digitaleng.news /////

DE

Digital Engineering

November 2017

FOCUS ON: Engineering Computing P.14

Engineers Harness Processor POWER

by Jamie J. Gooch

The Clouds are Clearing

THERE WAS A TIME, not long ago, when cloud computing was met with skepticism, and cloud-based subscription licensing was met with outright hostility. Those days are gone for startups and companies with startup mentalities.

With so many long-established companies scrambling to stay ahead of (or catch up to) disruptive upstarts, it seems like cloud computing is a foregone conclusion. It's hard to argue with many of the benefits, especially for smaller companies and startups.

"Moving to the cloud allowed us to get rid of distractions and focus on what we do best," said Jeff Walters, vice president of Engineering at Globe Trailers, during a panel discussion at Dassault Systèmes' 3DEXPERIENCE Forum last month. "We look at IT as a non-value add to the customer."

Fellow panelist Javier Glatt, CEO and co-founder of CadMakers, Inc., agreed. "When you bootstrap a startup, you want to spend every minute of every day with your hair on fire: find a customer, serve the customer, learn, get better, find another customer ... we don't want to be spending time with IT stuff ... Strategically, we want to play as a business where the world is going, not where the world was."

But even larger companies are finding value in the cloud. The third panelist, Kavi Parupally, is senior director of Business Applications at Rockwell Collins. The aviation solutions provider has 30,000 employees worldwide and is about to go live with its product lifecycle management (PLM) cloud migration.

"Before we start any initiative, we ask 'Why not cloud?'" Parupally said. "We can't compromise on the capabilities and the functionality we can offer to the business. User experience is something we consider to be paramount. With that said, we also need to take into consideration other things. What's the reach that you want to bring in? If you want to have a global reach, suddenly cloud brings a lot more flexibility ..."

Not So Fast

Earlier this year, at the NAFEMS World Congress in Stockholm, Sweden, Andrew Jones said "the marketing around cloud is, at best, unfortunate. It's brash, misleading, relentless. The promises are tempting: lower cost, easier to use, no vendor lock-in, unlimited resources, use what you need, no queues, etc. The reality is less attractive."

Jones is vice president of the high-performance computing (HPC) business for the Numerical Algorithms Group (NAG), a provider of computational software as well as consulting and HPC services. He was speaking to a room full of simulation experts about high-performance computing options.

"In many cases ... it's more expensive than doing an in-house cluster," he continued. "There is a setup effort and you do get vendor lock-in. Not through the inability to move your code or your workflow from one vendor to another, but the sheer fact that moving any volume of data from one cloud vendor to anywhere ... is a pain in the backside and infeasible in some cases."

Innovate Like a Startup, Do Your Homework Like a Stalwart

It's a good point to keep in mind, whether you're looking to move simulation runs to the cloud, simply collaborate on designs or migrate an entire PLM workflow to the cloud. Not all cloud implementations are created equal, and the cloud is better at some applications than others.

"The cloud is definitely becoming an option and must be considered for an increasingly wide range of HPC simulation use cases," Jones said. "It is very valuable in many cases. My comment, I guess, really, is explore the marketing around cloud with some skepticism."

Even startups and small companies should compare the long-term costs of the cloud vs. on-premise solutions for various applications. However, I got the sense that the type of caution and skepticism Jones advised wasn't going to be followed by the two young companies on the 3DEXPERIENCE Forum panel.

"The 2017 Millennial culture—the fickle world of software consumers—people want what they want right now," Glatt said. "It should work like Snapchat works. Something very simple. In AEC (architecture, engineering and construction), cloud is not a benefit; basically you have to do this."

Likewise, Walters referred to the cloud as "cool." He said his team loves configuring new servers and playing with the latest technology, but the cloud lets them get to work designing products. It wins the cool new toy competition.

One thing all of the speakers agreed on was summed up by Jones: "It's not really a question of if you're going to use HPC, it's a question of when and how." DE

Jamie Gooch *is editorial director of* Digital Engineering. *Contact him via jgooch@digitaleng.news.*

DE
Digital Engineering

TECHNOLOGY FOR OPTIMAL ENGINEERING DESIGN

INSIDE

digitaleng.news

November 2017
Vol.23 • No.2

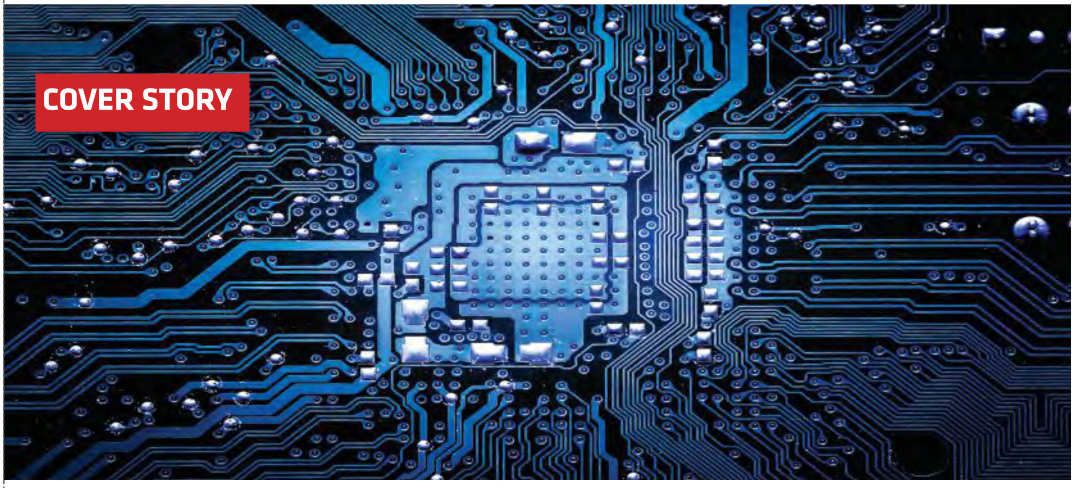

COVER STORY

FEATURES

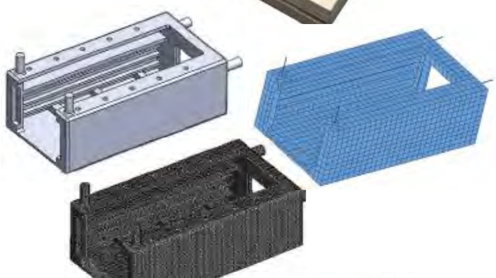

ON THE COVER: Image courtesy of Thinkstock

FOCUS ON ENGINEERING COMPUTING

DEPARTMENTS

PUBLISHER
Tom Cooney

EDITORIAL
Jamie J. Gooch | Editorial Director
Kenneth Wong | Senior Editor
Anthony J. Lockwood | Editor at Large
Stephanie Skernivitz | Associate Editor
Sarah Petrie | Copy Editor

CONTRIBUTING EDITORS
Tony Abbey, Brian Albright, Mark Clarkson, David S. Cohn, Tom Kevan, Randall Newton, Beth Stackpole

ADVERTISING SALES
Tim Kasperovich | East Coast Sales
440-434-2629
Tom Cooney | Midwest/West Coast Sales
973-214-6798

ART & PRODUCTION
Darlene Sweeney | Director
darlene@digitaleng.news

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

ADVERTISING, BUSINESS, & EDITORIAL OFFICES
***Digital Engineering*® magazine**

Peerless MEDIA, LLC

Peerless Media, LLC
111 Speen St., Suite 200,
Framingham, MA 01701
Phone: 508-663-1500
E-mail: de-editors@digitaleng.news
www.digitaleng.news

Kenneth Moyes | President and CEO, EH Media

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

BPA WORLDWIDE

Digital Engineering*®** (ISSN 1085-0422) is published monthly by Peerless Media, LLC, a division of EH Media 111 Speen St., Suite 200 Framingham, MA 01701. Periodicals postage paid at Framingham, MA and additional mailing offices. *Digital Engineering*® is distributed free to qualified U.S. subscribers. **SUBSCRIPTION RATES:** for non-qualified; U.S. $108 one year; Canada and Mexico $126 one year; all other countries $195 one year. Send all subscription inquiries to MeritDirect, *Digital Engineering*®, PO Box 677, Northbrook, IL 60065-0677. ***Postmaster: Send all address changes to MeritDirect, *Digital Engineering*, PO Box 677, Northbrook, IL 60065-0677.

Address all editorial correspondence to the Editor, *Digital Engineering*. Opinions expressed by the authors are not necessarily those of *Digital 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.

Cloud Services, 2017

$260.2 Billion
Public cloud services

$58.6 Billion
Software as a service

$34.7 Billion
Infrastructure as a service

The worldwide public cloud services market revenue is projected to grow 18.5% in 2017 to total $260.2 billion, up from $219.6 billion in 2016.

Software as a service (SaaS) revenue was far greater in 2016 than expected, reaching $48.2 billion. SaaS is also growing faster in 2017 than previously forecast, leading to a significant uplift in the entire public cloud revenue forecast.

The highest revenue growth will come from cloud system infrastructure services (infrastructure as a service [IaaS]), which is projected to grow 36.6% in 2017 to reach $34.7 billion.

— Gartner, Inc., Oct. 12, 2017

Semiconductor Sales

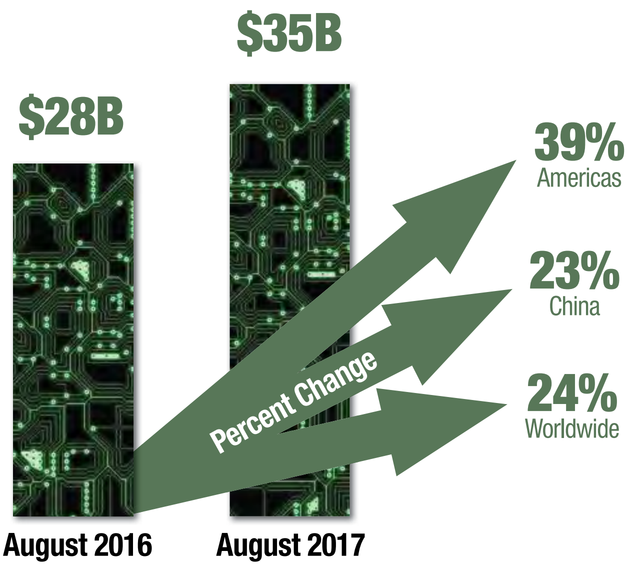

Worldwide sales of semiconductors reached $35 billion for the month of August 2017, an increase of 23.9% compared to the August 2016 total of $28.2 billion and 4% percent more than the July 2017 total of $33.6 billion. The Americas market led the way with growth of 39% year-to-year and 8.8% month-to-month.

— The Semiconductor Industry Association

Personal Computing

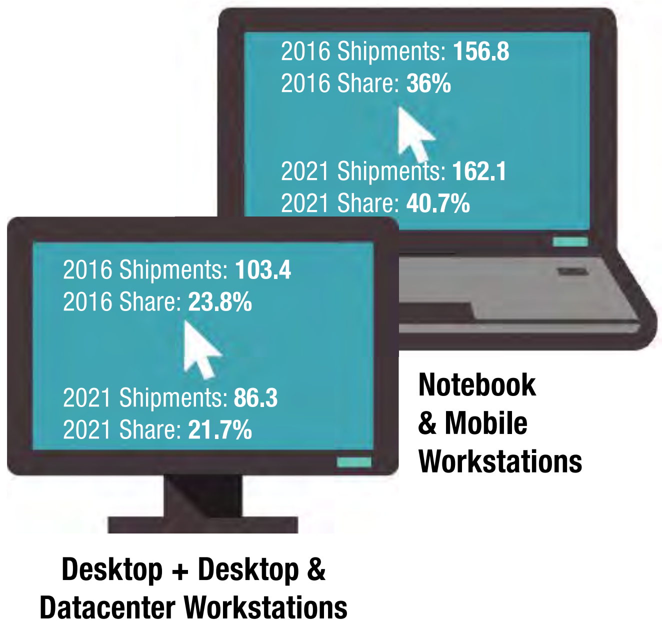

Personal computing device forecast, 2016-2021 (shipments in millions). The commercial segment is forecast to stabilize in 2017 and show growth in 2019 and beyond.

— IDC Worldwide Quarterly Personal Computing Device Tracker, Aug. 24, 2017

CAD in the Cloud

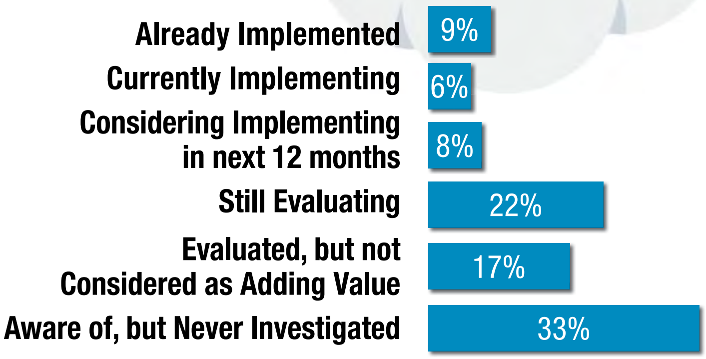

CAD in the cloud was defined as a "model where software applications are delivered to the user over the internet i.e. software is not stored on the company's network, desktop PC, laptop, etc."

Who Do You Know?

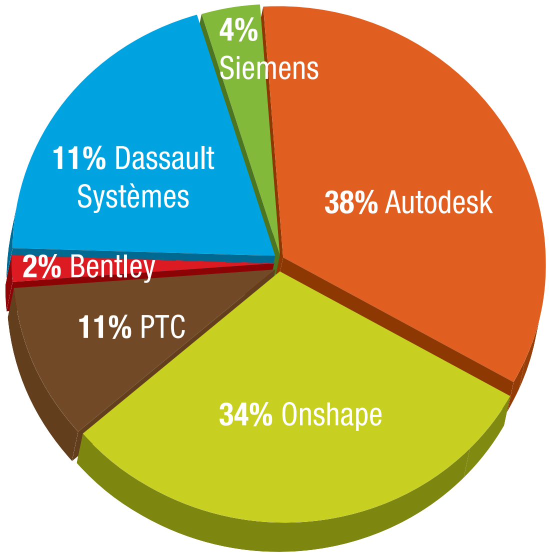

Top-of-mind awareness of cloud-based CAD among manufacturing respondents.

— *"CAD in the Cloud: The CiC Market Trends 2017 report," September 2017, Jon Peddie Research and Business Advantage Group*

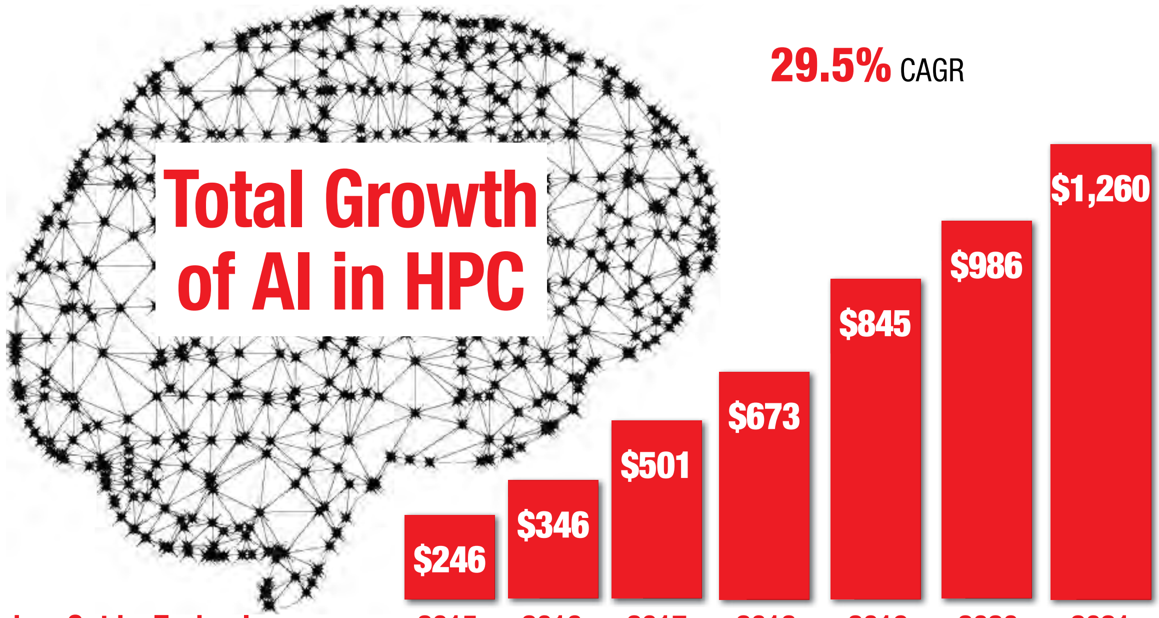

Broken Out by Technology	2015	2016	2017	2018	2019	2020	2021	CAGR 16-21
Machine Learning in HPC	$203	$282	$373	$478	$548	$546	$594	16.0%
Deep Learning in HPC	$15	$31	$93	$159	$258	$395	$618	81.5%
Other Artificial Intelligence In HPC	$28	$32	$34	$36	$39	$44	$48	8.4%

Worldwide high-performance computing based revenues, in millions, for machine learning, deep learning and artificial intelligence show a total compound annual growth rate (CAGR) of 29.5% between 2016 and 2021.

— *Hyperion Research, 2017*

Siemens Doubles Down on Digitalization with Acquisitions

BY JAMIE J. GOOCH

MANY TECHNOLOGY vendors have been working phrases like "industrial internet of things," "digital twin," "digital thread" and "digital disruption" into their presentations for years now. It can be tough to know which are paying lip service to buzzwords and which are truly committed to the ideas of connecting digital designs and physical products up and down the lifecycle. One way to separate the hype from reality is to follow the money.

At the Siemens Industry Analyst Conference in Boston earlier this month, Siemens PLM Software President and CEO Tony Hemmelgarn said the company had invested $10 billion in digitalization efforts over the past 10 years. A big chunk of that change came from its acquisition of Mentor Graphics that closed this year, for what amounted to an enterprise value of roughly $4.5 billion. That was after last year's acquisition of CD-adapco for $970 million. Just before the conference, Siemens announced it is acquiring TASS International, which will boost its autonomous driving simulation and test capabilities, and just last week announced it is acquiring Infolytica, expanding its simulation suite into electromagnetics.

Siemens PLM Software's leadership on stage at the company's 2017 Industry Analyst Conference.

"The integrations are going extremely well," said Hemmelgarn. "CD-adapco, for example, has just been a great merger, bringing the team together to serve our customer base."

One result of that teamwork is Siemens PLM Software's Simcenter portfolio that was announced last summer. It combines system simulation, 3D CAE and test applications. CD-adapco products are part of the Siemens Simcenter Portfolio, which a company blog post called "a suite of simulation and test solutions for predictive engineering analytics (PEA). PEA is the application of multidisciplinary simulation and test, combined with design exploration and data analytics, to create a true-to-life digital twin of products."

The Race for a More Identical Digital Twin

The "digital twin" term is still fairly new. It's been around in the aerospace industry since at least 2012, but gained steam when it took center stage at PTC and Dassault Systèmes conferences in 2015. Siemens, PTC and Dassault Systèmes all have different approaches to the digital twin and everything the term entails. From Siemens' viewpoint, simulating an entire system, all along the product lifecycle, enhances the value of a virtual representation of a physical product.

MORE → digitaleng.news/virtual_desktop/?p=13258

Mentor Prepares Engineers for an Electric, Self-Driving Future

Mentor CEO Wally Rhines.

"Whenever there is a major technical discontinuity, there is opportunity for new entrants and there is risk for the establishment," said Mentor CEO Wally Rhines from the stage of the Integrated Electrical Solutions Forum (IESF) Automotive, on September 20 in Plymouth, MI. The auto industry is indeed in the midst of "discontinuity," thanks to consumer demand for what amounts to a connected computer on wheels, new electric drivetrains and the rapid advancement of assisted to autonomous driving. But Mentor, too, is looking for opportunity in disruption as it integrates with Siemens PLM Software, which acquired the company to extend its digitalization efforts.

To make his point, Rhines looked back at the beginning of the auto industry, when almost 300 companies rushed to take advantage of the disruption in the transportation market and compared it to the 300 companies now developing electric cars and trucks, and the almost 100 companies that have announced autonomous driving programs. Rhines' point was clear: Only a few companies will survive the current disruption, and those with the most efficient electronic engineering tools will have a distinct advantage.

MORE → digitaleng.news/virtual_desktop/?p=13276

Autodesk and America Makes Hold Event to Promote Design for Additive Manufacturing

BY KENNETH WONG

TRX+, A TWO-DAY EVENT organized by America Makes and Autodesk, included keynote talks by Terry Wohlers from Wohlers Associates on "The Future of Design for Additive Manufacturing and 3D Printing," and Erin Bradner from Autodesk on "Generative Design: Realizing the Future of Design for Additive Manufacturing."

The void the organizers are hoping to fill is the lack of education in design for additive manufacturing (AM).

3D Printing Knowledge Hunger

America Makes is a public-private membership organization that focuses on "advancing the additive manufacturing/3D printing industry," with "a portfolio of more than $100 million in public and private funds," according to its home page. It's managed by the National Center for Defense Manufacturing and Machining (NCDMM). Autodesk is a platinum member of America Makes.

Autodesk also belongs to the Digital Manufacturing and Design Innovation Institute (DMDII), a sign that the company is betting heavily on AM becoming an integral part of design and manufacturing.

"Technology maturity from academia to production environment takes many years, with lots of hurdles to overcome," explained Patti Vrobel, research engineer, Autodesk Research. "America Makes funds projects that will accelerate the adoption of AM. The other part is education, of course. We need to start training people to use these tools."

America Makes offers AM-related training sessions in the classroom and online. In-person classes are usually held in the organization's Youngstown Business Incubator office in Ohio.

AM still has many gaps to fill to reach maturity. Part of the evolution may be the development of AM-aware simulation tools and design programs.

"We know the way the filament are deposited affect the [3D-printed] part's quality, but it's very difficult to tell that to the finite element analysis (FEA) software," noted Dr. David Rosen of Georgia Tech, during his presentation on "A Design Guidance System for AM." He added, "We have to make the design tools AM-aware, and we should do it with the tools the designers are already using."

Autodesk Netfabb software allows you to simulate and analyze complex 3D printing projects. ***Image courtesy of Autodesk.***

MORE → digitaleng.news/virtual_desktop/?p=13301

Defining AI: The Ability to Evolve is Essential

At the World Mobile Congress America in San Francisco, speaking on the panel titled "The Future of AI," Gunnar Carlsson described himself as "a recovering academic." A retired math professor, he's now the president of enterprise AI technology supplier Ayasdi, based in Menlo Park, CA.

"AI isn't all about the sexy stuff that one hears about, like self-driving cars. There are some simple, pragmatic things one can do with it," Carlsson pointed out. "We worked with a global customer to solve a money laundering problem, to lower false positives dramatically."

Somar Velayutham, the artificial intelligence evangelist for NVIDIA, offered another example of seemingly mundane but practical use of AI. "At NVIDIA cafe, we use an image-recognition system. So if you just leave your soup, salad and lunch in front of it, it tells you what the calories are and what the price is," he added.

Recently in the engineering sector, design and engineering tool providers like Autodesk and Dassault Systèmes have begun exploring ways to incorporate AI into their 3D modeling and simulation products. Leading microprocessor makers—Intel, NVIDIA and AMD—also stand to benefit from the growth in AI, as the process of developing AI-driven systems requires tremendous processing power.

The moderator, Paul Hsiao from Canvas Ventures, posed a question to the panelists: What is AI, and what is not AI?

MORE → digitaleng.news/virtual_desktop/?p=13314

BRIEFINGS

News and New Products

Infosys, Renishaw Forge Partnership

The two companies report that they are combining their engineering expertise and global resources to help accelerate deployment of AM for volume production of end-use metal components. Infosys will apply its engineering processes and design for additive manufacturing knowledge to manage product development projects from concept to launch. Renishaw will support Infosys through its global network of Additive Manufacturing Solutions Centers, which provide access to Renishaw metal AM technology, backed by application engineering expertise, post-processing capability and metrology, Renishaw reports.

MORE ➜ digitaleng.news/de/?p=39658

Autodesk Toolkit Takes Pain Out of IoT Discovery

Autodesk, in partnership with Electric Imp, an internet of things (IoT) platform focused on secure connectivity to devices, has released the IoT Discovery Toolkit designed to take some of the sting out of early prototyping. The packaged solution includes a Wi-Fi-enabled gateway, grove connector cable, Autodesk Fusion Connect account with predefined business application templates and an Electric Imp developer account with a predefined IDE framework, enabling devices and cloud applications to be easily managed.

By teaming up on an IoT toolkit solution, Autodesk and Electric Imp are providing key pieces to facilitate and simplify IoT product discovery and implementation in a single purchase, supported through resources by both partners. Many companies starting down the path to IoT get stuck in the concept and prototyping phase as they struggle with long lead times to figure out how to build out and secure the connectivity stack from the device all the way up to the cloud.

MORE ➜ digitaleng.news/virtual_desktop/?p=13246

Project Management Tool Takes Aim at Complexity

Managing projects is a big pain in engineering circles, yet few turn to traditional project management software to get the job done.

LiquidPlanner, a maker of what it calls "dynamic" project management software, is teaming up with P3 Group, a global manufacturing consultancy, to try to change that scenario. Citing what it calls a "tidal wave" of complexity in light of Industry 4.0 trends like the internet of things, 3D printing and mass customization, LiquidPlanner says engineering, product development and manufacturing teams need a mechanism to

navigate relationships between people, priorities and effort. This set of dynamics, officials say, is just not possible to manage efficiently using standard project management tools, let alone the still popular spreadsheet.

MORE ➜ digitaleng.news/virtual_desktop/?p=13252

SolidWorks 2018 Debut Puts Spotlight on Collaboration

Dassault Systèmes SOLIDWORKS Corp. has launched the 2018 edition of its SolidWorks portfolio of 3D design and engineering applications.

With this release, Dassault Systèmes says, teams can collaborate concurrently to rapidly and cost-efficiently design a product or part, validate its function and manufacturability, manage its data and related processes, streamline and automate its fabrication as well as inspect it. With SolidWorks 2018, any changes in design or manufacturing are reportedly easy to manage and automatically flow to all related models, programs, drawings and documentation due to intellectual property embedded early on in the design process.

Powered by Dassault's 3DEXPERIENCE platform for engineering, the company says SolidWorks 2018 supports a complete design through manufacturing strategy with solutions that simplify the interactions between disciplines across the product development workflow.

MORE ➜ digitaleng.news/de/?p=39876

Altair Extends Strategic Relationship with HPE

Altair has entered into a multi-year original equipment manufacturing agreement with HPE. This agreement represents an expansion of the long-term partnership between HPE and SGI (who HPE acquired). HPE will now be able to include Altair's PBS Professional workload manager and job scheduler on all of HPE's high performance computing (HPC) systems.

PBS Professional gives HPE cluster users a solution for HPC workload management. As an HPE-integrated product, PBS Professional optimizes job scheduling on HPE Apollo and HPE SGI servers to achieve the highest levels of system use. PBS Professional is also integrated with HPE's HPC system management solutions: HPE Insight Cluster Management Utility for (CMU) for HPE Apollo and HPE ProLiant platforms as well as HPE SGI Management Suite for HPE SGI 8600 systems.

MORE ➜ digitaleng.news/de/?p=39482

➜ For more coverage of new products for design engineering teams, visit digitaleng.news/de/products csimsoft bolt2™

BECAUSE MESHING IS HARD

TAKE THE CHAOS OUT OF SIMULATION

MODEL PREP
WE'VE ELIMINATED IT.

MESHING TIME
WE'VE COLLAPSED IT.

COMPLEXITY
WE'VE SIMPLIFIED IT.

PROVE IT TO YOURSELF

CSIMSOFT.COM/CHAOS

| CONSULTANT'S CORNER |

COMPUTING

by Chad Jackson

Engineering's Compute Power Appetite is Insatiable

FOR YEARS, ENGINEERS NEEDED MORE and more compute power. More memory and compute power has been needed to open and manipulate huge 3D assemblies in CAD applications. More compute power has been required to solve increasingly large finite element analysis (FEA) models with thousands, or even millions, of degrees of freedom. Over the span of decades, the need for more compute power continuously grows. That, of course, has been driven by software providers exploring the very edge of what can be done.

Today the need for more compute power is no different: More is needed. However, engineering organizations are facing some new questions: Do you put more power in the desktop or buy more capability in the cloud? In some cases, it is not an either-or type of answer; it's both. To really figure out the right answer for your organization, figure out your software capability needs. Let's run down the options.

Automated Simulation

Simulation-driven design, the idea of conducting analyses to make more informed design decisions, has been a longstanding vision in engineering organizations. In the past few years, a new technological paradigm has been developed to enable that pursuit. The idea is to create a "templatized" simulation that only needs a few inputs. Non-expert users can plug their designs into that template, tagging the right inputs for their situation, and press the "go" button. Some time later, the simulation is complete and the results can be reviewed. This process can be repeated, and even automated, to assess many different designs as a single batch.

From a compute perspective, heavy use of this approach will increase demand for more powerful computing resources. However, automated simulation capabilities can come in simulation applications that run on the desktop, can leverage compute farms or can operate in browser-based apps that run in the cloud. Investing in more powerful desktops or compute farms comes as a one-time capital expense. Supporting more compute power for cloud apps comes as an ongoing operational expense.

Instant Simulation Results

An intriguing new simulation capability has emerged in some analysis software applications: instant simulation results. The concept is that as soon as a non-expert identifies or supplies the minimum information needed to run the simulation, the software starts showing the lowest resolution result. As the software continues to refine the results, it is shown as an animation on the screen. As geometry is changed, removed or added to the design, the results update. It allows users to see the impact of changes almost in real time. This approach provides some instant gratification for engineers who want to play what-if games with their design.

Although this new approach leverages some significant advances in analysis software, it also has some unique computing requirements. It uses graphics processing units in workstations to run dramatically more iterations of simulations.

Big Data in the IoT Era

No discussion on modern product development would be complete without the inclusion of the internet of things (IoT). Streaming data from sensors and software off of products allows for more insight into how products operate. That, in turn, provides engineers with information to design better offerings in the future. It also allows them to feed that data back into simulation models, effectively creating a digital twin of the physical product.

The IoT carries some unusual compute demands. The data that streams off products can pile up quickly. Soon enough, organizations find themselves with big data, a set of information so large it is hard to even host it, much less analyze it. Desktop resources need to be dramatically improved if local software applications are used. Alternatively, some organizations employ cloud-based analytics tools that reside next to the data. Organizations must plan carefully when considering these options, as many frequently underestimate their needs.

Takeaways

Engineering will always need more compute power. However, the fundamental questions about computer resources are changing: Do you invest in desktop-based or cloud-based assets? The key to making the right choice is awareness of the tools the engineering organization is planning to use. With that, you have the right context to make the right choice. DE

Chad Jackson *is president of Lifecycle Insights (lifecycleinsights.com). Send email about this commentary to de-editors@digitaleng.news.*

ELEMENTS

by Tony Abbey

Who Are You Calling Dumb?

I LIKE USING ANALOGIES. When it comes to finite element analysis (FEA) elements, think of them as players on your football team who have a certain level of intelligence.

3D continuum elements are what I call "smart." They understand what is happening inside their volume. Mathematically, this is through their shape functions and internal displacement responses. However, there are degrees of smartness; parabolic second-order tetrahedral elements and first-order brick elements have an acceptably high IQ. However, distort those elements and they will produce poor results. These are team members who are not playing to their full potential. They need coaching to get organized into a better mesh distribution. First-order tetrahedral elements have an unacceptably low IQ. You do not want them on your team.

A CAD design, meshed properly with 3D solid elements, will create an FEA model with a well-defined response. The intelligence of the model is really helping us out. We do not have to do so much thinking about how the elements represent the displacement or stress distribution. But as the coach, we are going to be critical about preparation and performance.

FEA solvers do not know anything about CAD geometry; they understand nodes and elements. So, the raw input file describes the 3D geometry via these smart 3D spatial elements.

Now we come to the 2D shell elements and 1D beam elements. I invariably describe these as "dumb" when describing their characteristics.

So why is a 2D shell element dumb? It has no concept of the world outside its datum mid-plane surface. I cannot directly mesh a thin shell type CAD model with 2D elements. Techniques, such as mid-surfacing, explain to the poor old shell element where its mid-plane is. We even need to tell it what thickness it has and which direction is up. Smart preprocessing can associate the CAD geometry thickness to the element, as an FEA physical property. But even then, we need to check accuracy.

The input file created for the FEA solver describes the design component as 2D surfaces mapped into 3D space, with a parametric thickness. We, the analyst, work harder to idealize the structure so that this representation will work adequately.

The 2D shell element guesses what is happening through thickness, by extrapolating from the mid-plane. Pure in-plane load assumes a constant stress variation through thickness. That's OK for in-plane stresses, but shear stress through thickness is ignored for a plain stress element. Bending uses a linear variation of direct stress and a parabolic distribution of shear stress through thickness.

The B Team

So why have these dumb elements on our team? By now they should have been eliminated because of the limitations that we have seen. We can indeed view them as the second team, and they will struggle to perform well in a general, chunky, 3D component. However, compared with our star players, the 3D solids, they are dirt cheap. If the conditions are right, and we have thin shell-type design components, then our "poor performers" really start to shine. They have limited horizons, and will not overthink the problem, as their 3D colleagues would do in this case. Their approach is "I've been told it's a shell, so I will model as a shell." This is where our 3D superstar gets a little out of his depth, to excuse the pun. His approach is: "Space is space. What is a shell?"

So, selecting 3D solids or 2D shells is rather like picking the right team, for the right conditions. My professional football analogy is rather suspect, but there must be a comparison to pitching your multi-million-dollar star players against a scratch college team, on a muddy field with 300 spectators present. Being a Brit, maybe I should stick to rugby at this point.

So, where does that leave the 1D beam elements? They can only handle 1D concepts. They understand what happens along their axis, but have no clue about response in the cross-section. They follow simple rules to take a guess at what is going on out there. These guys are the bottom of the pay grade. However, we can afford them in large numbers; it will not put the smallest dent in our wallet. They also play to their strengths. The approach is familiar: "I've been told it's a beam, and I'm going to handle it like a beam." Their report back to you, the analyst, is straightforward. If we set the superstars onto this type of structure, we get a large salary bill, and every one of them tells a complicated story.

There we have it: another analogy. Use it to pick your team, but remember not to insult the players too much! **DE**

Tony Abbey *is a consultant analyst with his own company, FETraining. He also works as training manager for NAFEMS, responsible for developing and implementing training classes, including a wide range of e-learning classes. Contact tony.abbey@nafems.org for details.*

The Heart of the Engineering Workflow

New devices, new workstation classes and HPC-incorporated workloads expand the playing field for CPUs and GPUs.

BY KENNETH WONG

FIVE OR SIX YEARS AGO, engineering workflow was driven primarily by the workstation. Most engineers relied on the CAD-certified workstation as their primary hardware to create 2D layouts and 3D models. To the extent possible, they ran stress analysis and simulation on the same machines. And they harvested the GPU (graphics processing unit) in the machine to create photorealistic renderings. But advances in more accessible high-performance computing (HPC) and more powerful mobile computing options have changed the engineering workflow.

With the emergence of mobile workstations and even engineering-capable tablets, on-demand cloud services and on-premise HPC, the engineers' workflow now extends beyond the traditional desktop workstation. They may perform design reviews, collaborative discussions and light engineering on mobile devices. To speed up their complex fluid flow simulations and ray-traced rendering jobs, they may augment the power in their workstations with on-demand computing. Some may even move to virtual machines as their primary systems, using lightweight tablets or "dumb clients" as access points instead.

The switch is also reflected in the type of support now available from leading design software vendors. SolidWorks, one of the widely adopted CAD programs, is now certified for Microsoft Surface Pro, a mobile tablet. As part of Autodesk's subscription services, many Autodesk software products let you tap into on-demand computing services to speed up simulation and visualization tasks. This article focuses on, quite literally, the core of the new working paradigm. It looks at the role of the different microprocessors in this multisystem workflow.

Tiny Desktops, Mobile Tablets

Intel's processors fall into two distinct lines: the Intel Core i, often found in consumer and gaming PCs; and the Intel Xeon, designed for professional workstations and data center products. Generally, the design and engineering workflow demands hardware with workstation-class performance; therefore, CAD, simulation and visualization software users from the engineering community may naturally gravitate toward workstations with Intel Xeon CPUs.

However, under the new usage paradigms, some CAD-centric operations are now possible on mobile tablets, a market

NVIDIA's Quadro GPUs are the company's offerings for the professional workstation line. *Image courtesy of NVIDIA.*

The emergence of mini and tiny workstations, like the Lenovo P320 Tiny shown here, blurs the line between what is a desktop workstation and a mobile workstation. *Image courtesy of Lenovo.*

untouched by the Intel Xeon processor line. "Intel Xeon will not likely be found in a mobile, wireless, handheld tablet. Intel Core i is the CPU tailored for these devices," clarifies Steve Gabriel, Intel public relations.

Also emerging are the new, tiny or mini-class workstations, such as Lenovo ThinkStation P320 Tiny or HP Z2 Mini. In size and scope, they usually measure no more than a hardcover book, challenging the conventional classification of workstations. They're technically desktop workstations, in that they demand an external monitor, a keyboard and a mouse to operate. Yet, they're small enough to fit inside a lunchbox, a handbag or a briefcase—so they're also mobile. Because of the need to keep the power demand low, these new classes of workstations currently feature only Intel Core i CPUs.

"The ThinkStation P320 Tiny is currently available only with Core i CPUs because of the size limitation and the thermal envelope we have to work with," says Scott Rupert, Lenovo's workstation portfolio manager.

It's important to note that the Tiny and Mini workstations belong to the entry-level workstation class. These systems are designed for the conceptual design phase, with limited simulation and rendering. For this usage, Core i CPUs prove to be adequate. Those who routinely perform compute-intense simulation and visualization tasks may find them underpowered.

In March, BOXX launched the APEXX 1 1402, described as "the world's smallest workstation featuring an overclocked Intel Core i7 Kaby Lake processor." Overclocking— increasing the processor's capacity beyond the chip manufacturer's default setup—can be risky when done without proper safeguards. BOXX is one of the system suppliers that offers safely overclocked systems that give more performance from the processors.

Measuring 4.7x8.5x9.0 in., the APEXX 1 1402 features a four-core Intel Core i7-7700K processor overclocked to 4.7GHz. BOXX says it is "built to accelerate 3D modeling and CAD design workflows" and recommends it for professional applications such as SolidWorks, Autodesk Revit and 3D Studio Max.

Entry-Level and Professional Workstations

In 2012, HP released its first entry-level workstation, the HP Z210. Dell also released the Dell Precision T1600, later replaced by the T1650. These were part of the system vendors' efforts to nudge professional users away from top-of-the-line consumer PCs. With a memory boost and GPU upgrade, a high-end consumer PC could run engineering software; however, the risk of running such software on uncertified systems is significant, and system stability is not assured. With CAD-certified entry-level workstations priced around $1,000, system vendors hoped to convince budget-conscious buyers that they did not have to settle for consumer-class PCs.

"The Intel Xeon E3 processors target entry-level workstations with a compelling price-to-performance, for lower core-count use cases where system expandability is not a priority," explains Gabriel.

The professional workstation space, encompassing standard and high-end engineering workstations, is well-defined. They're usually configured with Intel Xeon W CPUs. According to Gabriel, the Xeon W CPU family is "a cost-optimized, one-socket platform with a core-count ideal for mainstream workstation and high frequency providing support for more memory and storage at a small footprint." He points out: "For users doing CAD, modeling, product design and VR (virtual reality), this is a good entry-level platform for those usages. Xeon Scalable Processor takes it to the next level with support for up to 56 cores, 3 TB memory and support for professional-grade graphics adapters, making it an ultimate content creation platform for VR."

A new challenge may be coming from AMD, a rival of Intel in the consumer and gaming PC market. In August, at the computer graphics conference SIGGRAPH, AMD introduced AMD Ryzen Pro, described as "workstation-class performance for premium desktops." According to AMD, it is "the first processor to offer up to eight cores for commercial-grade PCs, 16 threads and enables up to 62% more multithreaded performance on the Ryzen 7 PRO 1700 than other solutions."

AMD recently introduced the Radeon Pro SSG, a professional-class GPU based on its VEGA architecture. *Image courtesy of AMD.*

Engineering VR-ready features into microprocessors has become part of chip makers' competitive strategies. *Image from CES 2017 by Intel.*

Professional desktop and mobile workstations usually come equipped with GPUs, an essential feature that allows 3D CAD, rendering and simulation software users to tap into the GPU's parallel architecture for rendering, visualization and computing tasks. NVIDIA's offering for this segment is the NVIDIA Quadro GPU line. Select products from the Quadro family are branded VR-ready, indicating they are configured for VR content development and deployment.

Rival AMD's offerings in this segment are the AMD FirePro and Radeon Pro GPU families. The FirePro line is known for, among other things, support for multimonitor setup, a useful feature for those who want to deploy video-wall displays. AMD launched Radeon Pro WX 9100 and Radeon Pro SSG at the recent SIGGRAPH conference. Based on AMD's new Vega GPU architecture, these cards feature high bandwidth cache controller (HBCC), which the company says will allow "data to seamlessly move between onboard graphics memory and available system memory."

QUICK REFERENCE CHART

Entry-level workstations

- Intel Xeon E3 CPUs
- AMD Ryzen CPUs
- Radeon Pro WX GPUs (Radeon Pro WX 2100, WX 3100)
- NVIDIA Quadro GPUs

Professional workstations (mainstream to high end)

- Intel Xeon W CPUs (marketed for VR content creation)
- Intel Xeon Scalable CPUs
- AMD Ryzen Pro CPUs
- AMD FirePro GPUs
- AMD Radeon Pro WX 4100 to WX 9100, Radeon Pro Duo and SSG GPUs
- NVIDIA Quadro and Tesla GPUs

HPC, data center, virtualization

- Intel Xeon Scalable CPUs
- AMD EPYC CPUs
- AMD OPTERON CPUs
- RADEON Instinct GPU (especially for machine learning, AI, available Q4, 2017)
- NVIDIA TESLA GPUs (especially for machine learning, AI)

HPC, Data Center

The HPC space includes on-demand servers, on-premise servers and various forms of cloud computing. Workstation users who find the need to borrow additional computing capacity may now tap into any of these options, delivered from the browser or an enterprise network architecture. Although HPC is seldom needed by those dealing with conceptual design, it is routinely used by those using compute-intensive simulation and visualization programs.

In July, Intel launched what it described as "the new foundation for secure, agile, multi-cloud data centers." The chip maker writes, "The [Intel Xeon Scalable] processors deliver exceptional workload-optimized performance and hardware-enhanced security. Designed for trusted data service delivery, the processors are fueled by significant leaps in I/O, memory, storage and network technologies." The processors are said to offer a performance increase particularly in "modeling and simulation, machine learning, HPC and digital content creation."

AMD's offering for this space is AMD EPYC processors, launched in June of this year. EPYC includes "the first embedded x86 silicon-level data security on a server chip," according to AMD. Describing EPYC's single-socket advantage, AMD says, "Many IT organizations purchase two-socket servers and only populate one socket. Others purchase two-socket servers, not because they need the compute capability, but because they need more I/O and or memory capacity than what is available on current single-socket servers. AMD EPYC enables no-compromise one-socket servers with up to 32 cores, eight memory channels and 128 PCIe 3.0 lanes, enabling capabilities and performance previously available only in two-socket architectures."

Both AMD and its rival NVIDIA offer GPU-accelerated HPC products through their hardware partners. With researchers delving into highly parallel computing workflows (such as machine learning and big data analysis), the GPU's massively parallel architecture proves to be advantageous.

The NVIDIA Tesla GPU line is usually the centerpiece of NVIDIA's GPU-accelerated data center products. "Data scientists and researchers can now parse petabytes of data orders of magnitude faster than they could using traditional CPUs, in applications ranging from energy exploration to deep learning. Tesla accelerators also deliver the horsepower needed to run bigger simulations faster than ever before," NVIDIA writes.

The current HPC gold rush is the race to capture the emerging machine learning or artificial intelligence (AI) development, driven in a large part by the auto industry's pursuit of autonomous vehicles. AMD's new GPU, Radeon Instinct, is a sign of this. "The new Radeon Instinct MI25 accelerator, based on AMD's Next-Gen Vega architecture, with its powerful parallel compute engine, is the world's ultimate training accelerator for large-scale deep learning applications and is a workhorse for HPC workloads," AMD writes. It promises up to "24.6 TFLOPS FP16 or 12.3 TFLOPS FP32 peak GPU compute performance on a single board." Radeon Instinct is expected to be available in Q4. *(Editor's note: For more on acceleraotrs, see page 18.)*

The Blurred Lines

The GPU was originally a graphics-boosting coprocessor. About five years ago, when NVIDIA decided to refashion its GPU as a general-purpose computing device, it launched the GPU into many markets traditionally served by the CPU. As system makers introduce classes of computing platforms—tablet-style devices powerful enough to run CAD, and tiny workstations the size of a takeout lunchbox, to name but two—it gets harder to define what constitutes a "workstation." In developing smaller desktops, entry-level systems, and mobile tablets to attract new users, system vendors are now using the processor families designated to the consumer space, such as Intel Core i CPUs, for some of their offerings targeting the professional market. The changing environment is reflected at the silicon level, in the features engineered into the chips that power these new devices. **DE**

Kenneth Wong *is DE's resident blogger and senior editor. Email him at de-editors@digitaleng.news or share your thoughts on this article at digitaleng.news/facebook.*

INFO ➜ AMD: AMD.com

➜ **HP:** HP.com

➜ **Intel:** Intel.com

➜ **Lenovo:** Lenovo.com

➜ **NVIDIA:** NVIDIA.com

For more information on this topic, visit digitaleng.news.

Accelerate Data Processing

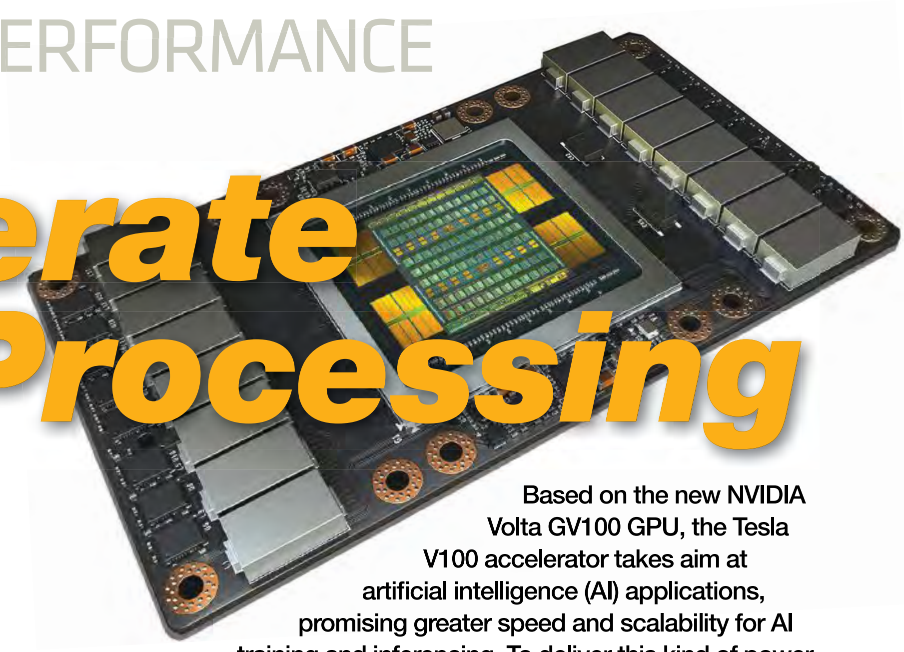

Based on the new NVIDIA Volta GV100 GPU, the Tesla V100 accelerator takes aim at artificial intelligence (AI) applications, promising greater speed and scalability for AI training and inferencing. To deliver this kind of power, the accelerator leverages an impressive capacity for parallel processing. *Image courtesy of NVIDIA.*

Machine and deep learning applications call for more parallel processing.

BY TOM KEVAN

DESIGN ENGINEERS' INCREASING RELIANCE on accelerators to enhance system performance has begun to alter the basic concepts of how devices channel and process data. Many leading-edge devices entering today's market incorporate technologies like high-resolution imagery and video, machine learning and virtual and augmented reality. Faced with demand for products that require rapid processing of huge amounts of data while consuming minimal energy, designers now turn to programmable logic chips that optimize processing workloads. This trend has been further driven by CMOS scaling's inability to keep pace with performance demands. As a result, accelerators—such as graphics processing units (GPUs), field-programmable gate arrays (FPGAs), digital signal processors (DSPs,) machine-learning accelerators and heterogeneous CPU cores that can perform tasks in parallel—have begun to take on new, high-profile roles in electronics designs.

Chipmakers often link the deployment of accelerators with efficiency and performance requirements. "Acceleration is generally chosen because it either improves energy efficiency (performance/W) or performance density (performance/mm^2)—often both," says Jem Davies, vice president, general manager and fellow at ARM.

At the core of these efficiency/performance issues lies a different type of compute workload. Applications like machine and deep learning and augmented and virtual reality require parallel processing hardware that can efficiently handle huge data sets, a feature that existing CPUs often cannot provide by themselves.

The Advantages

To better understand what accelerators bring to the table, compare their features with those of traditional CPUs. Typically, CPU cores are optimized for sequential serial processing. Processor manufacturers have pushed the performance of CPUs nearly as far as they can using conventional techniques, such as increasing clock speeds and straight-line instruction throughput.

Accelerators, on the other hand, leverage parallel architectures designed to handle repetitive tasks simultaneously. The primary benefit of accelerators is that they can offload and quickly execute compute-intensive portions of an application while the remainder of the code runs on the CPU.

"In the case of CPUs and GPUs, the contrast is large. A datacenter CPU can have eight to 24 cores," says Robert Ober, chief platform architect for datacenter products at NVIDIA. "A datacenter GPU has more than 5,000 cores that have the same order of floating-point capability. The result is orders of magnitude more usable capability and throughput from a GPU for these tasks. The result is more efficiency—fewer steps for the same work, and less power and energy spent on that work."

CPU giant Intel hasn't overlooked the need for acceleration and parallelism. When the company introduced its Xeon Phi processor family, it described it as enabling "machines to rapidly learn without being explicitly programmed, in addition to helping drive new breakthroughs using high performance modeling and simulation, visualization and data analytics."

Just last month, Intel launched a hardware and software platform solution to enable faster deployment of customized FPGA-based acceleration of networking, storage and computing workloads. In a press release, the company introduced the Intel Programmable Acceleration Card with the Intel Arria 10 GX FPGA enabled by the acceleration stack as the first in a family of Intel Programmable Acceleration Cards. It is expected to be broadly available in the first half of 2018. The platform approach enables original

equipment manufacturers to offer Intel Xeon processor-based server acceleration solutions.

A New Design Perspective

As engineers increasingly use accelerators to extract value from big data, design teams have to rethink the flow of data within the systems they are designing and reconsider where processing should occur. In doing so, they fundamentally change conventional design philosophy. For years, engineers focused on reducing energy consumption, turning off many of the cores in a chip to minimize power consumption. With changing consumer demands, many engineers are shifting their attention to improving performance. Rather than processing all functions in a single CPU, they have begun to use multiple heterogeneous types of processors, or cores, with specialized functionality.

"Machine learning workloads have very specific data flow patterns and computational requirements," says Davies. "By tailoring designs toward these workloads, great gains can be achieved in both energy efficiency and performance density. Because of the importance of machine learning, we expect all computing platforms—CPUs, GPUs and machine learning accelerators—to evolve to meet emerging machine learning requirements."

Proof of this approach can be seen in recent innovations. Data scientists have been using accelerators to make groundbreaking improvements in applications such as image classification, video analytics, speech recognition and natural language processing.

One Size Does Not Fit All

These advances, however, have not been enabled by just one type of accelerator. The implementation of machine learning and other specialized applications requires engineers to leverage a variety of accelerators. Although it's true that all accelerators improve performance, no one size fits all. Most of these processors involve some form of customization.

For example, GPUs perform well when accelerating algorithms in the learning phase of machine learning because they can run floating point calculations in parallel across many cores. But in the inference phase of machine learning, engineers benefit from using FPGA and DSP accelerators because these processors excel at fixed-point calculations.

Some companies have even gone so far as to create their own customized accelerators for machine learning. "You've probably seen the first of these new architectures, the tensor processing unit [TPU] recently announced by Google as their proprietary custom accelerator for machine inference," says Nigel Toon, CEO of Graphcore. "Startup Nervana—recently acquired by Intel—also claims they are working on a TPU. It's especially exciting to see Google advocating tailored processor design for machine learning."

Graphcore itself has introduced a custom accelerator for machine intelligence applications. "If we think of the central processing unit in your laptop as being designed for scalar-centric control tasks and the GPU as being designed for vector-centric graphics tasks, then this new class of processor would be an intelligence processing unit [IPU], designed for graph-centric intelligence tasks," says Toon. "When we started thinking about building a machine to accelerate intelligence processing at Graphcore, we knew that we had to look beyond today's deep neural networks … Our IPU has to outperform GPUs and CPUs at all these tasks. But perhaps more importantly, it has to provide a flexible platform for the discoveries yet to come."

One of the latest entries in this arena is Inuitive's new multi-core image processor called the NU4000. This chip supports 3D imaging, deep learning and computer vision processing for applications such as augmented and virtual reality, drones and robots.

Although these technologies give design engineers new options for enhancing advanced systems, they also introduce additional challenges into the design process.

Challenges

At the nuts-and-bolts level, the inclusion of accelerators in designs forces engineers to make tough tradeoffs in performance and flexibility. "The gains in energy efficiency and performance density come from targeting your design toward a given workload or set of workloads. This, in turn, will reduce the flexibility and sometimes the programmability of a system," says Davies.

These tradeoffs, as well as the increased emphasis on performance, also force designers to take a broader view. According to Davies: "In the past, an SoC designer may have focused primarily on a single benchmark or set of benchmarks. We're seeing the focus switch from benchmarks to use cases, where the use case may consist of multiple different workloads, running on a combination of accelerators, working together in unison."

Adding these new factors to the mix will make it essential that engineers adopt a new set of technology-specific tools. Not surprisingly, processor providers have already moved to meet this need.

New Architectures and Development Tools

Because of this, design engineers will find a broad assortment of tools and development environments tailored for the inclusion of accelerators in designs. Many of these go one step further and provide the means to use accelerators to develop leading-edge applications like machine intelligence.

NVIDIA offers its CUDA Toolkit. This development environment targets GPU-accelerated applications, providing libraries, debugging and optimization tools, a C/C++ compiler and a runtime library to deploy applications.

CUDA libraries promise to enable drop-in acceleration across multiple domains, such as linear algebra, image and video processing, deep learning and graph analytics. Using built-in capabilities for distributing computations across multi-GPU configurations, engineers can develop applications that scale across a variety of platforms, from single

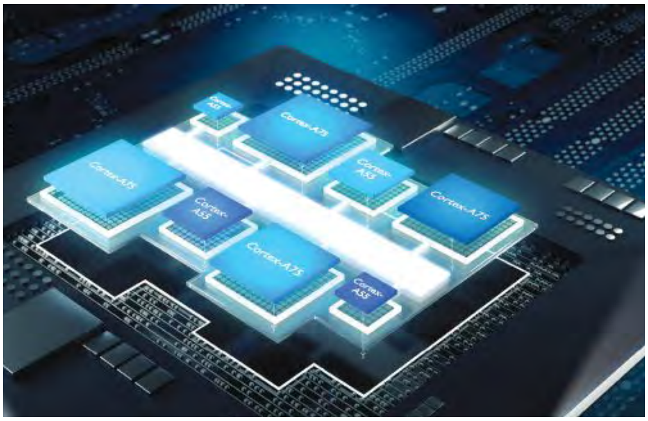

ARM's next-generation architecture, called DynamIQ, aims to improve integration of Cortex-A multicore processors and accelerators. The architecture promises to increase communications between the processing cluster and accelerators 10-fold, as well as provide new instructions to support machine learning and artificial intelligence. *Image courtesy of ARM.*

GPU workstations to cloud installations.

At the same time, Intel has been leveraging its nearly ubiquitous ecosystem of technology to position the Intel Xeon Phi family as being easy for developers. According to the company, it allows implementers to simplify code modernization and reduce programming costs by sharing code and a developer base with Intel Xeon processors. "Standardizing on a unified Intel architecture means you can use a single programming model for all your code," says Intel, "thereby reducing operational and programming expenses through a shared developer base and code reuse."

Some accelerator makers like Graphcore have developed processors customized specifically for machine learning. To help with the deployment of its IPUs, Graphcore has developed Poplar, a graph-programming C++ framework that abstracts the graph-based machine learning development processes from the underlying graph-processing IPU hardware. The framework includes a graph compiler that promises to translate the standard operations used by machine learning frameworks into optimized application code for the IPU. The graph compiler builds up an intermediate representation of the computational graph to be scheduled and deployed across one or many IPU devices.

ARM has unveiled its latest generation of processor designs named DynamIQ. The semiconductor giant contends that chips built using the new multi-core microarchitecture will be easier to configure, allowing manufacturers to connect a wider assortment of CPUs. This could allow for not only more powerful systems-on-chip but also processors that can better perform computing tasks like artificial intelligence.

DynamIQ builds on ARM's "big.LITTLE" approach, which pairs a cluster of "big" processors, with a set of power-sipping "little" ones. DynamIQ takes this flexibility one step further by supporting cores that fall anywhere in between—an approach known as heterogeneous computing. DynamIQ will let chipmakers optimize their silicon, allowing them to build AI accelerators directly into chips, which promises to help systems manage data and memory more efficiently.

Looking to the Future

It's clear after looking at developments like ARM's DynamIQ that the rise of accelerators will not diminish the importance of CPUs. On the contrary, the CPU will continue to play a vital role in future systems. "CPUs are good for general-purpose computing," says Ober. "Modern computers run dozens of applications and numerous background processes, so the need for powerful CPUs will likely continue to increase."

A look at current system frameworks bears this out. SoC applications have evolved to a point where there is typically a central traditional processor complex that orchestrates tasks at a high level, while complex, specialized tasks are distributed to heterogeneous accelerators sprinkled throughout the device.

That said, both traditional CPU makers and accelerator developers are mindful of forces like machine learning and augmented reality, which promise to play a big part in shaping the form and function of the next generation of electronic devices. To ensure their roles in supporting these new technologies, processor providers of all ilks continue to push the limits of their respective technologies.

ARM's new microarchitecture DynamIQ certainly shows the semiconductor giant's intention to provide the flexibility to support applications like artificial intelligence. NVIDIA's Quadro and Tesla GPUs continue to advance the power of accelerators.

But where you are likely to see the most dynamic change is in accelerators that have been built specifically for advanced applications like machine intelligence. A good example of this can be seen in Graphcore's Colossus IPU, a 16 nm massively parallel, mixed-precision floating-point processor expected to become available early in 2018. Designed from the ground up for machine intelligence, this new processor promises to be nothing like a GPU. What it is supposed to do is take the accommodation of new workloads one step further, pushing the envelope on how devices channel and process data. **DE**

Tom Kevan *is a freelance writer/editor specializing in engineering and communications technology. Contact him via de-editors@digitaleng.news.*

INFO → ARM: ARM.com

→ Google: blog.google/topics/google-cloud/google-cloud-offer-tpus-machine-learning

→ Graphcore: graphcore.ai

→ Intel: Intel.com

→ Intuitive: inuitive-tech.com

→ Nervana: intelnervana.com

→ NVIDIA: NVIDIA.com

For more information on this topic, visit digitaleng.news

Rethinking Remote Rendering

The introduction of an external Quadro GPU gives those who prefer thin, lightweight devices the option to use a detachable GPU when needed. *Image courtesy of NVIDIA.*

External GPUs, virtual machines and remote workstations broaden visualization options.

BY KENNETH WONG

THE IMPORTANCE OF PHOTOREALISTIC RENDERING is evident in the CAD-embedded rendering and visualization tools that have, over time, become a standard part of mainstream design software packages. Some, like SolidWorks and Autodesk Inventor, have integrated rendering tools that let you visualize the CAD model in a ray-traced rendered mode with the push of a button. Other packages let you incorporate rendering into the workflow through plug-ins to popular renderers like Luxion Keyshot or Chaos Group's V-Ray.

Because rendering is a compute-intense operation, a rendering job can effectively disable an underpowered system, making it unusable for the duration of the rendering job. To avoid this, users need to have sufficient graphics-processing firepower. The requirement usually translates to a robust, bulky desktop system with multiple processors and GPUs (graphics processing units).

This demand is often in conflict with the user's desire to work from thin, lightweight, connected devices, preferable not just for their portability but also for their aesthetic appeal. In this article, we look at various remote rendering and visualization options available. By detaching the access device from the point of compute, these methods allow users to harness workstation-class performance from thinner, lighter devices.

External GPUs

This August, at SIGGRAPH (Los Angeles, CA), NVIDIA unveiled external GPU (eGPU) offerings, to be delivered by partners Magma, Sonnet, AKiTiO and Bizon. Just like an external modem or storage drive, the independent, external GPU can be attached to a computing device. It requires a Thunderbolt connection.

External GPUs are not new, but the addition of NVIDIA Quadro into the available selection is a significant development. Previously, NVIDIA Quadro was available only as build-in GPUs in preconfigured systems—mobile and desktop workstations. The eGPU option brings Quadro to those who might prefer a thin, light, mobile machine, either for stylistic reasons or for costs, but still need the visualization power of a professional-class GPU.

The detachable option is particularly useful for those who need rendering and visualization occasionally. The eGPU gives such users the ability to plug or unplug the GPU as dictated by their needs.

"For those who have invested in thin and light notebooks, this is an easy way to get the power of NVIDIA's most capable professional GPU. In doing so, you instantly supercharge your prosumer or consumer device to work with larger, more complex 3D models, run programs with interactive ray-tracing, run simulation, even create VR content," says Sandeep Gupte, director of professional visualization.

Workstations as Remote Renderers

Usually people think of remote rendering as rendering on an offsite server or in the cloud, but perhaps a powerful workstation could be used as a remote renderer. "The amount of GPUs you can fit into our high-end desktop chassis, the ThinkStation P920, lends itself to being a mini-HPC (high-performance computing) system on your desktop. If you don't have a data center or IT resources to build a cluster, you could still use [the workstation] as your remote renderer, shared among several employees," says Scott Ruppert, Lenovo's workstation portfolio solutions manager.

The ThinkStation P920's architecture can accommodate as many as three high-end GPUs, linked together via NVIDIA's NVLink technology. NVIDIA describes NVLink as "a high-bandwidth, energy-efficient interconnect

The rendered automotive interior shown here is created with SolidWorks Visualize. The software, SolidWorks Visualize Professional, offers a network rendering license to speed up the job. *Image courtesy of SolidWorks.*

that enables ultra-fast communication between the CPU and GPU, and between GPUs."

Lenovo also offers the ThinkStation P320 Tiny, roughly the size of a hardcover book. The small device features a GPU, the NVIDIA Quadro P600. It's adequate to power CAD visualization, but not for rendering complex scenes. Ruppert proposes such a small device can be the everyday device you use to access the more powerful ThinkStation P920 for remote rendering and computing.

HP's professional workstations, such as the HP Z, usually come with HP Remote Graphics Software (RGS). You can use the software to "access, share and broadcast your Windows and Linux workstation apps in amazing, high-speed clarity from any remote PC, Mac or Windows tablet with HP Remote Graphics Software (RGS)," HP says.

The software allows you to access and control a powerful GPU-accelerated desktop workstation from thin, lightweight devices; therefore, you could use this method to remotely render and visualize complex scenes from smaller connected devices.

Dell also offers remote computing options. For example, Teradici PCoIP (PC over internet protocol) remote access software is compatible with all Dell Precision tower and rack workstations. BOXX Technologies also offers Teradici solutions. In his review of the solution (Virtualization: "Access Your Workstation from Anywhere," *DE*, July 2015), Contributing Editor David Cohn wrote "Seeing SolidWorks run on the inexpensive Dell laptop and the iPad was simply astounding."

Virtual Machines for Remote Visualization

In August, NVIDIA launched the Quadro Data Center Workstation (vDCW), a new virtualization solution. The vDCW software can "turn NVIDIA Tesla GPU-accelerated servers into powerful workstations," NVIDIA explains.

The reference to Quadro in its name warrants further explanation. NVIDIA's GPU-accelerated servers are built on NVIDIA Tesla GPUs, not Quadro. Quadro is the GPU brand for discrete workstations. But with the help of vDCW software, you can create and support virtual machines (VMs) that work and behave like NVIDIA Quadro-powered workstations, though they're in reality running on Tesla-powered hardware. In other words, the VMs will have Quadro characteristics, but are running on Tesla under the hood.

vDCW provides up to "24GB of GPU memory for working with large, immersive models," NVIDIA states. Because VMs are accessible from mobile devices, it offers users a way to get the same visualization and rendering power found in Quadro-powered machines without having a physical desktop workstation.

The vDCW, NVIDIA says, "addresses the increasingly compute-intensive workflows—with their exponential growth in data size and complexity—associated with new technologies for 3D, photorealistic rendering, virtual reality and deep learning."

Network and Cloud Rendering

Many rendering applications, such as Luxion's CAD-friendly Keyshot and SolidWorks' CAD-embedded rendering function, include a network rendering mode, which gives users the ability to compute the rendering job on a pool of machines instead of a single machine. Because it harnesses the power of a greater number of processor cores than what is available in a single machine, this option speeds up the rendering job significantly. However, it requires that the users have access to an internal server or an on-premise cloud.

With some rendering applications, users have the option to tap into on-demand cloud resources. As part of its subscription, Autodesk offers cloud services, which give users the option to send rendering jobs to remote servers maintained by Autodesk. The cloud-rendering feature is available to subscribers of the Autodesk architecture collection, media and entertainment collection, and product design collection, among others.

This means users of Autodesk Revit (part of the architecture collection) or 3DS Max (part of the product design collection) may send complex ray-traced renderings of architecture projects or automotive assemblies to the cloud, leaving their own desktops or laptops unaffected by the rendering jobs' burden.

With network rendering and cloud rendering, it is important to know the underlying renderer's technology. To accelerate jobs on a CPU-based renderer like Luxion Keyshot, you would need to designate more CPU cores to the network cluster or the cloud-hosted cluster. By contrast, GPU-accelerated renderers like Octane, Iray, Chaos Group's V-Ray and Nextlimit's Maxwell are significantly accelerated by GPU cores.

Remote rendering is also available in browser-based options from vendors such as Clara.io and Lagoa. The software-as-a-service style renderers usually let you upload your 3D CAD model to the server, apply materials and backgrounds, then finish the rendering job online. With this approach, the computing power of the user's own device is irrelevant. Therefore, those working from Google Chromebooks, iPads or tablets may remotely render complex scenes without being hampered by the lack of cores in their own system.

Whereas vendors like Clara.io and Lagoa offer robust model-preparing and editing tools, material libraries and high-definition backgrounds to build a scene from the ground up (so to speak), vendors like Rendercore offer on-demand, remote hardware without any online rendering software. The company's tiered pricing is based on priority—the speed with which you want your job to be completed. Premium (high priority) jobs are priced $0.25 per core per hour; economy (low priority) jobs are $0.15 per core per hour. With monthly and weekly passes (monthly, $379; weekly, $139 per node), the per-core pricing can be as low as $0.01 per GHz hour. **DE**

Kenneth Wong *is* DE*'s resident blogger and senior editor. Email him at de-editors@digitaleng.news or share your thoughts on this article at digitaleng.news/facebook.*

INFO → **Autodesk:** Autodesk.com

→ **BOXX Technologies:** BOXX.com

→ **Clara.io:** Clara.io

→ **Dell:** Dell.com

→ **HP Remote Graphics Software:** www8.hp.com/us/en/workstations/remote-graphics-software.html

→ **Lagoa:** home.lagoa.com

→ **Lenovo:** www3.lenovo.com

→ **Luxion Keyshot:** Keyshot.com

→ **Rendercore:** Rendercore.com

→ **SolidWorks:** Solidworks.com

Simulation-as-a-Service On-Ramp

How network infrastructure affects small- and mid-size business access to cloud-based simulation.

BY BRIAN ALBRIGHT

SIMULATION IS TRANSITIONING from a highly specialized operation performed by experts at the end of the design process to a more ubiquitous activity that can help optimize products throughout the design cycle. That means companies are doing more simulation and asking for results much faster.

Cloud-based solutions that enable simulation-as-a-service—either through complete outsourcing of simulation or opportunistically accessing simulation tools on an as-needed basis—have made this easier.

"Whereas a few years ago clients were really not averse to bringing simulation tools in house, development schedules and markets are moving at such a fast pace that clients are willing to pay a premium to leverage outside expertise and resources," says Scott Herndon, manager of simulation client development for CFD (computational fluid dynamics) and FEA (finite element analysis) at IMAGINiT Technologies. IMAGINiT offers cloud-based Autodesk products, and has also developed its own mini-cloud resources on its own supercomputers.

"The costs of simulation were too big of a barrier for smaller companies," says Ray Leto, president of TotalSim. "The idea of simulation-as-a-service is to take all of the expertise and resources [and] put it in a black box with an easy-to-use interface so that customers can walk through it and use it as they need it."

"With smaller companies that have not used HPC in the past, by having access to the cloud they are able to run bigger jobs once they outgrow what they can do on a workstation," says Gabriel Broner, vice president and general manager of HPC at Rescale.

Although the availability of cloud-based simulation-as-a-service solutions is expanding, in some cases these operations require the transfer of large amounts of data or the use of simulation tools that require very low latency. This can prove to be challenging, particularly for small- to mid-size businesses (SMBs) that may not have access to direct, high-speed connections to cloud providers or other network infrastructure.

"The reliability of the network is critical because the last

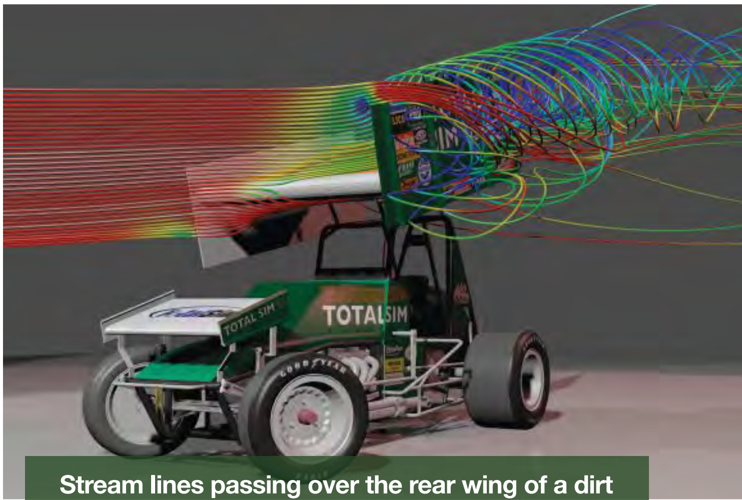

Stream lines passing over the rear wing of a dirt track vehicle. *Images courtesy of TotalSim.*

thing you want is to put a lot of work into putting a solver into the cloud and have it fail because of a bandwidth or data connection issue," Herndon says. "We rarely have issues with that, though, and it happens much less now than it did in years past."

The availability of new high-speed networks can enable new types of services. Simulation services company TotalSim, for example, has been able to leverage access to a local gigabit fiber network at its home base in Dublin, OH, as well as the resources of the statewide OARnet (Ohio Academic Resources Network) 100G regional network and the Ohio Supercomputer Center.

"I can't overstate how important it is for us within the city of Dublin to have access to these networking infrastructures," says Leto.

Leto's company was involved in a project to make simulation applications available to customers using those high-speed network resources. Working with researchers at Ohio State University and other technology partners, TotalSim was able to offer simulation capabilities to customers with very little latency. "We're not networking people, so we're not always thinking about the challenges on that side of the problem," Leto says.

The push for municipalities to expand their gigabit network infrastructure, and then offer access to those networks ilike they offer access to water, sewer and other utilities, will help make these cloud-based applications easier to use and access.

That is because the network plays an important role in making these applications work for users that need real-time responsiveness, says Prasad Calyam, assistant professor of computer science at the University of Missouri College of Engineering, who worked with TotalSim on the project while he was acting as research director at the Ohio Supercomputer Center/OARnet at Ohio State University. "Integrating the various desktops and HPC systems, storage and other collaborators, and providing connectivity across them can provide a huge improvement in workflows," Calyam says. "It helps improve the gains in time, cost, effort and convenience. Networking really drives this transformation to the cloud."

Companies are also finding other ways to tap into cloud-based simulation resources. Workload data can be placed directly in the cloud so that the data is not moving back and forth between the user and the cloud infrastructure.

At Rescale, customers can connect via the public internet, while Rescale manages connections to cloud services providers. "We also have the ability to offer dedicated high-speed links," Broner says. "For customers that already have some systems on premise, we can work collaboratively with them to determine what jobs are good candidates to move to the cloud."

High-Speed Access in Ohio

In Dublin, TotalSim teamed with the city, Ohio State University, the OARnet network and other entities to create an app-based approach to providing access to simulation and compute resources to local businesses.

The team at OSU worked with TotalSim, which uses HPC resources to test virtual prototypes for clients in the aerospace, automotive and manufacturing markets. The problem TotalSim wanted to solve was to find a way to provide access to data-intensive services without overwhelming the public network and bogging down simulation activities for its clients.

The project was funded in part through an award from US Ignite, a nonprofit focused on helping to create services and applications that leverage advanced networking technology. Launched in 2012 by the National Science Foundation and White House Office of Science and Technology Policy, the initiative has helped launch and support Smart Gigabit Communities (funded by the NSF) that offer access to high-bandwidth networks (among other things). Those cities include Cleveland, Austin, Chattanooga, Kansas City and others.

The project was named "Best Application for Advanced Manufacturing" at the US Ignite Next Generation Application Summit in Chicago. Researchers from the Ohio Supercomputer Center (OSC), OARnet, Ohio State University, the City of Dublin, Metro Data Center (MDC) and the University of Missouri-Columbia (MU), in partnership with TotalSim, VMware and HP, were awarded $25,000 to develop "Simulation-as-a-Service for Advanced Manufacturing."

The application allows users to remotely access the software and compute resources through a virtual desktop-as-a-service system for manufacturing. Users access results of simulations via a thin-client connection to a virtual desktop, whereas the heavy lifting of the large data sets is handled on the DubLink and OARnet fiber networks.

Dublin rolled out its own gigabit network called DubLink, which serves as the backbone of the project. It connects to Metro Data Center, a regional supercomputing facility in the city and runs parallel to OARnet, Ohio's statewide 100 gigabit network.

Calyam says the project built on existing networking resources that had not, to that point, been operationalized with any applications. Funds from winning a Mozilla competition helped launch the prototyping phase, and Calyam was able to obtain a donation of cloud data center GENI racks and last-mile fiber connections.

"They created a modeling and simulation service that people can access as if they had their own solution because of the speed and low latency of the fiber network," says Glenn Ricart, founder and CTO of US Ignite. "They can rotate models in real time."

By reducing latency, the application allows better real-time collaboration between TotalSim and its customers, which can speed up iterations.

Having local infrastructure was critical, as distance between the companies using the app and the compute resources can affect latency. "Even if you have a gigabit network all the way from here to Oregon, you can't necessarily provide the responsiveness you'd get having the application hosted in Dublin and distributed over local Dublin fiber," Ricart says.

Network Alternatives

Although high-speed fiber networks are expanding, and the number of Smart Gigabit Communities is growing, not every small or mid-sized company has an on-ramp to this infrastructure. Fortunately, there are other ways to successfully access cloud-based or hosted simulation solutions.

At Rescale, customers submit simulation jobs and are able to match them to the most effective compute architectures for that particular project. "Having more bandwidth is always going [to] help things for SMBs, but even without a gigabit connection, there are a lot of workloads that can run in the cloud really well," says Ryan Kaneshiro, chief architect at Rescale. "The sweet spot is probably CFD jobs that have small input file sizes and generate a lot of output data. You can do post-processing remotely through a remote visualization node or batch load processing script to whittle down the amount of data that needs transferred back to the local workstation."

For other jobs with larger data transfer needs, more bandwidth is crucial. "That said, it isn't the only option," Kaneshiro says. "If you are talking about static data sets, a lot of that information is already sitting in the cloud. If it's a one-time transfer, then there's also the option of shipping hard drives to Amazon or Microsoft or whatever provider you are using."

More direct connections are also available. "We see direct connections filtering down to SMBs as cloud usage starts to grow," Kaneshiro says. For example, services like Megaport can help companies establish those direct connections to cloud services.

US Ignite's Ricart says that high-speed networks will always have some latency limitations because of the speed of light and the way the fiber networks are designed (which is typically not in a straight line and requires several hops). "Every time you take a signal and put it through a router, that creates a delay," Ricart says. "Both the speed of light and the number of times you have to make a connection run the clock when it comes to latency."

TotalSim has successfully worked with clients across the country by using its app design to work around the latency issue. "What we've found is that the way we've designed the UI (user interface) side of the web application, and the latest software stack, is that it has allowed rendering and visualization and transfer of data to be faster than it used to be," Leto says. "A person in California that is accessing the apps is still looking at a web page coming out of the Ohio Supercomputer Center. We're just sending images and charts and graphs to be rendered in a browser, which is very lightweight and fast."

Where the company does see a challenge is with interactive 3D post-processing or manipulation. "But every year the tools are getting better, and people are figuring out how to make the rendering work faster and the remote visualization capabilities through the browser are getting better," Leto adds.

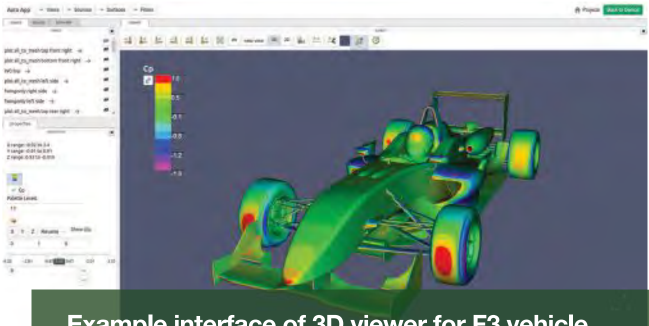

Example interface of 3D viewer for F3 vehicle using TotalSim's results web application hosted at Ohio Supercomputer Center.

Taking advantage of solutions that offer different access options is important for making simulation-as-a-service work for smaller companies. Leto says that "bare metal" HPC installations provide the best bang for the buck when it comes to CFD simulations. Public cloud services like Amazon or Google can provide greater scalability.

It's also important that the various high-speed networks typically available in larger cities can be coordinated and integrated as these high-speed networks expand, and that's been a big part of US Ignite's efforts. "We are rapidly seeing the 'gigafication' of the American internet," Ricart says.

With a local hub that connects those networks, communities can reduce latency; otherwise, traffic may flow hundreds of miles away before it can be relayed and exchanged between network providers. Cities can establish digital "town squares" where a variety of applications and services are available for use. That presents an opportunity for companies with heavy simulation needs to access compute resources they would otherwise be unable to afford, while helping cities attract and retain high-tech businesses.

"In the case of Dublin, the city is already invested in data centers and has resources they can use as incentives for companies," Calyam says. "If TotalSim is using this, then other companies will look at it and be able to do a cloud-based transformation. Those companies are much less likely to relocate to another city because they have these resources." **DE**

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

INFO ➜ **HP:** HP.com

➜ **IMAGINiT Technologies:** imaginit.com

➜ **Metro Data Center:** metrodatacenter.com

➜ **Ohio Supercomputer Center:** osc.edu

➜ **Rescale:** rescale.com

➜ **TotalSim:** totalsim.us

➜ **US Ignite:** us-ignite.org

➜ **VMware:** vmware.com

For more information on this topic, visit digitaleng.news.

A VR-Ready POWERHOUSE

New MSI WT73VR mobile workstation offers tons of power to go.

BY DAVID COHN

TAIWAN-BASED MSI (Micro-Star International) recently sent us yet another well-appointed mobile workstation. However, unlike the thin, lightweight WS63 we recently reviewed (*DE*, September 2017), the mighty WT73VR system is quite large and meant to be a desktop replacement fully capable of running design, engineering and virtual reality applications.

The new VR-ready MSI WT73VR 7RM-648US mobile workstation is big and heavy, and expensive, but delivers exceptional performance. *Image by David Cohn.*

The WT73VR measures 16.85x11.3x1.93-in. (WxDxH) and weighs a hefty 8.91 pounds. Its large (6.56x3.38x1.38-in.) 230-watt external power supply adds more than 2 pounds. The brushed black aluminum case looks great, with its glowing green MSI workstation logo centered between two glowing white decorative blades, and immediately sends the message that this system is fast.

Unlike many other original equipment manufacturers, MSI does not offer custom configurations. Instead, the company sells only preconfigured systems based on specific combinations of components. For example, the computer we received—the $4,699 7RM-648US—is based on a 2.90GHz quad-core Intel Core i7-7820HK CPU and a 1920x1080 display. MSI also offers the 7RM-687US ($4,999), powered by a 3.0GHz quad-core Intel Xeon E3-1505M processor and sporting a 4K (3840x2160) IPS display.

Although both of those Intel CPUs include integrated graphics, MSI also equips both versions of the WT73VR with a powerful NVIDIA Quadro P5000 graphics card. This discrete GPU includes 16GB of dedicated GDDR5 graphics memory and 2048 CUDA (compute unified device architecture) cores.

Both MSI WT73VR systems also come with 64GB of RAM, installed as four 16GB DIMMs (dual inline memory modules), although the Xeon-powered system uses ECC (error-correcting code) memory. And both versions of the WT73VR also include a fast SSD (solid-state drive) primary drive as well as a standard hard drive. Our evaluation unit came with a 512GB Samsung NVMe SSD and a 1TB Hitachi 7200rpm SATA drive.

Plenty of Ports

Lifting the lid reveals the 17.3-in. display and a very nice SteelSeries keyboard with 102 backlit keys, including a separate numeric keypad. MSI's Dragon Center keyboard manager app lets you adjust the backlight color. A 1080p webcam is centered above the display with a single microphone to one side. An LED adjacent to the webcam glows white when the camera is active.

A 4x2.75-in. touchpad with multitouch capabilities is centered below the spacebar and includes two dedicated buttons. To the right of the keyboard are five buttons, including a V-shaped power button that glows white when the system is powered on. The other buttons below this enable you to toggle between the discrete GPU (graphics processing unit) or integrated Intel graphics, toggle the cooling fan speed, access the keyboard manager app and launch a user-defined application. A pair of 3-watt stereo speakers is located above the keyboard, with a 5-watt woofer on the underside of the case.

The right side of the case provides a pair of USB 3.0 ports, an SD card reader, a ventilation grille and a Kensington lock slot. The left side of the case includes three additional USB 3.0 ports, four audio jacks (headphone, microphone, line-in and line-out) and another ventilation

Mobile Workstations Compared	**MSI WT73VR** 17.3-in. 2.9GHz Intel Core i7-7820HK quad-core CPU, NVIDIA Quadro P5000, 64GB RAM, 512GB PCIe SSD and 1TB 7200rpm SATA HD	**MSI WS63** 15.6-in. mobile 2.8GHz Intel Core i7-7700HQ quad-core CPU, NVIDIA Quadro P3000M, 32GB RAM, 512GB PCIe SSD and 2TB 5200rpm SATA HD	**Eurocom Tornado F5** 15.6-in. mobile 3.6GHz Intel Xeon E3-1270 quad-core CPU, NVIDIA Quadro M4000M, 32GB RAM, 2TB PCIe SSD	**Lenovo ThinkPad P50s** 15.6-in. mobile 2.6GHz Intel Core i7-6600U dial-core CPU, NVIDIA Quadro M500M, 16GB RAM, 512GB PCIe SSD	**MSI WT72 6QN** 17.3-in. 2.9GHz Intel Core i7-6920HQ quad-core CPU, NVIDIA Quadro M5500, 32GB RAM, 256GB PCIe SSD RAID 0 and 1TB SATA HD	**Lenovo P40 Yoga** 14.1-in. 2.6GHz Intel Core i7-6600U dual-core CPU, NVIDIA Quadro M500M, 16GB RAM, 512GB PCIe SSD
Price as tested	$4,699	$2,599	$5,450	$1,427	$4,999	$1,705
Date tested	6/28/17	4/3/17	2/13/17	10/10/16	9/15/16	7/27/16
Operating System	Windows 10	Windows 10	Windows 10	Windows 10	Windows 10	Windows 10
SPECviewperf 12 (higher is better)						
catia-04	157.84	96.83	85.32	21.75	128.73	19.98
creo-01	129.89	87.28	80.21	25.34	103.28	24.34
energy-01	12.56	11.59	6.36	0.52	16.25	0.61
maya-04	100.99	66.22	60.58	13.27	81.64	12.25
medical-01	59.31	39.09	27.39	9.68	61.03	14.03
showcase-01	67.53	54.80	48.46	6.97	58.88	6.81
snx-02	185.13	71.52	78.14	31.85	120.83	26.46
sw-03	160.26	103.08	100.19	37.24	118.06	35.31
SPECapc SOLIDWORKS 2015 (higher is better)						
Graphics Composite	4.95	4.38	7.60	2.67	5.99	2.65
Shaded Graphics Sub-Composite	3.06	2.71	4.14	1.96	3.69	1.78
Shaded w/Edges Graphics Sub-Composite	3.89	3.50	5.46	2.52	4.84	2.40
Shaded using RealView Sub-Composite	3.54	3.14	5.64	2.01	4.77	2.00
Shaded w/Edges using RealView Sub-Composite	4.27	3.81	9.20	3.43	7.80	3.42
Shaded using RealView and Shadows Sub-Composite	4.07	3.61	6.44	1.96	5.16	2.03
Shaded with Edges using RealView and Shadows Graphics Sub-Composite	4.51	4.03	9.56	3.14	7.97	3.22
Shaded using RealView and Shadows and Ambient Occlusion Graphics Sub-Composite	13.46	11.77	16.22	3.02	9.15	3.38
Shaded with Edges using RealView and Shadows and Ambient Occlusion Graphics Sub-Composite	13.17	11.53	23.22	4.53	13.57	5.07
Wireframe Graphics Sub-Composite	3.91	3.33	3.65	2.61	3.20	2.20
CPU Composite	4.28	3.97	4.23	1.89	2.39	1.95
SPECwpc v2.0 (higher is better)						
Media and Entertainment	3.12	2.80	2.96	1.04	2.64	0.99
Product Development	3.13	2.78	2.49	1.28	2.65	1.11
Life Sciences	3.60	3.27	3.05	1.25	3.08	1.25
Financial Services	2.90	2.81	3.10	0.49	1.24*	0.49
Energy	2.94	2.74	2.60	0.96	2.61	0.87
General Operations	1.45	1.37	1.37	0.87	1.37	0.85
Time						
Autodesk Render Test (in seconds, lower is better)	67.00	52.90	78.30	172.50	73.20	149.00
Battery Test (in hours:minutes, higher is better)	2:55	4:20	3:20	11:44	3:09	9:10

Numbers in blue indicate best recorded results. Numbers in red indicate worst recorded results. ** Results provided by MSI.*

Price/Performance chart based on SPECwpc Product Development benchmark dataset.

grille. A panel on the rear of the case provides an RJ45 LAN connector, a mini-DisplayPort, a Thunderbolt/USB Type-C port, an HDMI video port and the connector for the external power supply, centered between two more ventilation grilles. Three small LEDs across the front of the case indicate Wi-Fi, battery and hard drive status. The bottom of the case is essentially one large ventilation grille.

As has become quite common, the battery is not removable and the manual does not address any user-serviceable components. During our test, the MSI WT73VR remained cool but was hardly quiet. Although its 35-decibel noise level was barely audible at rest, this climbed to 57dB when running our benchmarks, and 65dB (equivalent to vacuum cleaner) when the fan was switched to its high speed. The eight-cell 75Whr Li-ion battery kept the system running for just two hours and 55 minutes on our battery rundown test.

Exceptional Performance

The MSI WT73 7RM-648US performed extremely well on all of our benchmark tests. Thanks to its Pascal-based NVIDIA GPU, it outperformed other mobile workstations we've reviewed recently on almost every dataset in the SPECviewperf test of graphics performance.

On the SPECapc SolidWorks 2015 benchmark—which is more of a real-world test—the MSI WT73VR also did very well, although it did lag the two mobile systems we previously tested that were equipped with much faster CPUs. On the demanding SPECwpc benchmark, the WT73VR also did extremely well, again garnering top results in almost every category. On our own AutoCAD rendering test, the 67-second average rendering time was respectable, but certainly didn't break any records for mobile workstations.

MSI preloads Windows 10 Professional 64-bit and backs the system with a three-year limited warranty that includes a one-year global warranty. And unlike many other lesser-known brands, MSI mobile workstations are ISV (independent software vendor) certified for major CAD/CAM software from Autodesk, Siemens and SolidWorks.

The MSI WT73VR 7RM-648US definitely delivers great performance, but with its $4,699 price tag, it is likely to appeal to a more limited set of potential users. **DE**

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

INFO ➔ Micro-Star International Co, Ltd.: msi.com

MSI WT73 7RM-648US

- **Price:** $4,699 as tested
- **Size:** 16.85x11.3x1.93-in. (WxDxH) notebook
- **Weight:** 8.91 pounds as tested, plus 2.15-pound power supply
- **CPU:** 2.8GHz Intel Core i7-7820HK quad-core w/8MB cache
- **Memory:** 64GB 2400MHz DDR3
- **Graphics:** NVIDIA Quadro P5000 w/16GB memory and 2048 CUDA cores
- **LCD:** 17.3-in. diagonal (1920x1080), non-glare
- **Hard Disk:** 512GB M.2 PCIe SSD and 1TB 7200rpm SATA
- **Optical:** none
- **Audio:** line-in, line-out, microphone-in, headphone-out (with SPDIF); plus built-in microphone and speakers
- **Network:** integrated Gigabit Ethernet (10/100/1000 NIC) with one RJ-45 port, 802.11ac wireless LAN and Bluetooth 4.2
- **Modem:** none
- **Other:** five USB 3.0, one USB 3.1 (Type C) Thunderbolt port, mini-DisplayPort, HDMI-out, 1080p webcam, SD card reader
- **Keyboard:** integrated 102-key backlit keyboard with numeric keypad
- **Pointing device:** integrated touchpad

Make Your Move into Metal Additive Manufacturing

Take a holistic approach to implementing metal additive manufacturing by partnering with experts.

GIVEN ALL THE HYPE surrounding 3D printing, it's easy to forget that the technology isn't some relative newcomer or unknown bit player. Rather, 3D printing, including metal additive technology, has increasingly become a principal actor on the manufacturing stage, used as a cost-effective production alternative for countless applications across myriad industries for the better part of a decade.

Metal additive manufacturing (AM), while commanding less of the current spotlight than thermoplastic 3D printing technologies, has become a proven production method in the last few years. The technology has gained significant traction in industries such as aerospace, automotive, medical, and dental, where there is outsized demand for relatively low volumes of parts, and where customization and the ability to combine parts for optimization purposes can play a significant role in differentiating a product or driving cost efficiencies throughout its lifecycle.

Metal AM systems vary based on the technology employed, but one of the most popular technologies, selective laser melting (SLM), spreads a layer of metal powder onto a substrate plate, and then a high-power laser selectively melts the powder to create the first layer of the build. A fresh layer of metal powder is then evenly distributed over the build surface, and the lasers melt each successive layer until the desired component is produced.

While AM technologies like SLM have been part of leading aerospace and automotive companies' manufacturing playbooks for the last few years, usage has been limited, in part because of the high cost of the additive manufacturing systems and materials, in addition to a general lack of knowledge and expertise in how to fully use additive methods. Thanks to advances in technology and manufacturers' growing familiarity with additive manufacturing, the industry is rapidly expanding. Worldwide revenues increased 17.4% in 2016, according to the 2017 *Wohlers Report*. Metal AM, in particular, is enjoying significant traction with nearly half of all service providers in the market running AM systems able to produce metal parts—an indication of increased customer demand.

Widespread interest in AM, and metal AM in particular, has been stoked by success stories from a range of industries, but particularly in the aerospace and defense sector. NASA, for example, leveraged selective laser melting 3D printing

Build an AM Team

Take the guesswork out of metal additive manufacturing and reduce costly trial and error by taking a holistic approach to implementing the technology.

Learn how to set yourself up for success by choosing a metal AM partner that has a knowledgeable team in place to help you. Four key players to look for in a metal AM vendor partner are the applications engineer, the systems engineer, the metallurgist and the service engineer.

"Making the Case for Building an AM Team" explains why these four key players are critical, how to move beyond AM manufacturers' speeds and feeds specifications to leverage the technology for optimal output based on your requirements, and real-world examples of how leading organizations are using metal AM to save time while reducing costs and product complexity.

Download "Making the Case for Building an AM Team" here: digitaleng.news/de/AMteam.

MANY INDUSTRIES are taking advantage of the benefits offered by metal additive manufacturing, such as shorter production times, lightweigthing and lower total costs. ***Images courtesy of SLM Solutions.***

technology to output turbo pumps, injectors, and valves (75% of the necessary parts for a fully 3D printed rocket engine) and tested all three individually with great results. By leveraging SLM technology, NASA was able to design each part with fewer components—for example, the injector was produced in just two pieces compared to the more than 200 parts for a similar, traditional component—and the team was able to construct the components in months, not years. More recently, NASA connected the 3D printed parts together as a traditional rocket engine, which produced more than 20,000 pounds of thrust while withstanding temperatures of 6,000°F.

GE also used metal AM to create its redesigned Leap engine fuel nozzle. Whereas it historically took 20 parts to complete the fuel nozzle's complex geometry, 3D printing allowed for a simplified fuel nozzle design that can be produced as a single unit and is five times more durable than its predecessor, according to the company. Others in and outside of the aerospace industry are catching on to the cost efficiencies and part optimization benefits of AM technologies. According to a study funded by the U.S. Department of Energy Advanced Manufacturing Office, aircraft weight can be reduced by 7% by replacing conventional manufacturing methods with AM technologies. With fuel expenses now ranging from 25 to 40% of total airline operating costs, every percentage of weight reduction equals huge cost savings.

AM technologies are also getting more air time as talk of industry 4.0 or digital manufacturing takes hold. As part of the vision for next-generation manufacturing, 3D printing joins technologies like big data analytics, simulation, autonomous robots, and the industrial internet of things (IIoT), to deliver new levels of intelligence and automation, enabling faster, more flexible manufacturing processes that allow companies to produce high-quality goods at reduced costs.

Perception vs. Reality

While the race is on, one of the biggest misperceptions is that 3D printing has evolved to a point where it is dead simple to use. Those with limited 3D printing experience who are focused primarily on prototyping are more likely to miscalculate the intricacies of tapping metal AM technologies for production purposes. Just like traditional manufacturing processes like casting, machining, forming, and injection modeling, metal AM is a specialized discipline that requires the expertise of trained professionals. While the primary metal AM technologies predominantly use the same basic melting process, there are critical differences in how the actual machines are designed and operated.

At the same time, manufacturers need to understand how to design parts for metal AM production so they can fully capitalize on the technology. For example, metal AM allows for customization of parts and complex lattice structures that are critical for lightweighting and aren't possible with traditional manufacturing methods. There are also design requirements specific to the behaviors and characteristics of the metal materials supported by particular machines. All of these considerations need to be factored in at the earliest design stages.

There's no question that metal AM technology has arrived at the point where it's a viable production alternative for companies of all sizes, across many industries. Manufacturers just need to ensure they choose the right metal AM technology for their application. They also need to align with a trusted and expert partner that can help navigate the inevitable twists and turns in what's likely to be a rewarding and transformative journey.

Learn more about taking a holistic approach to metal additive manufacturing by downloading "Making the Case for Building an AM Team," a complimentary paper produced by *Digital Engineering* in partnership with SLM Solutions. Download the full paper here: digitaleng.news/de/AMteam.

SLM SOLUTIONS

Industrial Issues in Additive Manufacturing

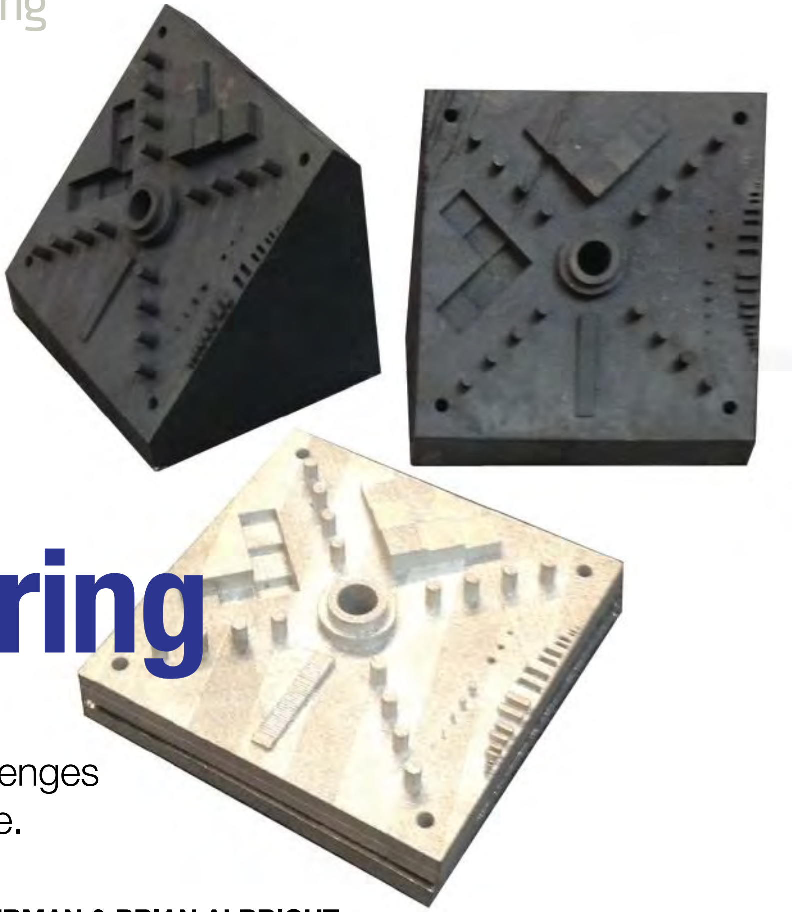

3D printing industry confronts challenges of consistency, education and more.

BY STEPHANIE SKERNIVITZ, PAMELA J. WATERMAN & BRIAN ALBRIGHT

THE 3D PRINTING INDUSTRY'S growth continues to accelerate at astronomical rates. According to figures from the "2016 Wohlers Report," in 2015, this industry, which includes global 3D printing products and services, had a compound annual growth rate of 25.9% or $5.165 billion.

The growth rate is even more pronounced on the metal side of 3D printing. According to Wohlers Associates, over 800 metal 3D printing machines were sold in 2015, a growth of 46.9% over 2014, when 550 metal 3D printing machines were sold.

All promising news, yet there's plenty of work to be done to keep the industry growing for years to come. In addition to industry-wide concerns on how to lower costs and maximize benefits for more users, there's another ongoing hurdle in churning out a quality product that is consistent, time after time.

Consistency is no small feat. It encompasses system reliability and process repeatability, especially when using 3D printing for production applications. To date, 3D printing system manufacturers are tackling these challenges head on with process monitoring and control software, but much work remains. Factors like lack of consistency across final products or lack of understanding of how materials work in 3D printing surface often. Not to mention the list of materials to use is still relatively short.

Material Issues

Industry leaders such as America Makes and plenty of big-name additive manufacturing (AM) companies aren't sitting idle waiting for someone else to land on a solution.

At the America Makes seminar, Tracy Albers, president and CTO of RP+M, didn't mince words when she suggested that those in the trenches of AM are working with approaches that are highly variable and that can't be reproduced consistently, not to mention the actual AM processes are often not even well understood. In the case of RP+M, the company has been running validation testing on Ultem 9085 on Stratasys Fortus 900mc machines to measure the level of variability in the process. The results? Albers reported that preliminary tensile strength tests revealed 30% coefficient of variance related to parts generated on the same equipment. These results prompted Stratasys to respond by redesigning the tip in Fortus machines, which led to reduced variance.

Getting AM-generated parts to perform consistently presents a massive hurdle for its widespread adoption in manufacturing. For instance, GE spent five years laboring to achieve FAA certification of its LEAP fuel nozzle, now

mass produced via AM. At first, only 30% of the parts met the company's specifications.

"There is a lot of variation among machines. How do we make one part, and make that part repeatedly so we can prove to the FAA that the material properties we've provided are true?" asked Deb Whitis, materials leader at GE Additive during that America Makes seminar.

With companies on the fence about using metal AM equipment to generate prototypes and production parts, quality and consistency have to be in place for them to deem the investment worthwhile. Currently being able to print the same part on the same machine time and again just is not the reality. Two parts, side by side, may be two different sizes, or there may be cracks in one, according to experts.

That's where the National Institute of Standards and Technology and Physical Measurement Laboratory (PML) enter the picture. Consistency, or lack thereof, in producing parts tops their priority list and dovetails nicely with their new joint project, the Additive Manufacturing Metrology Testbed (AMMT). This custom-created 3D printer is built to highlight how the whole AM process works.

The open platform testbed measures the size of a small car, according to NIST, and it is built to print in stainless steel, cobalt chrome and nickel alloy. According to NIST and PML, the testbed can track the temperature of the melt pool by measuring brightness of light reflected off of it while printing. A future goal for NIST would be to create a temperature map of the surface of the printed object over various light wavelengths.

Another significant additive manufacturing concern relates to lack of education about design-for-manufacturing, according to Star Rapid's President Gordon Styles. Star Rapid, like other service bureaus using AM, is not waiting for others to address the education problem.

He says a common recommendation is to "design for the process." "This may sound obvious, but not all plastic 3D printing rules can be applied to metal 3D printing, which can lead to parts failing during the build process," Styles says.

That's why he and his company have launched a video tutorial series to give potential solutions for product designers and engineers using powdered metal AM technologies. So far, topics have addressed familiar geometry/build flaws regarding wall sizing, gaps and holes, as well as how to create bridges and overhangs.

Other design and manufacturing companies are following suit with educational endeavors. "It's important to

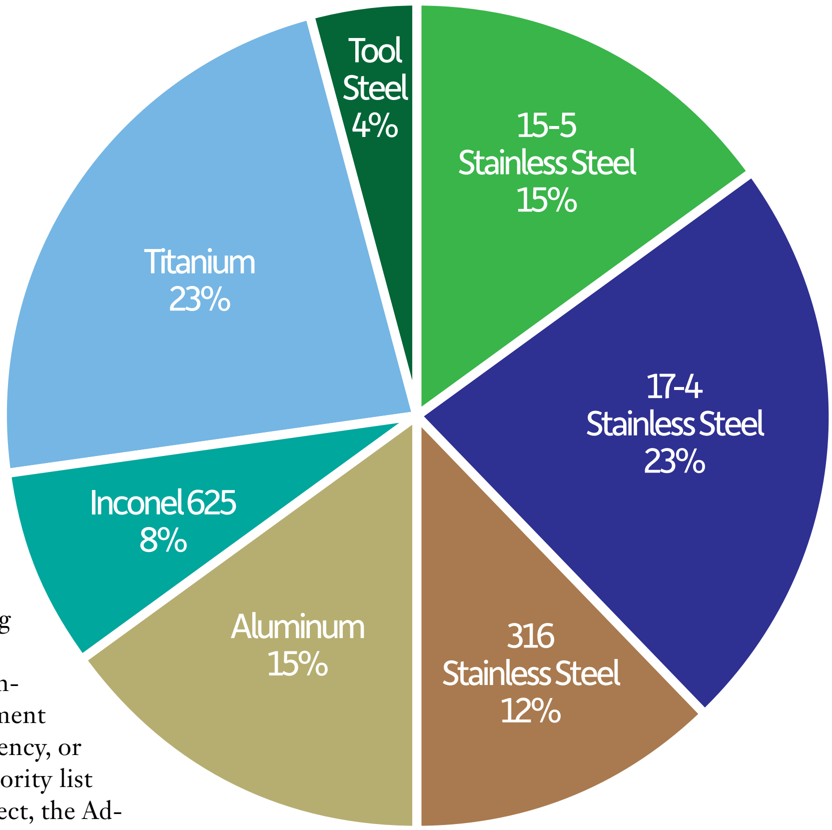

Percentage breakout by type of metal chosen for industrial 3D printing projects, averaged over past two years of parts procured through 3Diligent. ***Image courtesy of 3Diligent.***

recognize that (AM systems) are manufacturing tools just like any other, with different flavors and really different capabilities," notes Joe Manzo, CEO of Titan Industries, Tempe, AZ. His company uses electron beam melting (EBM), and operates two Arcam EBM Q20plus systems for partner company LAI International as part of its design engineering business. The systems can produce parts in Inconel, cobalt chrome and two titanium grades.

As a value-add, Titan Industries offers designers instruction about Arcam's AM capabilities. "E-beam and laser powder-bed sintering systems have different strengths, and a lot of that is based on the physics," Manzo explains. "With e-beam, you have a higher power source, which lets you melt thicker layers, which results in faster build times."

Obstacles for designers may include possible geometry-dependent problems, such as removing excess powder because residual powder in each layer can get semi-sintered. Manzo also points out that because e-beam systems operate with thicker layers and larger beam spot-size, the features they generate may not be as fine as what can be done on a laser system.

Despite AM's challenges, notable breakthroughs are occurring at a rapid pace. At the America Makes conference, Eric Wetzel, team leader at the U.S. Army Research Laboratory, elaborated on multi-material use, describing a multi-material process for creating optical fiber, medical microtubing and other structures by 3D-printing thermal drawing forms from ABS and polycarbonate. These forms could be manipulated into small structures while maintaining complex geometries. The team was also able to turn those structures into filament.

Over at Virginia Tech, Christopher Williams, an associate professor in its department of mechanical engineering, challenges the industry to design new AM-specific polymers. He also calls for research on printing with materials such as high-performance polymers that before now have not been usable within these systems.

"Materials must be redesigned for additive. Polymers can be tuned chemically for specific processes, but you have to understand the process-structure-property relationships.

Williams adds, "All of this will require designers to start thinking differently about how materials are used." **DE**

Metal's Merits as Viable 3D Printing Material

The metal subtopic of 3D printing is hot and getting hotter. Ever since the first metal additive manufacturing (AM) systems from the former DTM (now part of 3D Systems) made 3D printing with metal materials possible, developers have worked to make variations on the technology smaller, faster, safer, less expensive and easier to manage.

Vendors have made progress on every one of these goals, but for almost three decades the systems still required operating in a highly controlled, industrial environment. Limiting factors include the following: Raw metal powders present inhalation and explosion risks, and the complete build process generally includes post-build machining and heat-treatment (i.e., adding wire-EDM machines and industrial furnaces).

Now the industry has seen metal 3D printing announcements from both Markforged (January) and Desktop Metal (May), plus a "coming soon" teaser from HP. The Desktop Metal and Markforged office-friendly AM systems avoid the concerns posed by working with classic powdered-metal materials. At a top level, both approaches use plastic-encapsulated powders and extrusion systems, and are expected to produce parts at around a tenth of the cost of laser powder bed 3D printing. What could this mean for you?

"I think we will see opportunities for these new technologies," says Cullen Hilkene, 3Diligent CEO and co-founder, "especially in industries where certain (metal part) requirements like tolerances, porosity and surface finish aren't as stringent. I think you can draw parallels between the plastics market and the metals market. In plastics AM, you have powder bed and extrusion, and both have advantages."

Markforged's process is called Atomic Diffusion Additive Manufacturing (ADAM) while Desktop Metal has developed the Bound Metal Deposition (BDM) approach.

Learn more: rapidreadytech.com/?p=11770

— *Pamela J. Waterman*

Editor's Note: *This article combines articles written by* **Brian Albright** *and* **Pam Waterman** *for* DE*'s Rapid Ready Tech blog. View the links below for more info.*

INFO:

Articles:

- **AM Industry Faces Materials Challenges:** rapidreadytech.com/?p=11726
- **NIST Tackles Metal 3D Printing Quality:** rapidreadytech.com/?p=11715
- **Metal 3D Printing, Making the Unprintable, Printable:** rapidreadytech.com/?p=11754

Companies:

- **America Makes:** americamakes.us
- **Arcam:** arcam.co.uk
- **GE:** ge.com/additive
- **HP:** www8.hp.com/us/en/printers/3d-printers.html
- **LAI International:** laico.com
- **National Institute of Standards and Technology and Physical Measurement Laboratory:** nist.gov/pml
- **RP+M:** rpplusm.com
- **Star Rapid:** starrapid.com
- **Stratasys:** stratasys.com
- **U.S. Army Research Laboratory:** arl.army.mil
- **Wohlers Associates:** wohlersassociates.com
- **3Diligent:** 3Diligent.com
- **Desktop Metal:** desktopmetal.com
- **Markforged:** markforged.com

For more information on this topic, visit digitaleng.news

Modeling for Change

Take a closer look at agile approaches that move away from being locked in to traditional final signoff.

BY TONY ABBEY

Editor's Note: *Tony Abbey develops and teaches both live and on-line FEA classes. He also provides FEA mentoring and consultancy. Contact tony@fetraining.com for details.*

MANY OF YOU WILL HAVE BEEN FACED with a situation where a design is changing rapidly and finite element analysis (FEA) simulation is struggling to keep up. Changes in geometry have a knock-on effect on the mesh, loads and boundary condition strategy. In this scenario, it is difficult to keep FEA in a proactive role in providing design direction, rather than being used in a reactive checking mode. This article reviews some ways in which FEA modeling can anticipate design changes and keep ahead of the curve.

Synchronizing FEA and CAD Models

If CAD geometry and an FEA model are linked together, then it may be possible to update the mesh, loads and boundary conditions automatically as geometry is changed.

The methodology will depend on whether the CAD geometry is driving the FEA model development or whether the FEA is working independently on imported geometry. CAD-embedded products are strongly oriented toward the former. FEA entities do not exist outside of the framework of the geometry. All FEA manipulation will flow down from geometry. The geometry can be manipulated to suit the FEA requirements, but this will be a development of the geometry parallel to the main design intent.

Traditionally, independent FEA preprocessor tools create a limited subset of geometry, but the geometry tools are at least a generation behind current CAD technologies. There is a much heavier emphasis on using the mesh to define geometric definition of the component directly.

This clear demarcation has been blurred over the past few years with new tools such as MSC Apex from MSC Software and ANSYS Workbench linked to SpaceClaim. These types of products move the design re-analysis process away from a CAD-centric to a more FEA-centric perspective, with more powerful independent geometry manipulation tools linked to meshing.

A CAD-centric Perspective

Looking first from a CAD-centric perspective, if the geometry changes are variations on a theme, then the updating can be relatively straightforward. For example, if the controlling dimensions change, but the overall topology stays constant, then surface definitions will be persistent. Most meshing, load and boundary condition definitions will be relative to surface geometry. If the surfaces are maintained, then the relationship to the FEA entities can also be maintained.

If the geometry is updated and features are modified, then the relationship of the surfaces will change. There is a good chance that the original surface definitions will be destroyed or distorted so that they are no longer applicable. This means that the task of applying loads and boundary conditions will have to be repeated.

It is always tempting to use only automatic global mesh size and other global mesh controls. It is usually very cost effective to overlay local controls over this. But this takes some dedicated effort and requires going up the learning curve. If the local controls will be disrupted every time there is a significant change, then this becomes an unattractive option.

Keeping comparisons of evolving geometry and FEA responses can be very useful. However, at some point the decision may be made to abandon this approach and instead consider the model to be a fresh configuration. This implies rethinking the FEA strategy from scratch. Broad metrics such as maximum stress and deflections are more useful comparators across significantly different designs.

An FEA-centric Perspective

With a more traditional FEA preprocessor, there is naturally a lot more control

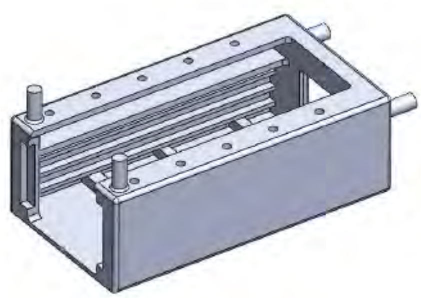

FIG. 1: The geometry model of the chassis.

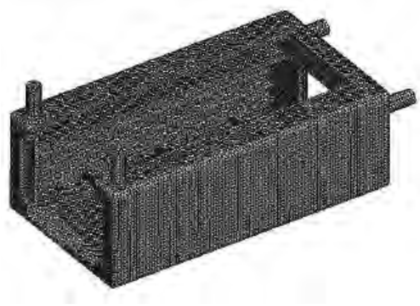

FIG. 2: The 3D element model of the chassis.

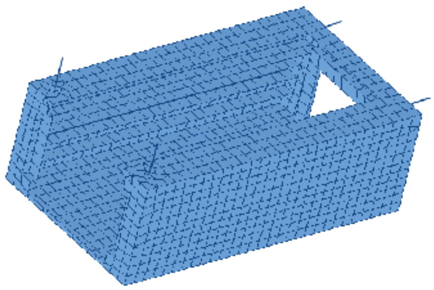

FIG. 3: 2D and 1D idealization of the chassis.

of the idealization as the design changes. The primary objective of these tools is to produce a mesh, and there are many, varied ways of achieving that. The challenge with these tools is to make them attractive to the design community, people who may not be thinking in terms of abstraction of FEA entities. As I have mentioned, there are new tools that focus on linking underlying geometry in a more powerful way to the evolving FEA model. These tools may be more attractive to the designer or part-time FEA user. The traditional dedicated FEA preprocessor uses geometry manipulation tools, which are unattractive to the designer.

The Dilemma

So, given a redesign, the big question is who will carry out the re-analysis? Is it more efficient to do this within a CAD embedded environment, or divert the task to an analysis-focused environment? The answer will depend much on the nature of the design and analysis process. I have talked to hundreds of clients about these issues. Their evolved solutions tend to have either FEA- or CAD-centric bias, but there are enormous variations within these central themes.

Using Modeling Idealizations

My background is FEA-centric, so I am familiar with driving design changes by rebuilding the FEA model. Years ago, FEA workflow originated with direct meshing with no underlying geometry. Later, 3D geometry techniques evolved ahead of CAD modelers, but they have never matured. Hence, their unattractiveness. However, dedicated preprocessors are intimately linked to the powerful concepts of idealization. Rather than modeling everything in 3D elements, there is an emphasis, where appropriate, on 2D shell and 1D beam elements. (*Editor's note: See page 13 for more on "dumb elements."*)

If using thin shell elements, a thickness variation can be achieved with one dialog box change. Similarly, beam cross-sections and their distribution can be updated rapidly.

An example of the rapid reconfiguration that can be created is shown in Fig. 1. A main chassis supports printed circuit boards and heavier trays with electromagnetic equipment. Only the main chassis is shown in Fig. 1.

The equipment had to survive a harsh vibration environment, which required normal modes, frequency response, shock spectra and random response analysis.

It was important to start scoping the dynamic response of the design to ensure the structure was in the right ballpark for meeting the loading environment. For example, if resonant frequencies strayed into critical frequency ranges of the shock spectra or the random response specification, then they would need to be redesigned.

The design was still evolving, primarily from the perspective of the electronic equipment and its distribution throughout the rack. The dimensions of the rack were also subject to variation. A full assembly had been schemed and the part count was around 300. Of this, only around 20 were structural parts or had significant mass.

The electronic components were smeared over the structure using non-structural mass. The heavier components were assumed to be very stiff. A combination of rigid spider elements and lumped mass elements were used.

Fig. 2 shows the solid element model of the main chassis. This had just under 1 million degrees of freedom. All of the structural components were modeled in a similar way and assembled using fully bonded contact. Total degree of freedom count was around 10 million. Normal modes analysis was fairly straightforward, but the dynamic response analyses became very tedious to run and check out, primarily because of the size of the model and the number of local modes.

At this point, the design started to change almost every day. To be responsive I changed tack to a 2D and 1D model. This is shown in Fig. 3. In this case I set up the walls of the tray to have a regular mesh. This formed the pattern for any variation that would be required. The change in height of the trays or PCBs just meant a change in height of that particular mesh. I also used a simple subset of the surfaces extracted from the CAD geometry to allow easy manipulation. Mass and center of gravity of components could be rapidly adjusted.

In retrospect, bonded contact could have been used more extensively to avoid the need for conforming meshes. Edge-to-surface bonded contact is now much more robust and opens up even more possibilities for fast meshing.

The initial analyses showed some worrying resonant frequency implications, particularly for the poorly supported printed circuit boards. One of the stiffer trays had a similar problem, due to the large components mounted on this. I was able to work rapidly with the designers to come up with stiffening schemes and mass redistribution schemes to alleviate the response to the shock and random response loadings.

At that point it was very tempting to start to work up the mesh into a better representation of the structure. However, this design was not frozen yet, and I wanted to wait until there was a definite decision to commit the configuration to the prime contractor. Indeed, at this point there were several significant configuration changes due to the electronic performance of equipment rather than the structural response. So, several more rapid evolutions of the resultant changes dynamic response were required. Compromises on structural response and electronic design expediency had to be carried out.

I do not have enough experience or skill working directly with CAD to be able to evaluate how quickly the changes could be affected in a full FEA/CAD embedded solution. I think the future holds a lot of potential, where the fidelity of the FEA tools and the subtlety of the CAD manipulation will be available within one program. With cross-training of analysts in a friendlier CAD tool and designers into a friendlier FEA tool, rapid and efficient working techniques will no doubt be evolving. I would also imagine that there would be many unique workflows developed, dependent on the particular talents of the engineers involved.

It is interesting to note the current emphasis on FEA democratization, putting FEA into the hands of the designer. I would also like to see CAD democratization, making CAD tools easy for analysts such as me to use! **DE**

Tony Abbey *works as training manager for NAFEMS, responsible for developing and implementing training classes, including a wide range of e-learning classes. Check out the range of courses available: nafems.org/e-learning.*

INFO → **ANSYS:** ansys.com

→ **MSC Software:** mscsoftware.com

For more information on this topic, visit digitaleng.news

Building Blocks of IoT Design

The rise of the internet of things could shift the designer's focus from system-on-a-chip to system-in-package.

BY TOM KEVAN

For the past 10 years, smartphone technology has been the primary driver of innovation in the electronics industry, because chipmakers have constantly pushed to produce smaller, denser, cheaper and more powerful electronic components. To accomplish this, PC and smartphone designers have leveraged system-on-a-chip (SoC) technology. This IC (integrated circuit) model greatly facilitates integration—addressing space constraints—and enables economies of scale by fabricating the same chip designs on common process platforms in production runs of millions of units. During this period, the advantages of cost and integration outweighed any of the SoC's shortcomings.

But this dynamic seems to be changing. Smartphone technology's footprint in the market has begun to shrink and, at the same time, internet of things (IoT) technology has taken on a more dynamic role. Research firm MarketsAndMarkets predicts that the global IoT market will grow to $561 billion by 2022, sustaining a compound annual growth rate of 26.9%.

In addition to changing market dynamics, the rise of the IoT also portends a fundamental shift in the hardware building blocks used by design engineers to develop products. But why are these changes occurring now? After all, PC and smartphone designers have long been integrating increasingly complex combinations of processing, communications and sensing components, pushing the limits of form, size, power consumption and functionality. What does this paradigm shift mean for design engineers? And what hardware building blocks will dominate the products they produce for the IoT?

The rise of the IoT brings with it a cornucopia of applications and disparate design requirements. This has triggered the development of greater chip-configuration diversity. As the IoT takes shape, a new hardware era will emerge, where design engineers will harness the power of all options to meet expanding demands for new functionality. *Image courtesy of Optimal Design.*

Catalysts for Change

The trigger for this shift lies in the fact that the design and market requirements of the smartphone and the IoT differ greatly. With the rise of the IoT, there is no longer a single semiconductor fabrication process that works for all of the components incorporated in a system. This runs counter to the parameters of SoC production, where these systems are locked into a single production process that essentially precludes using best-in-class components to build the systems. This means that one or more essential components must be compromised. The only way to implement the kind of customization required by the diversity of IoT applications is by reconfiguring the entire SoC fabrication line, which would greatly diminish one of its primary strengths—the fab's ability to deliver economies of scale.

The rise of the IoT, with its breadth of applications and disparate design requirements, is forcing product designers to seek greater chip-configuration diversity. And instead of being able to use the same chip design for millions of applications, the IoT places a higher value on customization of products manufactured in moderate-size production runs.

As a result of these shifting dynamics, chip makers have adopted different approaches to IC and component packaging and have begun to generate new chip designs. A major trend called system-in-package (SiP) combines one or more ICs of different functionality in a single package that performs as a system or subsystem.

The question then becomes: What chip models and configurations will become the building blocks of the IoT, and what roles will each of the variations play in the creation of new products? Will one approach dominate, or will the industry use a combination of approaches? And how will this affect product development teams? The answers to these questions become clearer after an examination of the strengths and weaknesses of each approach.

SoCs greatly facilitate integration by placing digital, analog, mixed-signal and RF subsystems on a single substrate. This IC model addresses space constraints and promises to speed development. But because all these elements are placed on the chip via a common manufacturing process, it is difficult to implement customization without compromising the approach's ability to leverage economies of scale. *Image courtesy of Silicon Labs.*

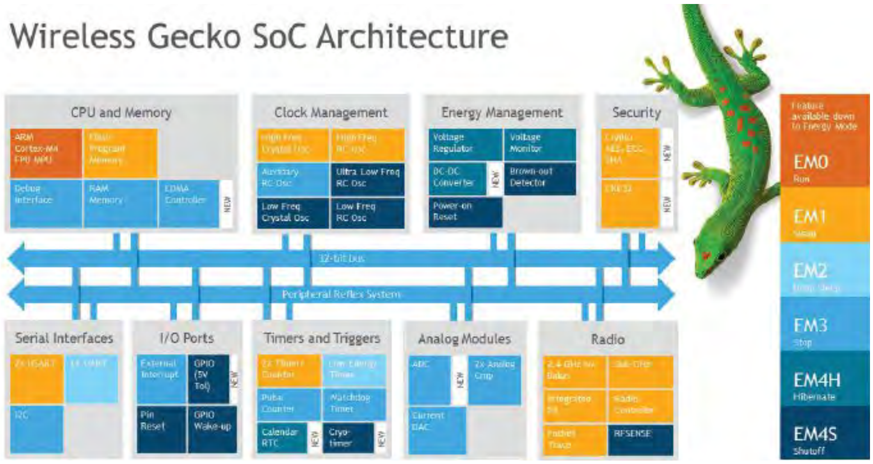

Silicon Labs' wireless Gecko system-on-chip (SoC) integrates an ARM Cortex-M4 core; 2.4 GHz radio and hardware cryptography. One of the SoC's greatest strengths is the wireless technology incorporated in the system, which includes Thread and ZigBee stacks for mesh networks, radio interface software for proprietary protocols, Bluetooth Smart for point-to-point connectivity and Simplicity Studio tools to simplify wireless development, configuration, debugging and low-energy design. *Image courtesy of Silicon Labs.*

SoCs, SiPs and Other Options

SoCs and SiPs each bring strengths and weaknesses to the table. It's important, however, that designers bear in mind that there are other options that may serve them well.

Looking at SoCs, economies of scale represent one of their greatest strengths, but there are other positive elements to consider. For example, they offer the smallest footprint and often provide the lowest power implementation.

Their drawbacks, however, are significant. "SoCs present serious development risks, such as mask errors and firmware bugs," says Todd Zielinski, senior director of electrical engineering at the Bresslergroup, a Philadelphia-based design firm. "They can also be difficult for novices to use and may be difficult to access because of minimal design support from manufacturers."

But when sizing up this chip model, engineers must also consider the technology against the backdrop of wireless design issues. "For companies with in-house RF (radio frequency) expertise and experience developing RF products, reusing R&D investments and designs based on SoCs often makes these devices a better approach," says Kamran Shah, director of corporate marketing at Silicon Labs. "Also, for any proprietary wireless design with unique antenna and other RF design needs, an SoC will generally be preferable."

Compared with these strengths and weaknesses, SiPs offer an interesting contrast. For example, SiPs often provide more capabilities for less volume of parts shipped than SoCs. This chip model also "is more accessible to the lower volume developer through normal sales channels, and has less risks in development—[because] they are usually built from a composite of well known and available parts in die form," says Zielinski. "They have slightly higher costs in production than SoCs, but [that] can be offset in development costs in lower volumes."

Designers should also be aware that SiPs are easier to customize for less money than SoCs. This directly relates to the manufacturing constraints posed by the fact that SoCs are locked into a single production process.

Considered within the context of wireless design challenges, SiP strengths can be stated by a few rules of thumb. "For wireless systems, the selection of an SoC or certified SiP module tends to come down to expected volume, time to market and in-house RF expertise," says Shah. "For fast time-to-market requirements, certified SiP modules have a significant advantage as they reduce both cost and time for regulatory certification. For companies without in-house RF expertise, certified SiP modules are also a much lower risk to embark on a project as the complexities of antenna design, PCB (printed circuit board) layout and RF testing and verification are addressed already."

After looking at the strength and weaknesses of the two technologies, it's clear that design engineers have their hands full, weighing all the considerations and tradeoffs involved. And to make matters even more complicated, there are also cases when an SoC and an SiP module can both be used in a design. For instance, a development team could use a certified SiP module as a network coprocessor device for wireless communication with an SoC.

But just how big a role can either of these technologies play in IoT product development? Is reliance on pre-built, plug-and-play SoCs and SiPs that incorporate prepackaged sensor, mixed-signal, signal processing and wireless subsystems a viable design approach?

The Plug-and-Play Option

The use of drop-in, prebuilt SoCs or certified SiP modules can be an ideal way for development teams to incorporate functionality quickly into designs. For most

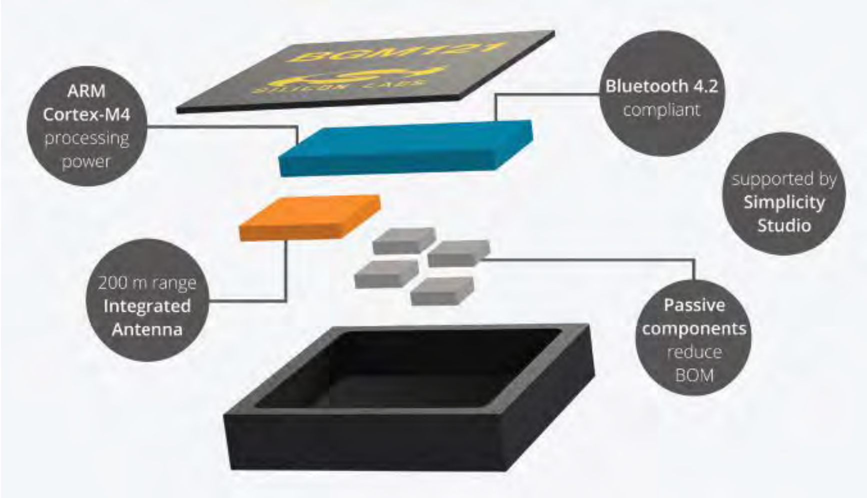

Silicon Labs' Blue Gecko BGM12x Bluetooth Smart SiP integrates a Bluetooth 4.2-compliant stack capable of running end-user applications onboard, as well as 256 KB of flash memory and 32 KB of RAM. The entire package measures 6.5x6.5x1.4 mm, making it well suited for space-constrained applications. ***Image courtesy of Silicon Labs.***

standards-based functionality, relying on commercial solutions can sometimes simplify the design process and speed development. But here's the catch: Prebuilt modules can save time, but they often increase unit cost as much as two-fold. They also limit flexibility in packaging, features and environmental constraints. The SoC-based approach also requires in-house RF expertise to ensure performance and integration requirements are met, and not all companies have these resources. Development teams adopting the SoC-based approach must address factors such as antenna design, PCB layout, component interactions, RF test equipment and worldwide regulatory approval.

Long term, for the IoT, an approach that relies solely on drop-in, prebuilt modules promises to fall short. "We need to remember that the IoT is not one application but a set of hundreds, if not thousands, of different applications," says Ron Lowman, strategic marketing manager for IoT at Synopsys. "So there is a broad range of differing specs required for the analog and RF components of these systems. This reduces the volumes and thus the potential profitability of embedding the specific RF and analog IP based on development costs. Also, the strategy to use plug-and-play SoCs or SiPs in IoT devices up to now has worked only because of these profitability hurdles."

To address this hurdle, Lowman says companies have sought to reduce their investment and maximize their capacity to get product to market faster. They do this via leveraging SiPs and plug-and-play prepackaged alternatives instead of analog and RF integration.

So for development teams, the decision to use a certified SiP module or an SoC comes down to expected volume, time-to-market requirements, aggressive design and RF requirements and available in-house expertise. Based on these conditions, it seems unlikely that developers will be able to rely solely on drop-in, prebuilt SoC and SiP modules, although they will definitely be part of the mix.

This raises the question: When would design teams be best served by specifying and combining individual subsystems (for example, microcontroller unit, memory, RF module) rather than relying on plug-and-play systems like SoCs or SiPs to handle the integration?

Building Your Own

The viability of custom development is greater for products that are expected to be in very high volume or have extremely complex and unique requirements for a market. "In the case of very high volumes, the potential cost savings in hardware with an optimized solution might outweigh the expense of hardware or RF expertise, additional test infrastructure and wireless certification," says Shah. "In the case of complex systems that may not be able to make use of standards-based technology, the only technical option may be a more custom implementation, where an SoC or certified SiP module may not meet all of the design requirements."

Specific technical challenges mandate custom development, such as when a systems form factor requires additional flexibility, as pointed out by Joe Kreidler, director of electrical and software engineering at Optimal Design, an Illinois-based product development and IoT solutions firm. "Optimal Design has found this [custom development] to be particularly the case with wearables. Even though a module may be a simpler solution from a board design and firmware perspective, a module may not be able to bend or be shaped around the curve of the body. In these cases, utilizing discrete components allows us to lay out the electronics to obtain the optimal size and shape necessary for the desired user experience."

Another instance where development teams are best served by custom development involves designs with tight power budgets. "In battery-operated devices, it's required that the subsystem be integrated to ensure extended battery life and cost reductions," says Lowman. "In a study by Microsoft, one of the largest differences from chip-to-chip power consumption comparisons was the length of time required to turn on the processor and turn off the processor. Integrating the wireless and analog components and sticking them on the AHB (advanced high performance bus) can save five to six cycles versus a serial interface connected to the plug-and-play systems."

Additional benefits include the increased time to turn on and off a second processor and second power management piece that may be on the external SoC/SiP, Lowman adds.

When is Expertise Required?

Even before development teams begin to consider what hardware building blocks to use in their IoT product design and whether to pursue a custom implementation, they have to determine if they have the right balance of expertise to carry out the project. With the complexity of the designs, ever-increasing incorporation of analog and RF subsystems, and the array of hardware components engineers have to choose from, having the necessary talent represents a critical factor.

A good rule of thumb states that the level of analog and RF expertise needed for a design project is largely a function of the type of product being delivered, the wireless communication protocol used and the level of integration of sensors and peripherals required. "A project relying on a standards-based wireless protocol, such as Bluetooth, can make use of certified modules and standardized protocol software without the need for extensive RF or analog design expertise," says Shah. "On the other hand, a company developing a proprietary wireless protocol may require optimization to meet system requirements in both software, hardware and antenna design."

Add the integration of wireless to a project developing a wearable, and the need for RF expertise grows. In this case, engineers typically do not have the luxury of using drop-in modules. Discrete components may be needed, and the device generally requires a custom antenna to meet size and performance requirements. For these devices, in-house RF expertise is essential to the success of the product.

"To build a successful product with embedded radios, RF simulation is needed from the beginning of the project," says Pete Nanni, lead RF engineer at Optimal Design. "The complete RF system consists of the module, printed circuit board and housing. An engineer cannot simply use data sheet performance specifications and radiation patterns because these patterns are measured in free space.

"When a module is built into a system, the radiation pattern changes, so simulation is key to optimizing RF performance. Many IoT devices have multiple radios, such as a Bluetooth Low Energy radio and a cellular radio," Nanni adds. "Off-the-shelf modules can be used, but the reference designs do not provide information about the interaction of the modules. An in-house RF team needs to perform simulations to understand how the radiation patterns are impacted by the multiple radios."

Challenges of a New Hardware Era

The hardware options of engineers designing products for the IoT have increased, facilitating the integration of digital, analog, mixed-signal and RF functions, and sometimes streamlining the development processes. But even with these building blocks in place, custom implementations will still be required in the most demanding of designs.

No single chip design will dominate development processes. In fact, designers will sometimes combine SoCs, SiPs and other multichip modules in a design. The decision to use a certified SiP module or an SoC will often come down to expected volume, time-to-market requirements, aggressive design and RF requirements, and available in-house expertise.

Expanded use of analog, mixed-signal and RF subsystems will place a premium on in-house expertise. Development teams with this expertise will have a greater opportunity to take advantage of the full spectrum of hardware options available to them, allowing them to better meet demands for smaller form factors, greater energy efficiency and enhanced functionality. **DE**

Tom Kevan *is a freelance writer/editor specializing in engineering and communications technology. Contact him via de-editors@digitaleng.news.*

INFO → **Bresslergroup:** bresslergroup.com

→ **MarketsAndMarkets:** marketsand markets.com

→ **Optimal Design:** optimaldesignco.com

→ **Silicon Labs:** silabs.com

→ **Synopsys:** synopsys.com

For more information on this topic, visit digitaleng.news

Each week, ***Tony Lockwood*** *combs through dozens of new products to bring you the ones he thinks will help you do your job better, smarter and faster. Here are Lockwood's most recent musings about the products that have really grabbed his attention.*

EDITOR'S PICKS

NVIDIA Helps Make Real-Time Simulation Real

NVIDIA technologies leveraged by ANSYS for new simulation tool for designers.

Discovery Live is ANSYS' new multiple physics design exploration tool for designers and other engineers who don't do simulation.

Easy to use, it includes direct-modeling tools and handles fluids as well as structural and thermal simulation jobs. Run all the what-if scenarios you want because Discovery Live gives results in near real time. Key to its performance lies in collaboration between ANSYS code and NVIDIA technologies.

MORE → digitaleng.news/de/?p=39517

Point-Cloud Converter Updated

Selectable surfaces function enables users to categorize objects within 3D scene.

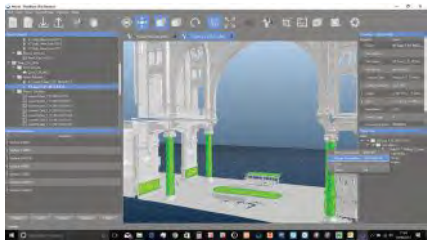

Arithmetica specializes in tackling computationally intensive jobs, evident in its release of version 3 of Pointfuse, its point-cloud conversion software. New to this version, Pointfuse V3 offers algorithmic enhancements that improve data processing performance compared with its predecessors.

In short, Pointfuse delivers a fast, automated and flexible way of converting large point-cloud data sets into high-fidelity vector models.

MORE → digitaleng.news/de/?p=39378

Short-Range Industrial Laser Scanner Debuts

FARO also unveils its next-generation FaroArm articulated arm.

Dubbed the FARO FocusS 70, this new addition to the FARO Focus Laser Scanner series melds a number of the capabilities of its sibling units, such as range and accuracy, but in a size that should be more convenient.

The skinny on it is that the FocusS 70 is a short-range, ultra-portable industrial scanner. Short range means from 1.97 to 229.65 ft. Ultra-portable means it tips the scales at about 9.25 lbs. and measures approximately 9x7x4 in.

MORE → digitaleng.news/de/?p=39276

Fusion 360 Now Has Sheet Metal

August update also sees design, sketching, simulation and CAM enhancements.

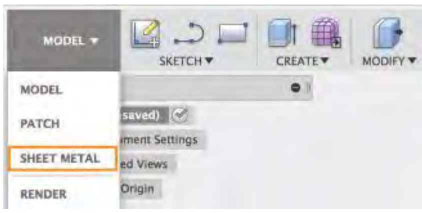

Autodesk's Fusion 360 cloud-based, subscription platform for 3D CAD, CAM and CAE adds a number of improved features. High on the list is a sheet metal design and manufacture workspace.

Features of the workspace include a single tool for making base, edge, contour and miter flanges. Your toolbox has a unfold/refold capability for checking out how features look as well as other tools for documenting and dimensioning flat patterns.

MORE → digitaleng.news/de/?p=39168

Next-Gen Engineers

Student Design Competition Profile: The Student Cluster Competition

Mastering a Real-World HPC-Driven Scientific Workload

BY JIM ROMEO

EACH YEAR, the Supercomputing Conference (SC) includes the Student Cluster Competition (SCC). It began 10 years ago as a part of the conference.

Stephen Lien Harrell, a scientific application analyst in Research Computing at Purdue University, began working with the Student Cluster Competition in 2010 as an advisor to a team from Purdue.

He has since been the lead staff advisor for seven Purdue teams that have competed in the U.S., Germany and China. His leadership has progressed over the years and he was recently named the chair of the competition in 2016 and the present Chair for 2017.

Digital Engineering: Can you provide an overview of the Student Cluster Design Competition?

Stephen Harrell: The Student Cluster Competition at SC17 will be the 11th competition that has [taken place] at the Supercomputing Conference. It began in 2007 as an exhibition and undergraduate challenge to show how easy it is to build and run a small-sized high performance computing (HPC) cluster. However, it has since turned into an exciting and educational HPC competition that immerses undergraduates in HPC as a concept and a career both as practitioners and scientists. Over the years [it] has spawned similar events in China, Germany and South Africa. One thing that differentiates the SC17 conference version is that the teams compete around the clock, while the others shut down overnight.

Working with hardware and software vendor partners, student teams design and build small clusters, learn designated scientific applications [and] apply optimization techniques for their chosen architectures. Then they compete in a nonstop, 48-hour challenge at the SC conference, striving to complete a real-world scientific workload while impressing conference attendees and interview judges with their HPC knowledge. However, there is a twist: all HPC clusters must fit within 3,000 watts.

DE: Cn you tell us about some of the designs that are part of the event?

Harrell: We expect the majority of designs to be centered around three major architectures: IBM POWER8 with NVIDIA (P/V)100s, Intel Skylake with NVIDIA (P/V)100s and Intel's Knights Landing. In addition we expect both an AMD Ryzen with Radeon accelerators and a system using ARM processors.

The designs are typically put together in collaboration with the students on the team, the team advisor and a vendor sponsor. Each team writes a proposal for participation that includes an architecture section and vendor selection.

Once a team is selected, they have until October to finalize their architecture. During this intervening time, students and advisors benchmark and profile the applications and talk to their vendor(s) about what the optimal architecture would be.

DE: Can you provide examples of what the event has produced?

Harrell: Our goal is to inspire students to join us in the HPC community. HPC is a large and varied field that lives on the leading edges of technology and science and needs smart and driven students to sustain it going forward. We have had success and regularly see our past competitors participating in the HPC community. Last year at SC16 we celebrated our 10th competition and invited past competitors to share the impact of the SCC on them.

DE: Who sponsors the program?

Harrell: Our supporters this year are SAIC (Science Applications International Corporation), Microsoft Azure, Geist and Alinea, now ARM. **DE**

Jim Romeo *(JimRomeo.net) is a freelance writer based in Chesapeake, Va.*

Student Cluster Competition (SCC) teams and advisors come from all areas of HPC. *Image courtesy of SCC.*

FAST APPS
Engineering Case Studies

Front-End Loader Linkage Durability Analysis Using Load Input from True-Load

CAD model of vertical lift loader assembly on chassis.

BY DAVID N. SLOWINSKI, CNH INDUSTRIAL

Finite element analysis (FEA) is a useful tool for the evaluation of the structural components of construction equipment. A common approach is to apply static loads to the finite element models to represent loading under normal field operations. The static loads applied often represent "worst-case" forces under those operations. The structure under investigation is then determined to be acceptable or not based on meeting target values for stress, strain, deflection and or some other criteria.

One concern with this approach is that it may miss potential fatigue issues. Estimated fatigue life for a given load can be determined using stress-life or strain-life calculations. However, the dynamic nature of the loading is difficult to reproduce. A single field operation may require numerous static load cases to simulate the complete range of forces applied, and cycle counts are difficult to estimate.

Another problem with representing dynamic loading with worst-case static loads is that it neglects the relatively low strain, but high cycle loads that can result in fatigue problems. Machines such as skid steer loaders and compact track loaders perform a wide variety of tasks that often involve highly dynamic loading. Accurately reproducing these dynamic loads in the finite element simulations greatly improves the confidence in the results of the evaluation and allows the analyst to compare design alternatives effectively.

The finite element modeling capabilities available in Abaqus/CAE allow the analyst to create simulation models that provide accurate response to the applied boundary conditions and loading. The key to using this technology for the fatigue analysis of construction equipment is applying loading that reflects the dynamics of the field operations.

True-Load software accurately duplicates time-history field loads and has a user interface that is directly integrated with Abaqus/CAE. A unit load model, in which any possible load path into the structure is represented by a unit load, is first solved in Abaqus to determine strain response to the unit loads. True-Load/Pre-Test then uses the strain response to determine optimal strain gauge locations that effectively turn the structure into its own load transducer. Strains are then recorded at the True-Load strain gauges during field test. The test strain data is then read into True-Load/Post-Test to determine the time history load functions by calculating the scaled unit loads that recreate the strains produced during the test.

Once the time-history loading is determined, the full field response of the structure is determined by applying the scaled loading in True-QSE, which superimposes the time-history loading on the unit load model. The time-varying stress, strains and deflections are determined at each node and element of the structure. True-Load/Post-Test also has built-in functionality to automatically output the time series loading in fe-safe format so that fatigue analysis of the structure can be completed.

In the case study presented in this paper, design alternatives for a front-end loader linkage were evaluated by the linear static FEA approach as described already and by an approach utilizing True-Load software to recreate field loading on the Abaqus FE model. Once accurate field data were applied to the model, Verity weld fatigue analysis was used to validate a design that met the simulation durability target.

MORE ➔ digitaleng.news/de/?p=39991

➔ For the full application stories, visit digitaleng.news/de/fastapps

AD INDEX

| **COMMENTARY** |

By Steve Chalgren

IoT Can Supercharge Engineering Development

THE INTERNET OF THINGS has the potential to create many headaches for product teams, but enabling the R&D team to see aggregated user behavior can be a big gain. Engineers could get answers to all kinds of questions about how a product performs in the hands of customers with everyday use, further enhancing product development in ways not considered in the past.

For example, with a coffee maker connected to the internet (which exists), engineers could determine, among other things:

- the ambient temperature of the room in which the coffee maker is stored, and
- how well various components are performing.

Information of this nature might lead to enhancements such as a more ruggedized design—a switch that can better withstand damp conditions, for instance—or result in the creation of a new coffee maker designed to make perfect coffee in locations with more variable environmental conditions.

In other industries, automated aggregated customer behavior analysis isn't new. In 1992, Walmart started collecting point of sale (POS) information from all their stores into one giant (at the time) data warehouse, creating a significant competitive differentiation that enabled it to disrupt and dominate retail.

In the early 2000s, cloud enterprise software entered the market. This is where I work. Today, we see anonymized, aggregated user behavior from all our customers. We can see product failures instantly and launch a non-conformance quality process to fix the problem quickly. We can target our customer research, improvement plans and R&D investment to continuously improve product for maximum customer benefit. Being this close to every customer creates a significant competitive advantage over traditional on-premise software solutions, which can only capture product usage from one customer at a time.

Turning Problems into Competitive Advantages

In the super-connected world of IoT, this wealth of new information doesn't mean that customer surveys and focus groups disappear. But, the IoT product company will know where problems are; the team will then be interviewing customers to get the context and verify understanding to make the next product better. In fact, in many cases, the IoT user data will uncover more problems to research—a good thing. These questions lead to insights that turn out to be competitive game-changers.

Even the process of identifying interesting questions will become automated, as organizations apply sophisticated analytics to product-captured IoT data. You can begin benefitting today by creating meaningful data flows to the R&D team. For instance, a team could create an analytic in its IoT product data warehouse that triggers when something fails. The trigger creates a quality review process in the product lifecycle management/quality management system with key data from the unit(s) for investigation. The R&D team solves the problem and releases an update. The release fixes millions of units automatically, even before the customers realize there was a problem.

To enjoy this engineering feedback loop, however, we need to think hard about two things before implementation: privacy and security. The medical device industry learned this lesson painfully when news broke that pacemakers were frighteningly easy to hack. Today, most medical devices are encrypted and locked down. We need to ensure that the information we collect is anonymized, encrypted and carefully controlled to protect the privacy of customers. Continual breaches in data security and control (e.g., Equifax's latest cybersecurity scandal, self-driving car vulnerabilities) demonstrate that while technology is a great enabler, benefits come with costs as hackers revel in the challenge every new data set and connected device brings.

Once this IoT engineering feedback loop is in place, however, it will become a major force for market disruption. Today, competition in most markets is predictable because companies can easily copy each other. However, the action one company takes in connecting its products and becoming the first-to-IoT market player can give it a formidable market position. As that company looks at the data coming in, it will discover previously unknown opportunities. Others will take notice and try to copy, but by the time they've done so, the original disruptor will have a new generation in place that takes advantage of the latest batch of data that keeps rolling in. It's a virtuous feedback loop that makes it difficult for others to achieve parity, all else being equal.

The lesson for product organizations? Incorporate IoT early and take advantage of the data that flows back as soon as possible to solve problems, improve products and innovate. The longer you wait, the harder it will be to catch up to your competitors who adopted IoT before you did. DE

Steve Chalgren *is executive vice president of Product Management & chief strategy officer at Arena Solutions (ArenaSolutions.com), which provides a product realization platform that enables companies to design, produce, and improve their innovative products quickly.*