, you had single-core processors that operated at 3GHz. Those

Ready to put Parallel Programming to work for you?

Understand the fundamentals of parallel programming in this new training series from Sun Microsystems.

As the world of high performance computing heats up, and multicore, multithreaded systems move into the enterprise, developers who have honed their parallel programming skills are ready to create applications that reach new levels of scalability, performance, safety and reliability.

To help software developers and engineers meet the challenges posed by parallel programming, Sun Microsystems is now offering a series of educational seminars called “An Introduction to Parallel Programming” where we’ll discuss parallel programming as a core fundamental of application development. Learn about the principals and tools available today at **sun.com/parallelprogramming**

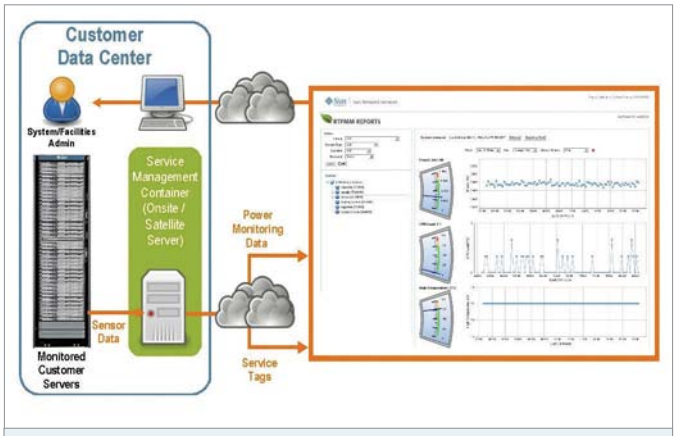

Figure 2: The Sun Intelligent Power Monitoring Service makes it possible for customers to improve power efficiencies via monitoring, reporting, and forecasting energy consumption. This subscription-based network service provides direct visibility via a Web interface. *Image courtesy of Sun Microsystems, Inc.*

ran very hot and consumed considerable energy. Now HPC clusters use dual-core or quad-core processors, and an average system has hundreds of cores, with the largest systems having more than 150,000 cores.

If you look at the frequency of each of the cores, however, it may be half the 3 GHz found in the older systems. That saves energy because you are tuning back each processor core. So it looks as if you are improving performance and saving energy.

Here's the catch. To do the same job, you are using more processors. So even with the cores running at the reduced frequency, you are using more energy. The paradox is that while clusters are becoming more energy efficient, in some ways they are moving away from energy efficiency.

Now think about what you have just done from a software standpoint. You have increased the parallelism of the system. You've just moved from the one-way parallelism of three years ago to four-way parallelism with today's quad-core processors. Don't forget that a lot of engineering software is written for single-core processors. Only some of the applications are easy to parallelize—meaning they can easily be subdivided into lots of independent small problems. If the commercial software you're using doesn't easily parallelize, you may not be able to do the job any faster, and you're using more energy. So you really need to balance performance and power consumption not in the abstract, but in the applications you are trying to run.

Start with the Source

What better place to begin effecting energy efficiency than with the power supply and distribution system. In many ways, this is the locus of the greatest energy loss and the place where the most significant improvement can be made. To appreciate the scale of the problem, consider the Power Usage

INFO

Cray
Seattle, WA
cray.com

Hewlett-Packard
Palo Alto, CA
hp.com

IBM
Armonk, NY
ibm.com

Sun Microsystems
Santa Clara, CA
sun.com

IDC
Framingham, MA
idc.com

For more information on HPC, please visit deskeng.com

Effectiveness (PUE) metric, a measure of the energy efficiency of a data center calculated by dividing the amount of energy entering a data center by the power used to run its computer infrastructure (see Figure 1).

According to several surveys, the average PUE of U.S. data centers is probably somewhere between 2 and 2.4. That means that if you have 2W going into the data center to power everything—the equipment, chillers, cooling towers, pumps, and UPSs—only 1W is going to the IT equipment. So half the energy is used before it gets to the computing systems. Then go one step further and ask: Of the 1W going to the IT equipment, how much is used to provide bits and bytes? The utilization rates of x86 architectures are between 5% and 12%. So if that 1W goes into the rack and you are getting only 5% utilization, only 0.05W is used to perform computing.

Figure 3: The Hewlett-Packard Modular Cooling System uses fans and air-to-liquid cooling heat exchangers to remove heat. The system is used in datacenters to eliminate hot spots without revamping the overall architecture.

Image courtesy of Hewlett-Packard

You can counter this energy loss and achieve greater efficiency by improving the power supply and distribution system on two levels. The first is the process of converting alternating current to direct current. The second is in the distribution of power within the HPC system itself.

You may have 13–15 kVA coming into the building, which you have to reduce so it can be used by the chips (about 1 V). The more steps you take in this process, the greater the energy loss through waste. So you reduce the rectification steps by bringing the high voltage down to only 480VAC, which many HPC clusters and computing racks can accommodate. In doing so, "you improve energy efficiency by 2 percent to 3 percent," says Roger Schmidt, IBM distinguished engineer.

Another way of saving energy at this level is to use high-quality equipment. "There are high-efficiency uninterruptible power supply systems, power distribution units, and transformers on the market," says Schmidt. "You can gain several percentage points by buying the better products on the market."

Next, you can improve energy efficiency by managing the distribution of power within the HPC system itself. A range of hardware and software tools can help you achieve

this on multiple levels (see Figure 2).

"There are tools now in place on some machines that allow you to go in and actually limit the power drawn by the computer," says Schmidt.

HPC system manufacturers are also beginning to implement power-management features that identify components that are not doing any work or processing and turn them off or put them in sleep mode. "Any workload requires the operation of only some of the components at any one time," says Giri Chukkapalli, HPC systems architect at Sun Microsystems. "So this fine-grain power management is critical. Well-integrated hardware and software is the key going forward."

Because HPC systems can have multiple power supplies within a system, it's also important to balance loading. HP uses an onboad microprocessor to do this. "A microprocessor we call the onboard administrator manages the loading of power," says Ed Turkel, manager of the High-Performance Computing Product and Technology Group at Hewlett-Packard. "It will actually distribute the power load across the multiple power supplies in a way that keeps them high up in their efficiency curve."

Additions and Subtractions

You can wring out even more energy efficiency by choosing a system in which the vendor has added extra components on the one hand and eliminated unnecessary elements on the other. Again, you are creating a balance based on the requirements of your application.

"We have an engineering team that works with the major ISVs, and they spend a lot of time with each new system generation, simply seeing how all the standard applications run in that environment," says Hewlett-Packard's Turkel.

A recent trend is the use of alternative processors, which in the case of HPC systems include general-purpose graphical processing units (GPGPUs) and specialized ASICs. At the moment, the alternative processors do not operate independently. They are tied to the x86 processors and function as coprocessors used to perform in PCs.

The primary advantage of using an alternative processor is accelerating the execution of the job. Overall, the processors can either make the job run faster or make it run with less power. With GPGPUs, you are really looking for a significant increase in performance that outweighs any increase in power used by the GPGPU.

"When GPGPUs get to their highest levels of performance, they are power hungry themselves," says Turkel. "They are drawing similar levels of power as the microprocessors themselves—in many instances, even more. So you don't automatically save power by using them, but you may be speeding up the application significantly."

But don't expect alternative processors to deliver their full advantage for a while. "The

software to program the accelerators is still in early stages," says Sun's Chukkapalli. "The libraries required to take advantage of accelerators are in the early stages. We are quite far away from an optimized situation."

On the other hand, you can reduce power consumption by eliminating unnecessary components and subsystems. "One of the aspects of power efficiency is simply not overloading the node with components that are not going to be used or that are not important for the application," says Turkel. "If you look at our ProLiant double-dense blade that we introduced last year, we've made some compromises. Instead of having two rotating disks, it has one. Instead of having a larger memory complement, it has fewer DIMM slots."

Cooling Technology

Because HPC cooling consumes large amounts of energy, the area is another prime place to reduce power consumption (see Figure 3). Generally speaking, there are two main cooling approaches: air and liquid. Liquid can be water or an inert substance. Smaller HPC systems primarily use air, while a growing number of larger systems are turning to liquid.

Cooling is achieved on multiple levels—the chip, board, cabinet, and room. HPC systems generate heat within the cabinets, and the heat is exhausted into the room; a fairly inefficient way to do it. Conventional wisdom is that the closer the cooling mechanism is to the origin of the heat, the more effective it is.

Liquid—water in particular—has much better heat transfer characteristics than air. By mounting cold plates using liquid cooling on top of processor modules and on the back of cabinets, you can effectively eliminate the heat before it circulates into the room. The lower temperatures allow you to raise the frequency of the processor, increasing the system's performance and improving the energy efficiency of the data center.

Compared to air-cooling technology, says IBM's Schmidt, you can realize about a 40 percent energy reduction using liquid. "If you look at the total load of the data center—air-cooled vs. water-cooled—you will save something like 10 percent up to 20 percent of the total energy in the data center," he says. "That includes the IT load and the infrastructure support."

Incremental Improvements

Working with your hardware vendor, you can improve HPC system energy efficiency on multiple levels, with a range of standard technologies. Effective techniques always factor in your application requirements.

Tweak performance at the microprocessor and system levels, using design best practices. Eliminate unnecessary components and subsystems. And make sure the system you select implements high-performance liquid cooling and energy management tools. All of that translates into better performance per watt. ■

***Tom Kevan** is a New Hampshire-based freelance writer specializing in technology. Send your comments about this article to DE-Editors@deskeng.com.*

Multicore Matters

Multicore processors will change the way we think about engineering computing, but only when software catches up.

BY PETER VARHOL

We're in the midst of another computing revolution akin to the advent of inexpensive PCs and the widespread use of 32-bit processors. We're still early in the revolution, and its real benefits may be years away from being realized. But with careful attention to systems and software, engineers can achieve some of these benefits today.

To understand this revolution of multicore processors and their dramatic impact on system performance, you must look back to the origin of Moore's Law. Named after Gordon Moore, one of the inventors of the microprocessor, Moore's Law loosely states that microprocessor density and performance will double approximately every 18 months.

Processor manufacturers such as Intel achieved this standard largely via two methods: higher clock speeds and smaller components. This enabled processors to cycle more quickly and to pack more circuits into the same space.

But neither method could be continued indefinitely. Very high clock speeds are associated with undesirable electronics characteristics as they pushed processors up into the radio frequency spectrum, while more and smaller components result in more challenging heat dissipation issues.

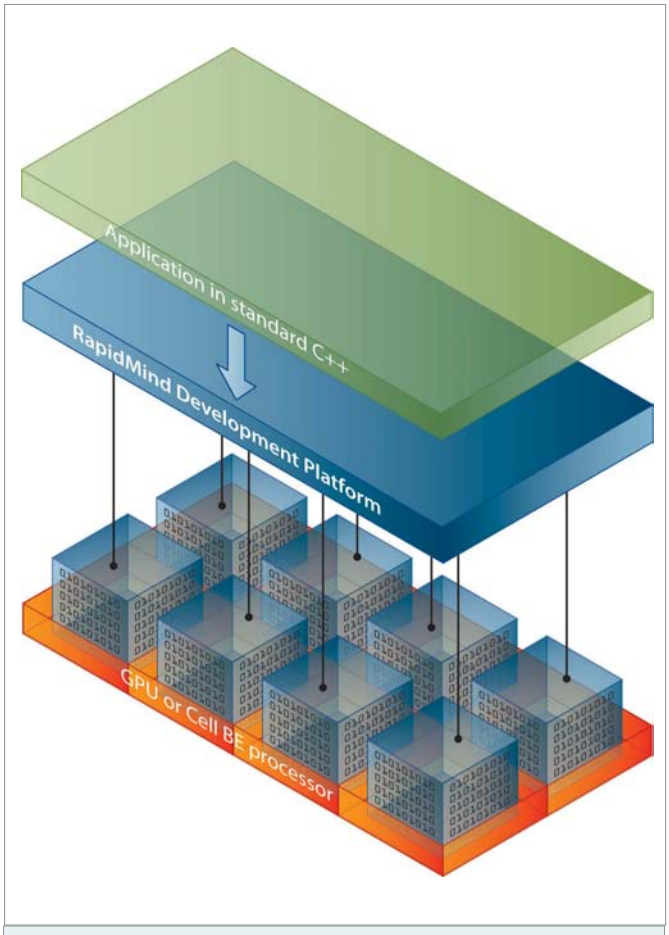

RapidMind lets you take an existing application, place it on top of its runtime platform, and enables the platform to automatically find and implement opportunities for parallel execution on multicore systems.

These limitations didn't kill Moore's Law,

More GFLOPS, Less WATTS

New AMD® and Intel® Technology

Improved Memory and I/O Bandwidth
Clusters and Servers Consume Less Power

Four Hot-Swap Servers in a 2U Chassis

- Up to 48 cores per 2U with Six-Core Opterons or Quad-Core Xeons
- 93% Efficient Hot-Swap, Redundant Power Supply

FasTree™ ConnectX® QDR and DDR InfiniBand Switches and HCAs

Fully Integrated Clusters

- Microway MCMS™ Cluster Management Software
- MPI Link-Checker™ and InfiniScope Network Diagnostics
- Intel Compiler Suite, PGI, GNU and Libraries Installed and Tested

Configure your next Cluster today!

www.microway.com/quickquote

GPU Computing

WhisperStation™

With 1 to 4 Tesla GPUs

Tesla C1060 GPU Performance:

- 1 TFLOPS per GPU
- 4 GB DDR3 per GPU
- 102 GB/Sec Bandwidth
- CUDA SDK

Run MATLAB® on Tesla with "Jacket"
FREE 15-day Trial Available

Clusters With Tesla™ S1070 - 4 GPU Servers

- 36 GPUs + 36 CPUs + 24 TB in 24U
- 40 Gbps FasTree™ InfiniBand
- InfiniScope™ Network Monitoring

AMD Opteron 64 | NVIDIA PARTNERFORCE MEMBER | intel Channel Partner Premier Member 2009

Microway®
Technology you can count on℠

508-746-7341
microway.com

GSA Schedule Contract Number: GS-35F-0431N

however. Microprocessors had already started down the road to using multiple execution pipelines as a way to better use clock cycles. They took the next logical step as the processor designers packaged up not only the pipelines, but also associated core execution circuits, making a single processor die in many ways equivalent to a multiprocessor system.

And now we all have them. This article was written on a dual-core laptop that is already three years old. If you're not already using one, your next computer will almost certainly have a quad-core processor. Consumers of HPC, such as design engineers, care about multicore because of the need for speed, especially for modeling and real-time simulation.

But it's not the clear win that it might appear to be at first glance. Neither operating systems nor application software take full advantage of multiple cores yet. And it may be a decade or more before we can fully take advantage of the power presented to us today.

But systems and software vendors are helping us get there. Intel is providing debugging and profiling tools to those who write and maintain their own code. Microsoft and the Linux community are working on better parallelism in operating systems.

Multicore processors rely on groups of individual cores for their power. A core is a set of circuits and functions that define the essential ability to execute instructions. It includes the processing pipeline, registers, arithmetic units, on-chip caches, and in general everything other than peripherals and buses that connect the cores to each other and to the world beyond the processor. Because the processor is driven by a single clock, each step in the pipeline is performed in lockstep across the cores, but that clock drives more than in single-core processors.

Because we can theoretically do a number of operations on the same clock cycle, there is less need to run at high clock speeds, which saves on heat generation and power consumption. And while the use of multiple cores means that the die size is typically larger, it results in a simpler motherboard than multiple individual processors.

Software—The Hard Part

While it's easy enough to incorporate these multicore processors into desktop systems, writing software to take full advantage of those processors is much harder. To have all of the cores engaged and running code from a single application is a difficult and error-prone process.

Writing multithreaded applications that

INFO

CriticalBlue
San Jose, CA
criticalblue.com

Intel
Santa Clara, CA
intel.com

RapidMind
Waterloo, Canada
rapidmind.com

Virtutech
San Jose, CA
virtutech.com

For more information on HPC, please visit deskeng.com

can run threads in parallel is one of the most technically difficult activities computing professionals can attempt. There are two types of bugs you can experience while writing multithreaded applications. The first is a race condition. If you have two separate threads, and one thread is dependent upon data produced by the other, what happens if the first thread needs the data before the second produces it, then the first thread faults in some manner.

Race conditions are especially difficult to debug, because they seemingly occur randomly. Sometimes the threads complete successfully, because the "race" is won by the thread providing the data. Other times it fails, because the data is not available when the first thread needs it.

The second type of bug is deadlock. If both threads require a resource held by the other, and do not give up that resource, then neither thread can move forward. Because threads are typically designed to hold onto resources until done with them, neither thread can complete.

Because Intel has a vested interest in ensuring the success of multicore processors, it is on the verge of releasing Intel Parallel Studio, a set of plug-ins for Microsoft Visual Studio that help build applications with parallel components. It consists of three components—Intel Parallel Composer, Intel Parallel Inspector, and Intel Parallel Amplifier. Intel Parallel Composer consists of a parallel debugger plug-in that simplifies the debugging of parallel code and helps to ensure thread accuracy. Parallel Inspector detects threading and memory errors and provides guidance to help ensure application reliability. Parallel Amplifier is a performance profiler that makes it straightforward to quickly find multicore performance bottlenecks without needing to know the processor architecture or assembly code.

Intel's Atom uses dual execution cores to deliver high performance in a small package for small computers or embedded systems.

Innovative third-party tools also exist to assist in the building of parallel applications. RapidMind has a platform that lets you take single-threaded applications and run them on top of the RapidMind platform for fully protected multithreading (you also have to make a few minor code changes).

Another company whose technology has the potential to address race conditions is Virtutech, whose Simics simulator provides for reversible execution and debugging. You can also vary the speed of each core on the simulated processor, to help locate the cause of the race condition.

Last, CriticalBlue incorporates two products to enable multicore developers to create

and improve parallel applications. Prism lets developers analyze code for parallel opportunities, perform what-if analyses on parallel strategies, re-code the application, verify its safe execution, and tune the result. Multicore Cascade adds cross-core software partitioning, task dependency analysis, and verification when optimizing software efficiency across architectures by mixing processors and Cascade coprocessors.

None of these tools support widespread multithreaded programming today. RapidMind is limited to C++ (although founder and chief scientist Michael McCool tells me that future versions will support .NET), while Virtutech targets embedded system development. But tools like these are the way we are going to take advantage of our growing processor power in the future.

Benefit to High-Performance Users

At its best, multicore processing provides the ability for a single application to have multiple threads of execution. If you are using a design application, for example, you can complete a design and run a simulation in the background, while at the same time annotating the original design.

At a higher level, multicore processors will also let engineers and other HPC users simultaneously run two active applications. Today, we often have multiple applications open, but are working in only one at a time. In the future, multiple processor cores will let us run two or more applications at the same time, with each actively executing.

What multicore processors won't do is help you get a single task done more quickly. If you have an application thread with dependencies throughout the thread, or have heavy user interaction with a single application, it is likely that you will see little or no improvement in your perceived performance.

Better use of multiple processor cores will happen gradually, in fits and starts. So the performance and productivity improvements of multicore processors are largely still down the road. But there are design software vendors who are delivering the parallel goods today. Check with your software vendor; if they are supporting parallel operation, upgrade to multicore machines immediately. If not, follow their product path, and time your hardware upgrades to coincide with the availability of hardware that can take advantage of it. But be prepared for an explosion of multicore applications sometime in the future. ■

Peter Varhol *has been involved with software development and systems management for years. You can send him comments about this article via e-mail to DE-Editors@deskeng.com.*

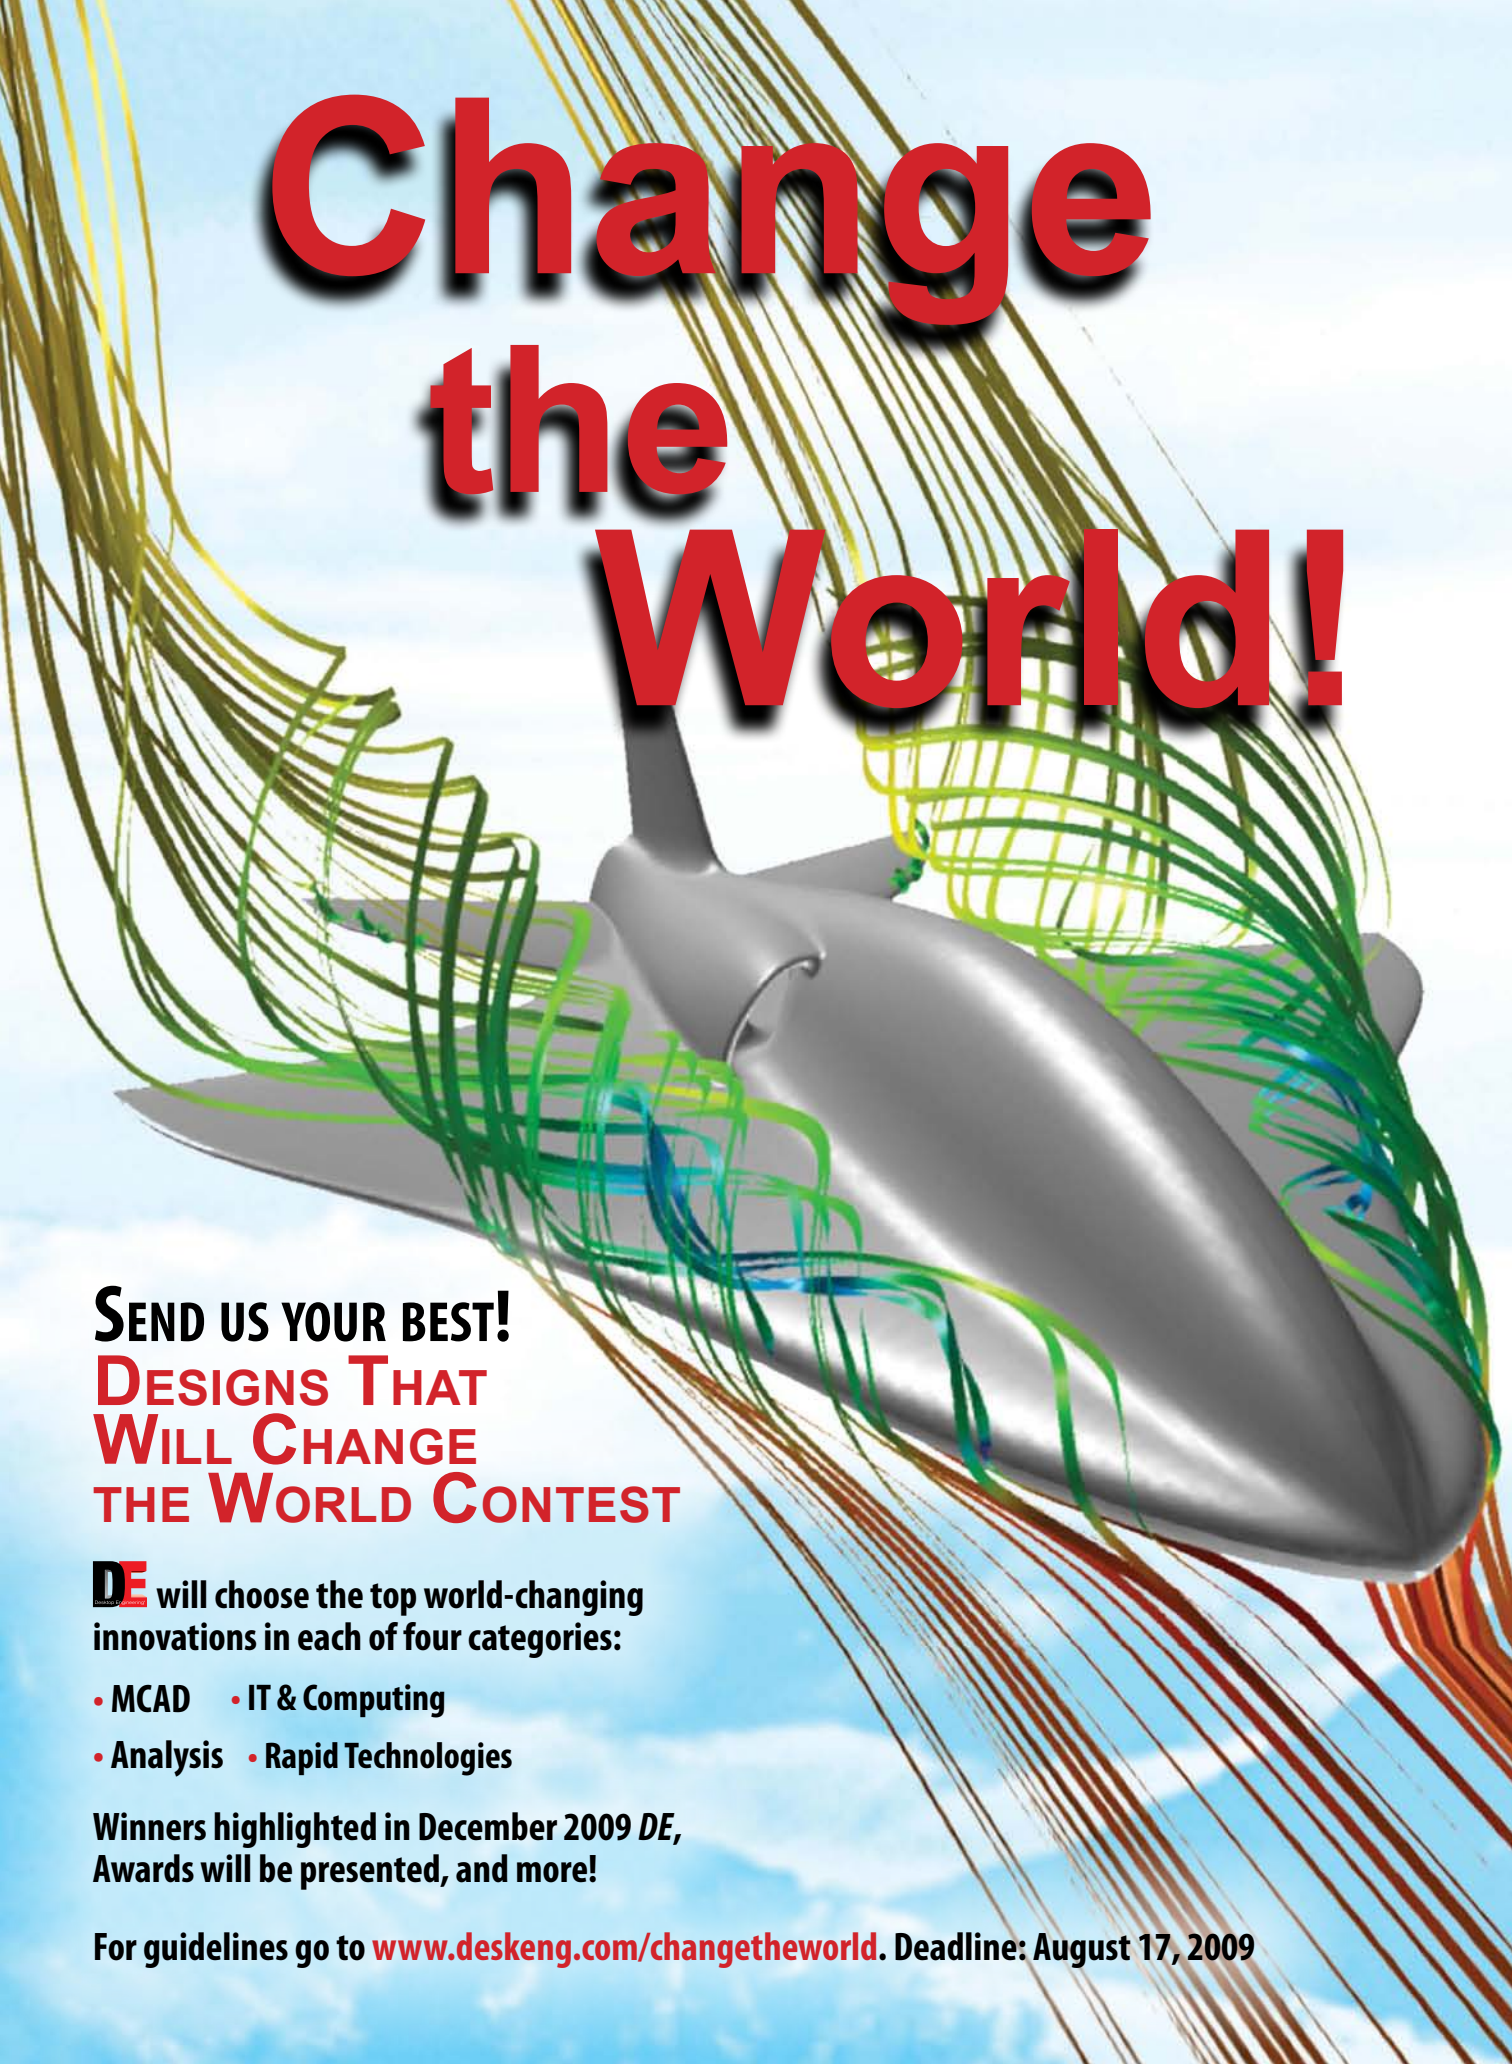
Change the World!
SEND US YOUR BEST!
DESIGNS THAT WILL CHANGE THE WORLD CONTEST
DE will choose the top world-changing innovations in each of four categories:
• MCAD • IT & Computing
• Analysis • Rapid Technologies
Winners highlighted in December 2009 DE, Awards will be presented, and more!
For guidelines go to www.deskeng.com/changetheworld. Deadline: August 17, 2009

SYNERGIS SOFTWARE RELEASES ENHANCED INTEGRATION WITH SOLIDWORKS FOR ADEPT ENGINEERING DOCUMENT MANAGEMENT SOFTWARE

Synergis Software (Quakertown, PA, synergissoftware.com) has announced the availability of Adept PDM 8.1, which features an enhanced integration with SolidWorks 3D CAD design software.

Adept's integration with SolidWorks includes critical document management operations inside the 3D CAD application. The integration provides an Adept Task Pane that lets users search for files by metadata or an Explorer-like browser, then drag and drop the files into the SolidWorks graphics window. It also enables users to sign in/sign out, open, insert, copy, and replace components; determine the status (In/Out/On Hold) of parts, assemblies, drawings and configurations; view thumbnail and parent-child relationship (references) for any selected file; and view any file's associated metadata using Adept's Library Cards. These capabilities are available without having to leave SolidWorks.Adept PDM 8.1 also features an Adept relationship tree—similar in look and feel to SolidWorks' FeatureManager design tree—with icons that show the status of a file; tool tips that display the file name, status, Library or owner; and check boxes to select files for sign out.

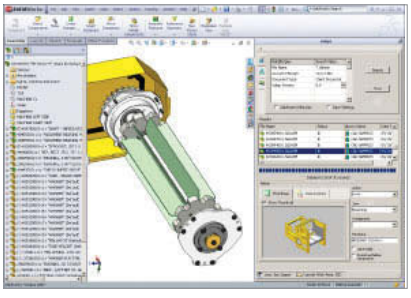

VX 2009 VERSION 14.01 UPDATE BRINGS MORE ENHANCEMENTS

VX Corporation (Melbourne, FL, vx.com) has announced the release of the v14.01 upgrade to VX 2009. This is the first point release, since the March production release of VX 2009 v14.0.

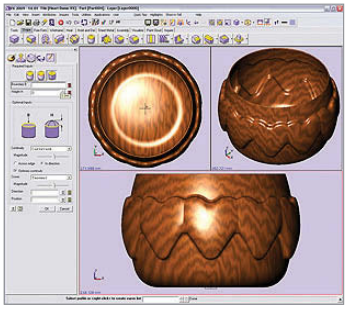

In addition to a new STEP translator, v14.01 has more modeling, drafting, and machining enhancements, as well as several GUI improvements.

New features include improved support of primitive surface types on import and export; a switch to export an assembly as separate components, or as shapes in one part; easy dynamic viewing with a more intuitive rotate and pan; easier access to run Show-n-Tell 3D Tutorials that users have downloaded from the on-line learning center; and updated QuickTips provides more helpful information to VX users. The update also includes new Intelligent Feature Machining (Mill2 Tactic) for automatic operation and tool selections based on 2D part features. It also features faster calculation times for 2-axis feature machining.

Milling and drilling enhancements include 5x sidecut enhancements tighten up the

tool path, new parameters give users more control with 5x Swarfcut, and hole Drilling now supports simultaneous spot drilling.

LASER DESIGN OFFERS SLP-500 HIGH-SPEED LASER SCANNING PROBE

Laser Design, Inc. (Minneapolis, MN; laserdesign.com) has announced that its line of 3D laser scanning probes has a new member, the SLP-500.

The lightweight SLP-500, with a laser line length of 50 mm, scans at high speeds with an accuracy of up to 20µ and is primarily used for mid- to large-sized objects such as molds, stampings, and sheet metal.

All the SLP models, which also include the SLP-250 designed for small parts, the SLP-400 for medium-sized objects, and the SLP-2000, can be used with the company's integrated WS-Series and DS-Series lines of CMM-based scanning systems as well as with Faro portable CMMs.

The new SLP-500 probe gathers data with higher densities than before, upping the resolution by 50 percent to 750 points along the laser line.

LFOUNDRY RELEASES PROCESS DESIGN KIT FOR TANNER EDA'S ELECTRONIC DESIGN SOFTWARE

LFoundry (Munich, Germany; lfoundry.com) has announced a process design kit (PDK) for A/MS electronic designs, developed using Tanner EDA's HiPer Silicon software.

Autodesk Releases New Upgrades for Info modeling

Revit Architecture 2009, Revit Structure 2009, Revit MEP 2009, AutoCAD Civil 3D 2009 and Autodesk NavisWorks 2009 now are shipping, according to **Autodesk** (San Rafael, CA). These building information modeling (BIM) software solutions provide a way to coordinate reliable information about a project through the design phase, following it through to construction and operations.

Using BIM, designers and other building stakeholders can create digital design information and documentation that can be used to accurately analyze performance and cost. The outcome of projects designed and managed by BIM has been that building teams can deliver projects faster and more economically, and can incorporate sustainable design information so that they can reduce the impact that buildings can have on the environment.

The 2009 Autodesk BIM Solutions include NavisWorks 2009 software solutions to enable users to pull together design models and data from documents created on both Autodesk and non-Autodesk software applications.

Revit and NavisWorks drew the most visitors to deskeng.com in May.

ADVERTISER ..... WEBSITE ..... PAGE

AMD-ATI ..... ati.amd/firepro.com CV2

AVEVA, Inc. ..... aveva.com CV3

CD-adapco ..... cd-adapco.com 27

COMSOL, Inc. comsol.com/conference2009 7

Maplesoft ..... maplesoft.com 9

Omega Engineering, Inc. ..... omega.com 1

Quickparts ..... quickparts.com 17

SolidWorks 3dudesgone3d.com 3

Stratasys-Dimension 3D ..... dimensionprinting.com/desk CV4

TecPlot, Inc. ..... tecplot.com 43

DE ELEMENTS OF HPC

Microway, Inc. ..... microway.com 65

Sun Microsystems ..... sun.com/parallelprogramming 59

DE MARKETPLACE

EME Corp. ..... largedocuments.com 47

Laser Design, Inc. ..... GKS.com 47

MEGALITH Ltd. ..... cadd.megalith.net 47

Okino Computer Graphics, Inc. ..... okino.com 47

Optimal Solutions Software LLC ..... optimalsolutions.us 47

Tormach LLC ..... tormach.com 47

DE PRODUCT SHOWCASE

Cadre Analytic ..... cadreanalytic 74

Concurrent Analysis Corp. ..... caefem.com 74

EASA, Inc. ..... easasoftware.com 74

Industrial Press ..... industrialpress.com 74

PDE Solutions, Inc. ..... pdesolutions.com 74

Saelig ..... saelig.com 74

LFoundry is providing Tanner EDA users with the LF150 modular 0.15 µm Low Power and RF CMOS process that provides up to six levels of aluminium interconnect, a polymide passivation, and I/O voltages of 1.8V, 3.3V and 5.0V. Optionally a MiM capacitor is also available.

HIGHLAND TECHNOLOGY ANNOUNCES ITS V420 RESISTANCE SIMULATOR

The V420 simulates resistors, RTDs, thermistors, strain gages, and load cells, and remains accurate for DC, AC, or pulsed excitation, according to **Highland Technology** (San Francisco, CA; highlandtechnology.com).

Users directly program resistances or simulated RTD temperatures. The company designed the V420's resistance changes to be fast and monotonic with no contact bounce or transient switching errors.

The V420 simulates eight two- or four-wire resistive sensors. Channels are individually programmable in four ranges, from 5 ohms to 65.5K ohms and operate from microvolts to 24V/50mA. Channels are fully isolated with up to 750V common-mode range with overvoltage protection up to 50V. The module provides channel calibration check capability and optional built-in self-test.

THERMOCOUPLE TEMP AND PRESSURE DATA LOGGER

OMEGA's (Stamford, CT; omega.com) OM-CP-PRTC210 series of pressure and thermocouple temperature data loggers accepts thermocouple types J, K, T, E,

R, S, B and N. Models are available with pressure ranges from 2 to 345 bar (30 to 50,000 psi).

The series features high-speed download, real-time operation, and programmable start time. The miniature-sized OM-CP-PRTC210 is CE compliant and has up to two years of battery life. This products are well-suited to chemical and wastewater industries

SHERBORNE SENSORS DEBUTS SERIES T233/5 SERVO INCLINOMETERS

Sherborne Sensors Limited (Hampshire, UK; sherbornesensors.com) has announced the North American debut of its Series T233/5, a family of compact, dual axis (x and y), gravity-referenced servo inclinometers, designed for use in military, commercial, and industrial applications.

Series T233/5 offers a ± 15 VDC input voltage, and ± 5 VDC output signal, with self-test on both axes. Units are capable to a resolution of 0.1 arc second. Each axis is fully conditioned, with total electrical isolation between axes. Models are silicone oil and electric damped, and available in ranges from ±1° to ±90°. Series T233/5 offers either solder pin leads or a connector as standard.

For more information, visit Sherborne Sensors.

Saelig Introduces New UM203 Mephisto Scope

Saelig's (Pittsford, NY; saelig.com) UM203 Mephisto Scope is a USB-powered multifunctional instrument that delivers the functionality of seven instruments in one box, according to the company. It features a 2-channel 16-bit 1MSa/s digital storage oscilloscope, an FFT spectrum analyzer, a 2- channel volt meter with pointer display (DC/AC/trueRMS), an analog data-logger, a multi- channel digital data-logger, a 16-channel 100kHz logic analyzer, and a digital 24-line I/O box.

The UM203 is housed in a metal case. With its low power consumption, the UM203 Mephisto Scope is suitable for use as a mobile toolkit connected to a notebook or netbook PC.

The Mephisto UM203 has the ability to perform stand-alone 100kSa/s data logging to an SD memory card, independent of a PC (a 256 MB SD card is included). In this application, an included external power supply is required. When connected to a PC, no power supply is needed because power is drawn from the USB port. Measurement commences by simply inserting an SD card. Analog and digital data can be logged to the SD card.

Windows software is provided to make the instrument quick to set up, with separate control/display panels for all seven devices. The Mephisto UM203 is also expandable with optional accessory modules that provide relay or opto-isolated outputs.

The Mephisto UM203 is available for $895.

To appear in the next Product Showcase, contact Jeanne DuVal at 603-563-1631, ext. 274 or jduval@deskeng.com

Multi-Physics Finite Element Solutions for Partial Differential Equations

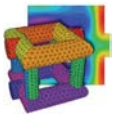

FlexPDE6 is the latest release of the pioneer multi-physics scripted finite element model builder and problem- solving environment for partial differential equations, and is still the best value.

Input your equations – output graphical solutions!

Ideal for end users or developers of niche products.

Present your equations, boundary conditions, domain description and graphics output requests in a readable, self-documenting script format. FlexPDE automatically builds a mesh, constructs a Finite Element model from your equations, and solves the system.

PDE Solutions, Inc.
pdesolutions.com

CADRE Pro 6 For Windows Vista/XP/2000

CADRE Pro 6 For Windows Vista/XP/2000 (32 & 64 Bit) with advanced modeling tools including CAD importing. A practical and affordable tool for engineers. CADRE Pro solves mechanical and architectural structures for internal loads, beam section or plate stresses, displacements, vibration modes, and natural frequencies. Advanced features for stability, buckling, shock and seismic loads. Eleven element types including specialized beams and plates.

CADRE Analytic
425.392.4309 • **cadreanalytic.com**

Got Spreadsheets?

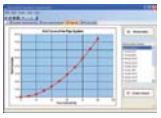

EASA enables rapid transformation of your spreadsheets into robust, collaborative, web applications. You can avoid costly mistakes due to usage errors, eliminate version control problems, and even integrate key spreadsheets with your CAE or PLM environment. EASA allows you to take spreadsheets beyond the limited role of a desktop solution and into the role of collaborative enterprise tool.

EASA
info@easasoftware.com • 1.800.711.5346
easasoftware.com/spreadsheets.html

Beginning AutoCAD 2010 Exercise Workbook!

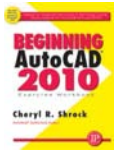

Just published! Updated for 2010/2010LT, including Parametric Drawing, Geometric Constraints, Dimensional Constraints, Application Menu, Initial Setup, and Realtime Search for commands. Comes with a 30-day trial version of AutoCAD 2010 and is ideal for classroom instruction or as a self-study tutorial. Also available … the latest AutoCAD Pocket Reference, providing essential "how to" information for the everyday use of AutoCAD.

For more information or to order, visit our website.

Industrial Press, Inc.
Phone 212.889.6330 • Fax 212.545.8327 • **industrial-press.com**

CAEFEM

Solution of huge problems at a fraction of cost. Analysis of thirteen million node problem in three and half hours. CAEFEM, developed in C/C++, handles linear static, normal modes, buckling, transient dynamic, modal transient, frequency response, response spectrum, nonlinear static and dynamic, nonlinear steady state and transient hear transfer analyses, interfaces to FEMAP, NASTRAN bdf, op2. For a limited time, CAEFEM is available for $995.

Concurrent Analysis Corporation
info@caefem.com • 805.375.1060 • **caefem.com**

Unique OEM Electronics!

Remarkable test equipment, software, and components with specs or prices you won't find ANYWHERE else: premium quality, economical LCD oscilloscopes, spectrum analyzers, USB and Logic analyzers, engineering software, pure RF sources, DMMs, wireless and wired data loggers, test equipment, touch-sense ICs, Ethernet hard-wired TCP/IP ICs and modules, USB I/O ICs, high-reliability PCs, wireless and I/O control, Ethernet-serial/USB-serial/RS232-485, software, etc. FREE Elektor magazine.

Saelig
info@saelig.com • 1.888.7SAELIG • **saelig.com**

Successful PDM Implementation for SMBs

Dave Chadwick
Siemens PLM Software

Small and medium-sized manufacturing businesses (SMBs) are implementing product data management (PDM) to improve efficiency and reduce time to market. Given limited resources, these organizations need to carefully plan and phase their PDM implementations to ensure success.

The market for PDM software and services among SMBs will grow faster than the market for enterprise PDM solutions, says CIMdata, and will represent 29 percent of the market by 2012. SMBs are increasingly implementing PDM solutions to address critical issues like the data explosion that results from the move to 3D CAD, reducing errors and rework, and collaborating more efficiently around the globe.

The benefits of implementing PDM in larger enterprises (such as reducing the time spent searching for data, improving accuracy of that data, and more efficient change processes) are equally important, if not more critical, to the success of SMBs, and yet many of them have not implemented PDM. These are companies typically concerned about the costs of implementing as well as maintaining and developing PDM systems. Typically they do not have IT departments, so implementing PDM requires careful planning and execution.

> With good planning, a PDM system will improve SMB competitiveness.

Several PDM systems are now available that specifically target SMBs with claims of being “out of the box” and “easy to implement.” These systems largely fall into two categories. First are those developed from “workgroup” or document management

solutions. Second are those from vendors of established enterprise PDM solutions.

One example of the latter is Teamcenter Express from Siemens PLM Software. This approach has key advantages for SMBs. There's the scalability of the solution to handle growing volumes of data and the possibility to easily add applications to meet needs not foreseen at implementation.

With either approach, SMBs can increase the probability of long-term success by adopting a phased approach to deployment. With a phased approach, an SMB can get an earlier return on investment. It is important to establish achievable goals for each phase of the implementation and to collect metrics before and after each. For example, if you can show management that one phase has trimmed the average time to complete a change order, it will be easier to get commitment for further phases.

There are four key areas to focus on for a successful PDM implementation:

- Manage the explosion of CAD data. PDM systems capture all the part, assembly, and drawing files, and all the links between files, versions, and revisions associated with 3D CAD systems automatically. They also manage this data, even from multiple CAD formats, resulting in improved reuse of it and more standardization.
- Identify the key issues that are impacting business. Is faster introduction of new products critical? Is improved management of engineering change critical? Further phases of implementation should focus on identified areas and should improve efficiencies.
- Train, train, train. In all phases, training users is critical, but SMBs are limited in resources. New training tools such as flash-based self-paced training guides can be of great benefit in this area. These utilities simulate the PDM environment and enable users to work through commonplace cases interactively without the need for the installation of a training database.
- Finally, plan for the future. With the soaring costs of raw materials and employment, applying PDM to the area of resource optimization will see increasing focus. Whether by reducing the amount of scrap and rework in the manufacturing process or by integrating project management into the PDM environment for early identification of bottlenecks, PDM will have an increasing role to play in improving SMB manufacturing competitiveness. ■

David Chadwick *is the product marketing manager for Teamcenter Express at Siemens PLM Software. Send e-mail about this to DE-Editors@deskeng.com.*

Lifecycle Solutions...
Powered by AVEVA NET

Do you have challenges around your operations, maintenance and safety data?

Does your information strategy rely on many different applications, that can't communicate with each other?

Are you concerned that lack of operational information is affecting your employees' ability to keep the business safe, compliant and profitable?

AVEVA NET is an application-neutral environment for all your operational, maintenance and safety data, no matter which program created it. AVEVA NET is also a platform to enable you to make data available and useful to a wide range of disparate applications, including ready-made lifecycle solutions from AVEVA and applications from other vendors.

Ask us about AVEVA NET and our lifecycle information management solutions. Email us at: sales.us@aveva.com or call: 713 977 1225.

AVEVA

CONTINUAL PROGRESSION

AVEVA Inc
10350 Richmond Avenue
Suite 400
Houston, TX 77042, USA

Tel +1 713 977 1225
Fax +1 713 977 1231
sales.us@aveva.com
www.aveva.com

STRATASYS
Make It Real™

FINALLY, THERE'S ROOM ON YOUR DESK AND IN YOUR BUDGET FOR 3D PRINTING.

Introducing uPrint™ Personal 3D Printers. Build affordable 3D models right at your desk.

The uPrint Personal 3D Printer turns your ideas into durable plastic models you can hold and test. With just a 25" x 26" footprint, uPrint makes 3D printing part of your personal workflow by producing models you can use for proof of concept, functional testing, product mock-ups and much more.

Watch your ideas take shape in a whole new dimension with the uPrint Personal 3D Printer. **Learn more online at www.dimensionprinting.com/desk**

ONLY

$14,900* USD

uPrint™ by dimension.

*Manufacturer's worldwide list price. Additional options, shipping, applicable taxes/duties not included. ©2009 Stratasys, Inc.