March 2020

DigitalEngineering247.com

DE Digital Engineering

Confronting Computing Complexities

- GPU-OPTIMIZED SIMULATION
- LARGE MODEL CHALLENGES
- HPC AND THE DIGITAL THREAD

Peerless MEDIA, LLC

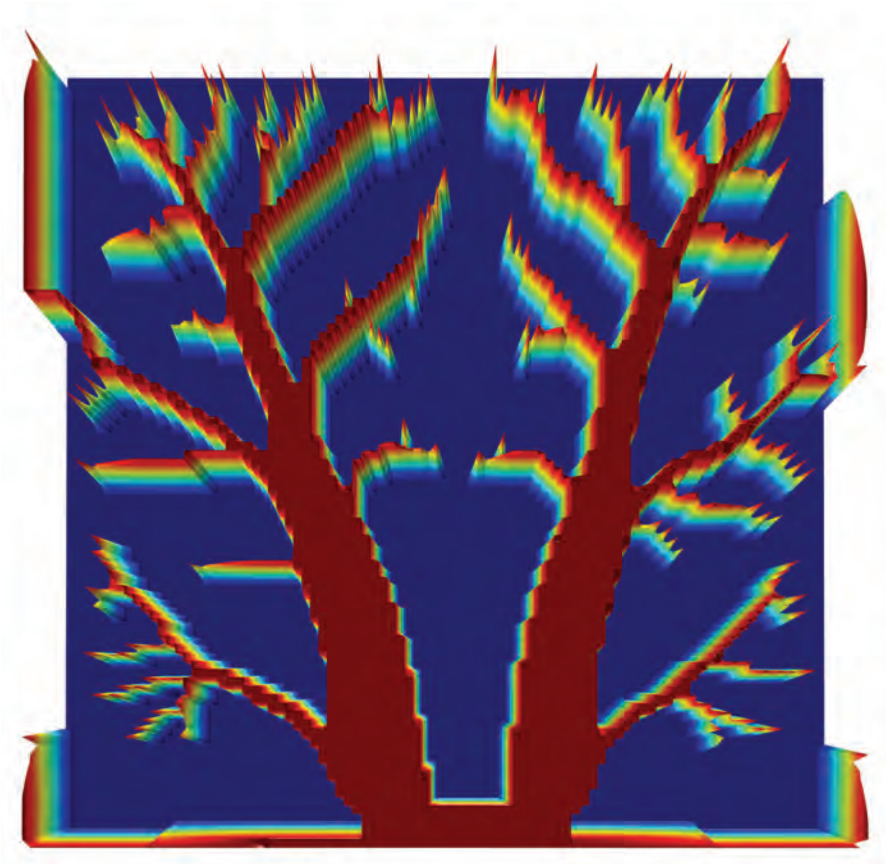

Topology optimization of a heat sink.

By Brian Albright

The Paradox of Choice

HIGH-PERFORMANCE COMPUTING (HPC) IS A KEY ENABLING TECHNOLOGY for advanced design and simulation. This issue, we're looking at how HPC, as well as cloud infrastructure and faster graphics processing units at the workstation level, are affecting every aspect of design. While that's our focus this month, HPC is far from the only technology that has opened up the design space to a nearly limitless field of possibilities.

Thanks to the convergence of additive manufacturing technology, generative design software, more user-friendly simulation tools and access to nearly unlimited HPC resources and cloud-based infrastructure, the sky really is the limit when it comes to design possibilities.

Artificial intelligence and generative design technology can use pre-loaded parameters to devise design options that no engineer could have imagined. Often these designs are wild and organic looking, like a product that was sprouted rather than built. Armed with new optimization tools and faster workstations, designers can also rapidly iterate those concepts.

But just because we can push the envelope on design doesn't mean that's always the best idea. In fact, too many options isn't necessarily a great thing. Psychologists tell us that the human brain can really only successfully manage around a half-dozen options at once. In engineering, technology helps whittle down the field based on whatever constraints we've placed on the design. Although these systems can optimize a design, they don't necessarily ensure we've selected the *right* design.

I thought about this a lot after Elon Musk revealed the Tesla Cybertruck, an electric pickup that was designed (at least partly) with input from several thousand of Musk's Twitter followers. The $40,000 truck looks like it was built for an old post-apocalyptic sci-fi film like "Logan's Run" or "Damnation Alley."

While the Cybertruck has some impressive stats when it comes to horsepower, towing and payload capacity, it … isn't really a pickup truck. I grew up in a rural community, where pickup truck brands run a close second to sports teams when it comes to fierce dedication and loyalty. For people who actually need and use pickups, truck preferences are handed down through families like favorite recipes and religious affiliations. These are not drivers who are going to be interested in anything that looks like the Cybertruck. When Ford finally releases the electric version of the F150, expect it to look pretty much like its gas-powered predecessor.

Of course, these aren't likely the buyers that Tesla is courting, any more than the company's cars are targeted at people currently driving 10-year-old Toyota Camrys. The reaction to the truck, though, is a good indicator that there are limits to what certain markets will accept.

I think that points to an important role that designers and engineers play, particularly as their software tools incorporate more AI, and as 3D printing enables different approaches to design. That role is to function as an advocate for potential end users. What does a real person want out of this product? How will they use it?

Electric vehicles are an excellent example of how this could play out. EVs aren't beholden to the same design requirements as cars with gas-powered engines and drivetrains. Take those design restraints away, and an EV could potentially be designed to not look much like a traditional vehicle at all. Most of them still do, however, because there is a limit to how far away from that standard design consumers are willing to stray. There are other considerations, too, like the size of parking spaces, accommodating infant car seats, ergonomics, safety and dozens of others.

Technology can optimize a design. Designers and engineers will increasingly have to set analytics aside and apply their more subjective human judgment to making sure those designs are right for their applications and their target audience. **DE**

Brian Albright *is editorial director of* Digital Engineering. *Contact him via balbright@digitaleng.news.*

DE
Digital Engineering
TECHNOLOGY FOR OPTIMAL ENGINEERING DESIGN
March 2020
Vol.26 • No.2

INSIDE

COVER STORY

FOCUS ON HPC

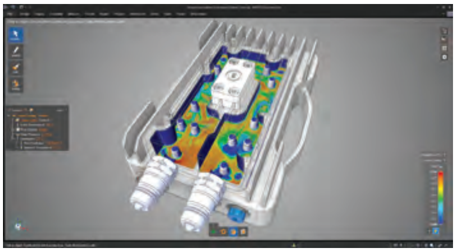

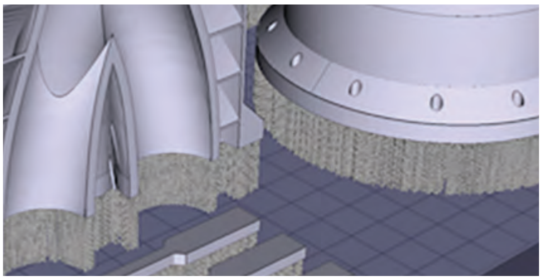

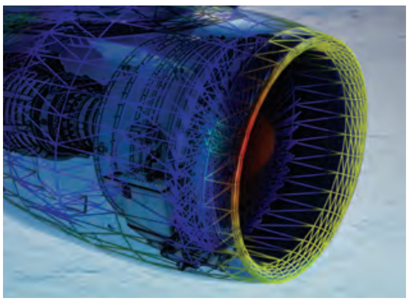

DigitalEngineering247.com

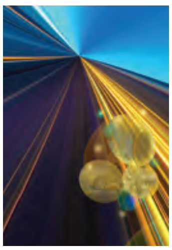

ON THE COVER: Designed with Getty Images and supplied images

FEATURES

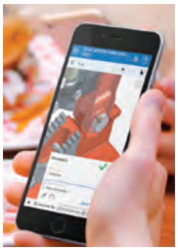

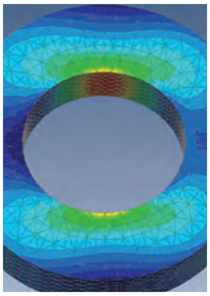

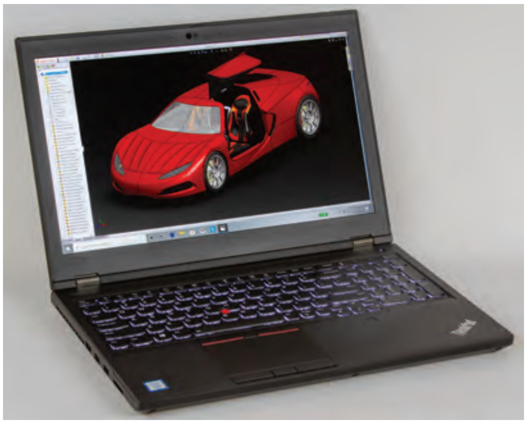

DEPARTMENTS

Editor's Pick, Page 40

PUBLISHER
Tom Cooney

EDITORIAL
Brian Albright | Editorial Director
Kenneth Wong | Senior Editor
Stephanie Skernivitz | Associate Editor
Jess Lulka | Copy Editor

CONTRIBUTING EDITORS
Tony Abbey, David S. Cohn, Kip Hanson, Tom Kevan, Randall Newton, Beth Stackpole

ADVERTISING SALES
Len Pettek | Western U.S. Regional Sales Manager
Phone: 805-493-8297
lpettek@digitaleng.news

Tom Cooney | Eastern U.S. and International Sales Manager
Phone: 973-214-6798
tcooney@digitaleng.news

CREATIVE SERVICES
Wendy DelCampo | Senior Art Director
Polly Chevalier | Art Director
Kelly Jones | Production Director

A PEERLESS MEDIA, LLC PUBLICATION
Brian Ceraolo | President & CEO

EDITORIAL OFFICES
Peerless Media, LLC
111 Speen St., Suite 200,
Framingham, MA 01701
Phone: 508-663-1500
de-editors@digitaleng.news
www.DigitalEngineering247.com

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

Peerless MEDIA, LLC
peerlessmedia.com

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.

BY THE NUMBERS

High-Performance Computing

The World's Fastest Supercomputers

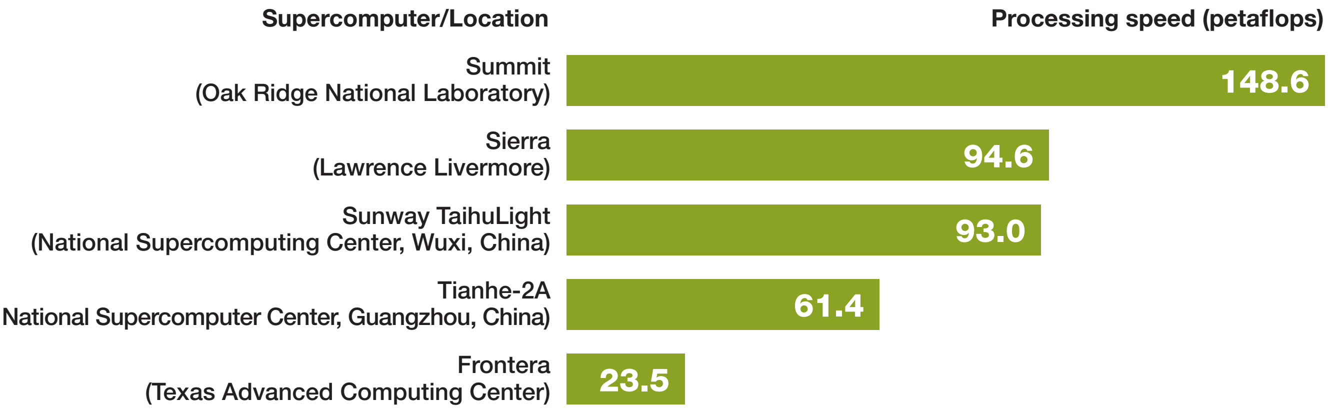

Source: TOP500, November 2019

$125 billion

The size of the computer graphics hardware market in **2018**.
The market is expected to grow to **$133 billion** by **2022**.

Source: Jon Peddie Research, July 2019

Total Computer Graphics Market ($ billion)

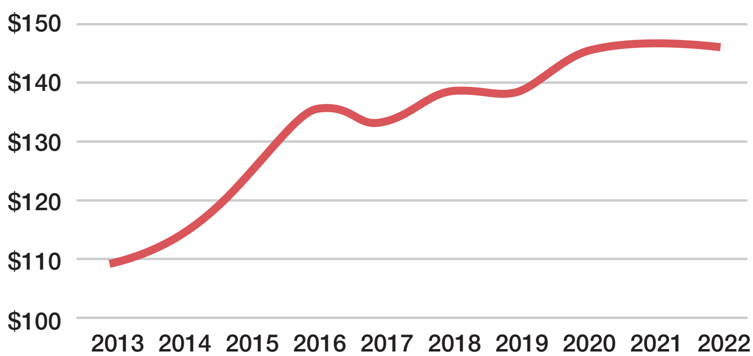

Source: Jon Peddie Research, July 2019

$4 billion

Hyperion Research expected cloud spending on high-performance computing work to reach **$4 billion** by the end of **2019,** up from **$2.5 billion** in **2018.** The company expects HPC cloud revenue to reach **$7.5 billion** in **2023.**

The estimated compound annual growth rate of computer-aided engineering (CAE) applications utilizing HPC in the cloud from **2018** to **2023.**

24.7%

Source: Hyperion Research, June 2019

$13.7 billion

The size of the HPC server market in **2018.**

Source: Hyperion Research, June 2019

$20 billion

The size of the HPC server market by **2023.**

Source: Hyperion Research, June 2019

Continued CAD Growth

According to a report from Jon Peddie Research (JPR), the CAD software market grew **9%** from **2018** to **2019,** and is expected to reach **$9 billion** (U.S.) in **2022.** JPR estimates the market will have a growth rate of more than **2.4%** over the forecast period of **2018** to **2022.**

Total CAD Revenue (worldwide)

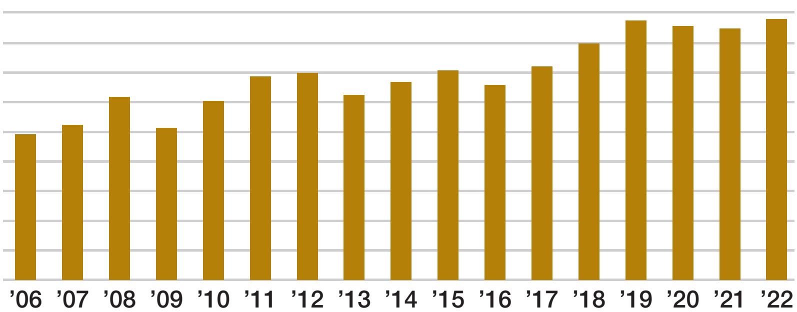

Source: JPR Worldwide CAD Market, January 2020

ROAD TRIP

'Crossing the Chasm' Author to Keynote CAASE20

Geoffrey Moore Will Discuss Disruptive Innovation at the Indianapolis Event, June 16-18

BY BRIAN ALBRIGHT

Innovative disruption is a fact of life for design engineers, given the push to more rapidly produce and iterate new designs. Simulation and design software vendors are also facing multiple disruptions with the advent of cloud-based software and flexible licensing models, as well as a push for a greater democratization of simulation systems.

That's why author and speaker Geoffrey Moore (www.geoffreyamoore.com) is such a good choice to present one of the keynote sessions at the upcoming CAASE20 (the Conference on Advancing Analysis & Simulation in Engineering) event in Indianapolis, June 16-18, 2020. The event will be held at the Indiana Convention Center. The conference is presented by NAFEMS Americas and *Digital Engineering*.

CAASE20 brings together the leading visionaries, developers and practitioners of CAE-related technologies in an open forum to share experiences, discuss relevant trends, discover common themes and explore future issues. The conference will cover a wide range of topics, including every aspect of engineering analysis and simulation.

Author Geoffrey Moore will keynote CAASE20. *Image courtesy of Geoffrey Moore.*

Moore is the author of "Crossing the Chasm," a best-selling book (most recently revised in 2014) focused on disruptive innovation and the "chasm" that exists between early adopters of new products and more widespread adoption. His most recent book is "Zone to Win: Organizing to Compete in an Age of Disruption" (2015). Though his first book focused on start-up companies, the most recent addresses issues that large enterprises face when embracing disruptive innovations, even when it is in their best interests to do so.

The Technology Adoption Lifecycle

According to Moore, there are four basic stages of adoption of new technology that can affect how a company markets new technology. The first stage involves early adopters. "They believe what you believe," Moore says. "They get what you are doing, they know it's early and they will have to do extra work, but they will buy in. They are fun to have as customers, because they are on your side of the table."

But once early adopters have embraced a product, that is when companies often face a chasm before reaching the rest of the market. How can you convince these other customers, who aren't "true believers," to get on the bandwagon?

"You target a niche of customers who are in pain," Moore says. "These are the pragmatists. They can't solve their problems with conventional solutions, and they are under increasing pressure to improve. They look to technologies, and although they don't believe what you believe, they think they need what you have. They will take a chance even if they aren't 100% sure."

If the solution is successful, these pragmatists provide valuable references for other customers and potential partners.

Next is what Moore calls the "bowling alley" stage, where companies move across adjacent niche markets. "The simulation market is likely what we call a 'bowling alley forever' market because it is so specialized and so technical," Moore says. "It's not a horizontal application for average people on the street, but there is clear value in specific markets.

The next phase is referred to as the "tornado," when customers begin adopting the technology just to play catch-up to the early adopters. "They may not believe what you believe, but they want what those companies have," Moore says. "That creates huge spikes in demand that drive market caps in technology sectors through the roof. All of a sudden, a huge amount of budget comes out of nowhere."

This phase is when the technology providers try to grab as much market share as possible to become dominant in a particular space. After that comes the "main street" phase, when customers begin to need the solution to stay competitive.

"Each of these four phases requires a different go-to-market strategy, a different set of positioning ideas, and that is key to crossing the chasm," Moore says.

According to Moore, engineers tend to be biased to the front of the lifecyle because they are closest to these new innovations. "That's where it helps to have this framework," Moore says. "If you are beating your head against the wall trying to expand the market, you might be using the wrong play."

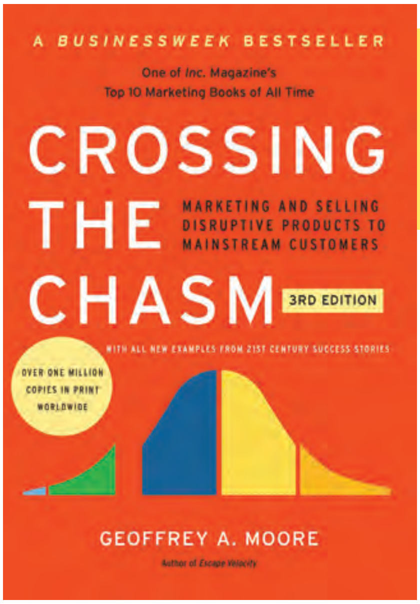

During CAASE20 in Indianapolis, Moore will discuss some of the disruptive innovation concepts from his book, "Crossing the Chasm." *Image courtesy of Geoffrey Moore.*

From Product-Focused to Consumer-Focused

The biggest change in the current market, Moore says, is that the product is no longer king. "Before, the company sold a product, and when the consumer bought it they owned it, and it was their problem to get value out of it," Moore says. "In the 21st century, the customer has more choices, and consumer power has trumped product power, because there is more supply than demand. When the customer is the scarce ingredient, you have to design backward from the customer base."

This could be a critical notion as the simulation industry searches for ways to expand its footprint beyond specialist engineers. "It's not that simulation solutions weren't good enough," Moore says. "It's that use cases require more generalists to be able to interface with the technology. The power has to be in the hands of the many instead of the few. That puts demand on the engineering team that they are not used to."

The products that engineers design and the software solutions that vendors are providing to those engineers have to be designed from a user point of view. "Just dumbing it down doesn't work," Moore says. "You have to pay attention to use-case based design. The tool has to be designed in the context of a use case that will pay off that tool."

To register for the event, visit the CAASE20 and NAFEMS website at www.nafems.org. **DE**

CAASE20 Highlights

The CAASE20 conference will feature a number of keynote sessions, including:

- **Geoffrey Moore**, consultant and author of "Crossing the Chasm."
- **Amir Husain**, inventor and author of "The Sentient Machine," as well as the CEO of AI company SparkCognition.
- **Dr. Maria Klawe**, president of Harvey Mudd College, who will discuss increasing diversity in the STEM workforce.
- **Monica Schnitger**, president and principal analyst at Schnitger Corp. and occasional *DE* columnist, who will discuss the evolution of simulation technology.
- **Scott Leemans**, creative engineering generalist at X.Company (The Moonshot Factory), who will discuss democratizing technology at the K-12 and college levels.
- **Bill Pappas**, VP of competition & engineering at IndyCar.

In addition, CAASE20 will feature live training courses on site, including:

- Nonlinear Analysis
- Dynamic Finite Element Analysis
- Structural Optimization—Topology Optimization and Generative Design
- Fatigue Analysis
- Joints and Connections
- CFD for Structural Designers and Analysts
- Introduction to Practical CFD
- Elements of Turbulence Modelling

www.nafems.org/events/nafems/2020/caase20

NEWS

CES SURVEY

Consumers Are Hungry for Personalized Products

Survey uncovers increasing appetite for product personalization, raising the stakes for engineers.

BY BETH STACKPOLE

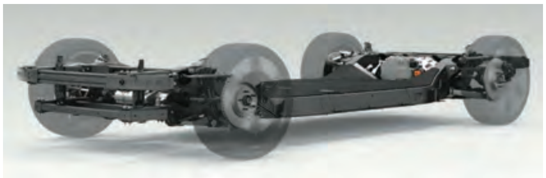

Electric vehicle startup Canoo is shaping its entire business model around personalization. *Image courtesy of Canoo.*

Driven by the preferences of a younger generation, consumers are leading the push for more personalized products, putting pressure on design teams to deliver the goods while maintaining speed and flexibility in their engineering processes.

Dassault Systèmes, in partnership with CITE Research, recently published the results of its CES 2020 survey on personalization (bit.ly/38PDule), canvassing 3,000 consumers in the United States, China and France to inquire about trends in personalization across the healthcare, mobility and retail segments as well as home and city environments.

The findings revealed a growing desire among consumers for personalization in products—one-quarter of respondents, on average, confirmed they were willing to pay more for goods tailored to their specific needs. Yet consumers do expect something in exchange for their personalization data—specifically, an average 25.6% savings.

There's more: While consumers are definitely willing to shell out the big bucks, they have no interest in any extra wait time to get their hands on customized goods. The survey found that the majority (83%) expect customized product deliverables in a matter of moments or minutes. Only 21% of those surveyed said they would wait four days or more for delivery of a personal product, and 20% aren't willing to wait any extra time at all for personalized items to come their way.

Canoo's modular "skateboard" is a universal underbody upon which many different vehicle "cabins" can sit. *Image courtesy of Dassault Systèmes.*

Consumers' increasing demand for personalization capabilities, coupled with their refusal to incur any extra wait time for delivery, sets up a major design challenge for engineers. "Demand for personalization requires a product design process that's more nimble and responsive to consumer needs," explains Olivier Sappin, CEO, CATIA, at Dassault Systèmes. "To keep up with consumer expectations, a business needs to produce and deliver personalized products with almost the same speed as an item 'off the rack.'"

Canoo, an electronic vehicle startup, is one Dassault Systèmes customer leaning on the 3DEXPERIENCE platform to create a nimble design process tuned for personalization. The company has architected its business model to cater to users who want customized transportation options.

To do so, Canoo built a modular "skateboard," which serves as a universal vehicle underbody upon which many different vehicle "cabins" can sit. The vehicle chassis, which functions as a blank canvas, will allow its customers to build any transportation experience on top.

"It could be a food truck, a rolling pop-up retail store, or a mobile medical center," Sappin explains. "By selling the mechanical foundation of an electric vehicle, Canoo is enabling other vehicle manufacturers to meet the personalization demands of their end users."

To effectively deliver personalized products, companies like Canoo and others need to embrace new technologies and skill sets to up their product development game. Competency in Internet of Things (IoT) capabilities as it relates to design is critical as is experience with additive manufacturing, which can provide a clear advantage for producing short runs of customized products.

"Businesses will need to invest in the composite materials to meet their item requirements and the 3D printing infrastructure needed to produce custom items at the desired production scale," Sappin says. **DE**

DESIGN

Radica Software Releases Electrical CAD Software

Electra Cloud provides anywhere access for designers via workstations, smartphones.

BY BRIAN ALBRIGHT

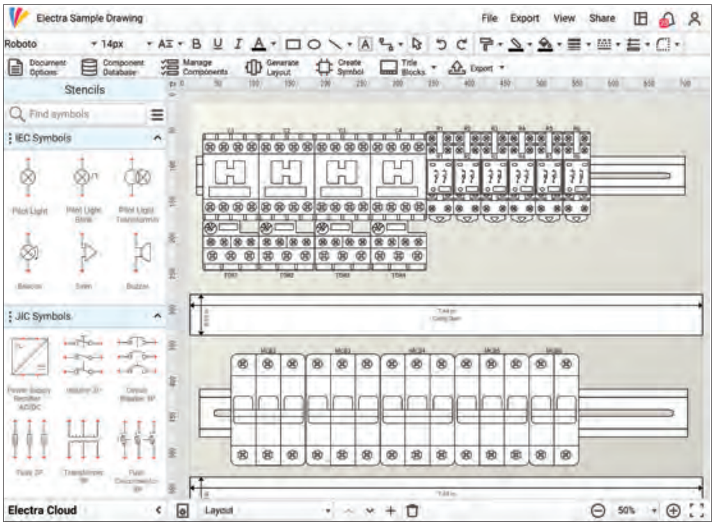

Electra Cloud from Radica Software is a browser-based electrical cloud tool. *Image courtesy of Radica Software.*

Radica Software, based in Malaysia, has released a fully browser-based electrical CAD software, Electra Cloud, that the company says will allow engineers to create, edit and share 2D electrical, pneumatic and hydraulic circuit drawings with team members.

The software works on all computer operating systems including Microsoft Windows, Linux and Apple Macintosh, as well as on any mobile devices such as smartphones. The company claims that the solution will help engineers complete circuit drawings between 300% and 500% faster than traditional software currently available.

Because it is cloud-based, the software can help users collaborate on their drawings with team members located anywhere in the world in real time, without requiring a software download.

"Software in the electrical CAD industry has not seen many changes," says Thomas Yip, CEO and founder of Radica. "So, we are thrilled to offer Electra Cloud, a breakthrough solution for engineers around the world. We're saying this because Electra Cloud can be accessed from anywhere using a web browser and users no longer need to have software installed on their computer to access electrical drawings like previously. This sets Electra Cloud apart from ECAD products in the market."

According to Yip, the cloud-based software has been in development for several years. "This was not an easy decision, because we are a small company, but we know that the cloud is where the market is heading," he says. "The product is fully browser based and available for all platforms, and unlike our competitors, because we don't have a lot of technology baggage in our software, we could jump on this right away."

Although customers have concerns about cloud reliability and security, Yip says that the majority of customers are embracing the cloud. "The majority of our beta testers are embracing the cloud positively, and they have told us that this is a far better experience," Yip says.

The Team Dashboard in Electra Cloud allows team members working on the same project to always have access to the latest version of drawings. It enables easy sharing, commenting, chatting and collaborative editing.

"With the Team Dashboard, you can invite colleagues or managers to view all of the drawings, and each drawing can be customized for sharing," Yip says. "It makes everything easy and accessible. Collaboration is the main value proposition here."

In addition, Electra Cloud also includes a File Version Control feature that saves changes made to the drawings in real time, which allows collaborators to revert to previous revisions without worrying losing any work.

Yip says the solution can also easily integrate with other systems in order to share data for simulation, digital printing and other applications.

The software features automated tooling including automatic tagging, circuit reuse, real-time cross reference, an extensive set of NFPA/JIC and IEC symbols, PLC modules, the ability to automatically generate bill of materials, terminals, panel layout and others items.

According to Yip, over the next year users of Electra Cloud can expect to see improvements that include a full-fledged component database, and capabilities that will allow the solution to serve as a central place for the engineering team to share specifications, PDFs, Google documents and spreadsheets.

"It is a great concept and it is great that you then have access from anywhere," said Joseph Stieha, an Electra Cloud Beta user from PowerTech Water Inc., based in Lexington, KY. The company has more than 150 beta customers currently testing the system.

Electra Cloud is now available directly at cloud.radicasoftware.com. It is available for free for 30 days, after which customers can use the system for $99 (U.S.) per month or $999 annually for Electra Standard (without hydraulic and pneumatic tools). The Electra Professional version (with hydraulic and pneumatic tools) is available for $1,499 per year.

"We also have a unique licensing option," Yip says. "Only users are required to license the software to use all of the tools. Marketing staff and managers do not require licensing, but can still view files and have permission to comment and approve. That can lower the total cost of ownership."

Radica Software also plans to release an offline version of Electra Cloud along with the ability to use the solution in a private cloud. DE

Large-Scale Models Take A Bite Out of Engineering Productivity

Hardware vendors and design tool makers are stepping up with solutions aimed at optimizing workstation performance when dealing with large-scale modeling.

BY BETH STACKPOLE

Engineers are building larger multidisciplinary models to flex their innovation muscle. Although these complex, high-fidelity 3D models are crucial to bringing sophisticated products to market faster, they can wreak havoc on workflows and productivity by grinding workstation performance to a halt.

Time-to-market pressures, surging consumer demand for highly personalized and intelligent offerings, as well as software and electronics content commanding a larger portion of product real estate are just some of the factors driving the need for larger and more complex 3D models. An increasing embrace of systems modeling concepts along with pervasive simulation use, including new modalities and more widespread adoption among a broader audience, is also boosting model fidelity to a point where it can be taxing for older workstations to maintain effective performance.

"Customers are trying to add more value for their customers and many are doing that by moving to system-level design," says Jon den Hartog, director of product management at Autodesk. "That means the scope of what they're modeling increases as a result. At the same [time], the scope is increasing because they are trying to create a more accurate digital representation of what the design is before they build it."

This means engineering teams newly empowered by multidisciplinary simulation and systems engineering workflows are also taking a productivity hit unless they recalibrate what constitutes an optimal workstation configuration.

Lenovo's P53 mobile workstation claims to be the first to market with NVIDIA RTX 4000 and 5000 GPUs.
Image courtesy of Lenovo.

"Without enough horsepower, it affects their quality of life working with a massive model," den Hartog says. "Every change will require a significant amount of time to calculate and propagate the math throughout the model."

Snow and sport bike rendered in SolidWorks Visualize 2020 with AMD RadeonW5700 GPU and Radeon Prorender. *Model courtesy of Adi Pandžic, CGP Design.*

Hardware Advances Drive Model Performance

Just as 3D models are gaining in complexity, advances on the hardware front are continuing to help engineering organizations solve the large-scale model performance problem. New solid-state storage options, CPUs with faster clock speeds and more powerful graphics processing units (GPUs) are just some of the hardware advances being integrated into next-generation engineering workstations to help with large-scale 3D model management and processing.

GPUs, in particular, are a technology bright spot that help boost workstation performance for complex CAD and CAE models. NVIDIA has been leading the charge here, backstopped more recently by activity with silicon leader AMD.

NVIDIA has a full family of GPUs, but its Quadro RTX family (powered by the NVIDIA Turing architecture) sets a new bar. The Quadro RTX line integrates RT Cores, accelerator units dedicated to performing ray tracing operations with high-level efficiency, along with high-end memory and artificial intelligence capabilities.

The architecture and core combination is designed to optimize performance of sophisticated applications like virtual reality, ray tracing, photorealistic rendering and simulation, all of which require massive compute horsepower and real-time performance, according to Andrew Rink, NVIDIA's head of marketing strategy.

The choice of RTX platform depends on the use case—the RTX Quadro 4000 hits the sweet spot for engineers immersed in photorealistic ray tracing applications, while the higher-end Quadro RTX 8000, which is equipped with 48GB of GPU memory and ability to pair two GPUs to double system memory and performance, is the highest end option.

"It really depends on the type of workflow you have and whether you're working on simple parts or complex large assemblies," Rink explains. GPU horsepower can also be targeted to help huge models load faster. "If it takes 10 seconds for a rotation to happen because the system is sputtering as you are turning a model, that is pretty frustrating and is a hit on productivity."

For its part, AMD is boosting CAD and simulation performance via multitasking and multithreading capabilities that come via more CPU cores. AMD is also busy advancing the range of its GPU line.

Last November, the firm announced the ADM Radeon Pro W5700, a professional PC workstation graphics card that features the high-performance, energy-efficient AMD Radeon DNA (RDNA) architecture and state-of-the-art GDDR6 memory specifically tuned for handling large models and datasets. The card is among the first to support high-bandwidth PCIe 4.0 technology, which doubles the transfer speeds between the CPU and peripherals card attached, including GPUs.

Thanks to accelerated CPU/GPU multitasking functionality, the AMD Radeon Pro W5700 delivers up to 5.6x the application workflow performance with GPU loads compared to the competition in the SPECviewperf 13 benchmark, according to AMD. Certification with the leading design, manufacturing and architecture, engineering and construction (AEC) applications and new capabilities to improve virtual reality (VR) workflows are among the new GPU's other notable enhancements.

"The AMD Radeon Pro W5700 is the perfect graphics card for large assemblies as it's designed to deliver real-time visualization and multitasking capabilities," notes Antoine Reymond, senior manager, ISV alliances for AMD. "You can be designing an engine block and doing computing in the background, and the visualization and computing can use the GPU power at the same time without degrading per performance."

Although GPUs are certainly a crucial tool for boosting large-scale model performance, they aren't the only hardware

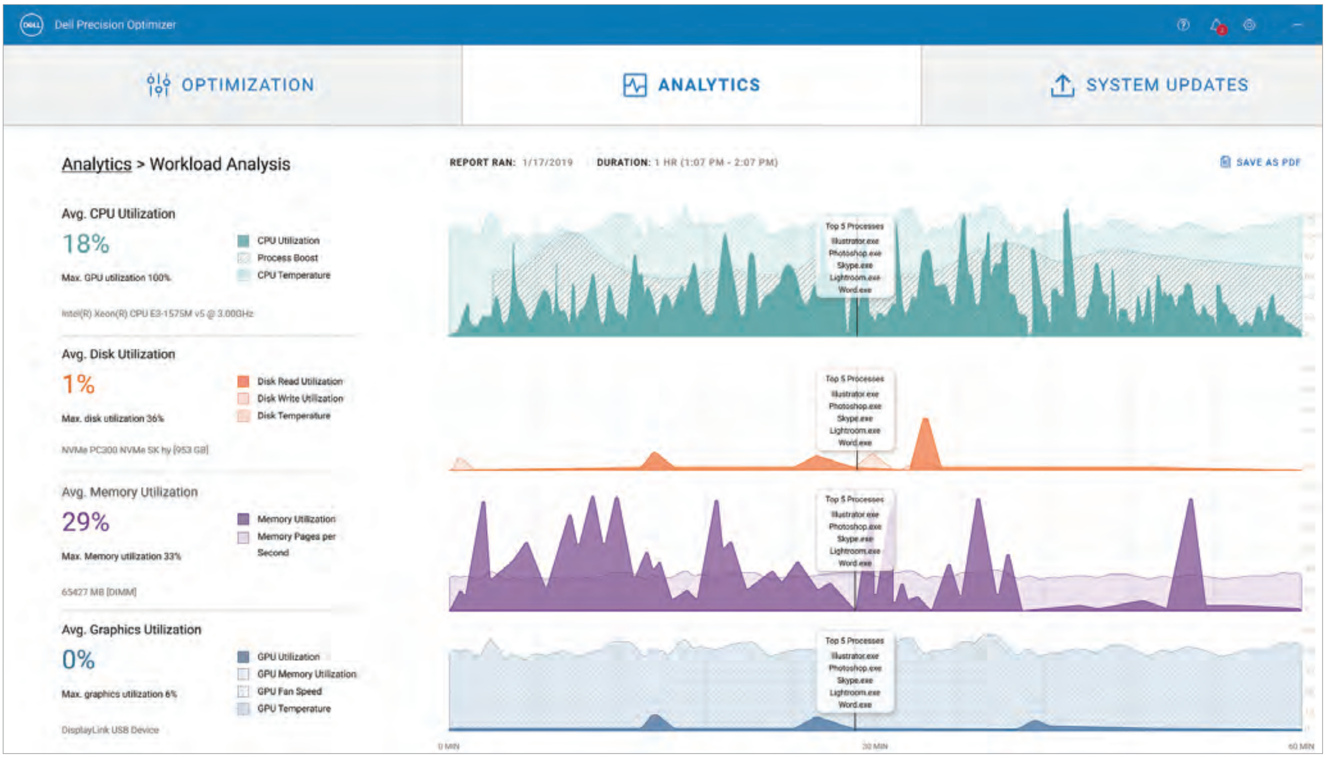

Dell Precision Optimizer software uses AI and machine learning to optimize performance of critical design tools running on Precision workstations. *Image courtesy of Dell.*

solution and they aren't perfect—sometimes users can have trouble getting data into the cache for the first time, notes Ken Versprille, executive consultant at CIMdata.

"Typically, the big complaint we hear from CAD users is that it takes so long to initiate the processes and activate the assembler and they struggle with that, even with newer GPUs," he says.

Beyond GPUs, the right choice of CPU, depending on the application, along with solid-state storage and memory options also play a big role in optimizing performance and addressing large-model complexity.

"GPUs are helping to revolutionize the engineering workflow, but you have to look at where the bottlenecks are in the system today," says Chris Ramirez, strategist for engineering, manufacturing and AEC, at Dell. "You have to look at the system in a holistic fashion—simply removing the bottleneck from the GPU when you still have problems at the storage subsystem level means even though you spent a lot of money, you might not have a workstation that was faster than before."

To help engineers better understand how and where their design software might stress their workstation, Dell offers Dell Precision Optimizer, a tool that provides a real-time view of system performance. Besides flagging potential system bottlenecks, Dell Precision Optimizer can tweak hardware so the Precision workstation model runs the slated application faster than it would with default settings.

"You can get two times faster performance on the same workstation running [Dell Precision Optimizer]," Ramirez says. "It can recognize that an app needs more GPU power and not as much CPU power so it will throttle down the CPU, transfer more power to the GPU, turn up the fans for cooling and tweak memory to allow the system to run faster. It's real-time optimization for workloads."

Software Vendors Find their Own Rx

Just as the hardware vendors are doing their part to advance GPUs and storage, CAD, simulation and design tool providers are also working to identify and re-architect areas in their code to enable the software to truly leverage GPUs and other optimization advancements.

"The software has to be re-architected to take advantage [of GPUs] … and that's a huge undertaking," explains CIMdata's Versprille. "The vendors are tacking on certain capabilities to improve performance, but they can't afford to redesign their complete architecture to support multiple threads."

Nevertheless, software providers are making headway addressing the performance and visualization challenges related to large-scale models. In addition to expanded GPU support, vendors are introducing data management features that allow engineering teams to holistically work on large models and render complete assemblies by limiting what must be loaded into memory or directing more processing

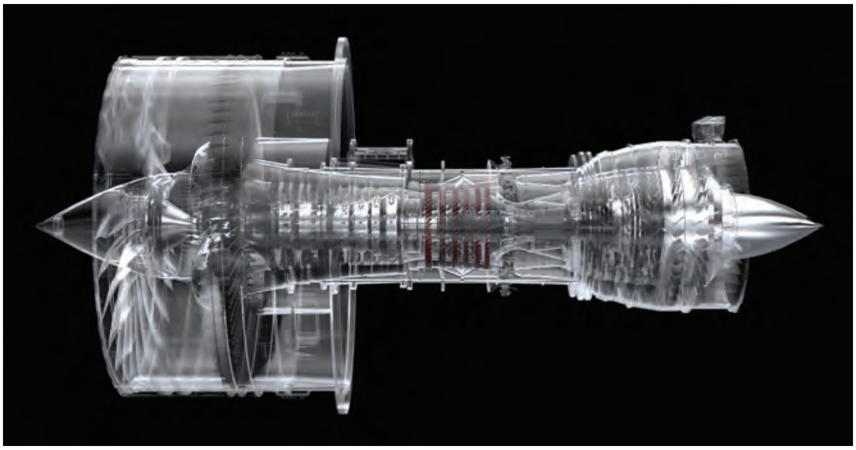

NVIDIA's Quadro RTX GPU platforms offers a range of capabilities to help facilitate large-scale model performance as well as bolster real-time rendering performance. *Courtesy of SolidWorks Visualize.*

work to the GPUs.

At Dassault Systèmes, for example, the SolidWorks team released beta capabilities in SolidWorks 2019 dubbed "RenderPipeline Project, [which is] aimed at delivering a rendering engine that is in line with modern programming paradigms and that can make complete utilization of a GPU," explains Siddharth Palaniappan, senior development manager, graphics applications at Dassault Systèmes.

"We store our user's model data as well as many other associated data on the GPU and do more work on the GPU than on the CPU for rendering," he explains. "This results in performance scaling across GPUs so if you have a high-end graphics card with a greater number of GPU cores, you get better performance. We also work very closely with vendors like AMD and NVIDIA to make sure their graphics drivers work nicely with SolidWorks."

SolidWorks is also focusing on accelerating model load times in addition to improving large model rendering performance. The team created a proof of concept for the February 3DEXPERIENCE World in Nashville, TN, to showcase how large SolidWorks models could be loaded using multiple cores and threads.

Although it's common for current-generation machines to feature quad-core or hex-core CPUs, AMD's new Ryzen Threadripper 3970X has 32 cores (64 threads), which makes handling and loading large assemblies much easier, Palaniappan says.

"More importantly, we were able to showcase that our model load time performance scaled with the number of CPU cores," he explains. "The machine was coupled with a Radeon Pro W5700 GPU, which allowed us to showcase buttery smooth [virtual reality] experiences for large assemblies. This is an excellent use case for large design reviews or modeling in resolved mode in SolidWorks."

At Autodesk, there has been a significant uptick these last few years in rewriting portions of code to better take advantage of parallel processing. For example, if an engineer makes a model change in Inventor, the system can divvy the calculations into multiple parts that are computed independently and brought back together at the end, den Hartog explains. Ray tracing operations are another use case where code parallelization can deliver extreme performance increases, he says.

Autodesk is also investing in real-world model testing to understand workflows and real-world bottlenecks and a feature called adaptive graphics. The latter is a specific capability that detects if the software refresh frame rate dips below a certain threshold when working with huge models; if so, it adjusts by not rendering the some of the smaller parts while rotating a model.

"As soon as the operation stops, the software redraws everything—the goal is to cut out the graphics lag time," den Hartog says.

With software vendors innovating new ways to leverage GPU horsepower, and hardware makers continuously pushing for new advancements, there are definite signs that the graphics lag for large-model performance is on the wane.

But it's up to users to know their software and fully leverage the capabilities to fully put the problem to bed. "Vendors are putting in a lot of bells and whistles to determine what you're working on so that only the graphics data required is fully loaded, which ensures much faster performance," says CIMdata's Versprille. "Engineers need to explore their options in this area and try to understand them as best they can." **DE**

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

→ MORE INFO

- **Autodesk:** Autodesk.com
- **AMD:** AMD.com
- **CIMdata:** CIMdata.com
- **Dassault Systemes:** 3DS.com
- **NVIDIA:** NVIDIA.com

For more information on this topic, visit DigitalEngineering247.com.

The Expanding Role of **GPUs in Simulation**

Legacy code is slower to move to the GPU; a hybrid strategy is recommended.

BY KENNETH WONG

In 2009, graphics processing unit (GPU) maker NVIDIA hosted its first GPU Technology Conference (GTC) in the Fairmont Hotel in San Jose, CA. Stepping up to the stage in the hotel's ballroom, NVIDIA CEO Jen-Hsun "Jensen" Huang unveiled what he called "the soul of a supercomputer in the body of a GPU." He was talking about the new GPUs built under the Fermi architecture, the basis for the company's subsequent GeForce and Quadro GPUs.

Fermi is "a brand new architecture designed from the ground up to be a computer first … [It treats] computer graphics and parallel computing as equal citizens," Huang explained. This was a watershed moment in the GPU's evolution.

Up to this point, the GPU was primarily a coprocessor to boost graphics performance—a piece of special hardware that made your explosions more spectacular in PC games. From this point on, the GPU would take on more and more of the type of computing usually done by the CPU. NVIDIA's compute unified device architecture (CUDA) programming language for GPU computing also laid the foundation for the transition. That opened the door to GPU-accelerated simulation, which harvests the GPU's superior number of processing cores to tackle massively large simulation problems.

In 2009, the premier GTC launched with 1,500 attendees. A decade later, GTC 2019 was attended by 9,000, NVIDIA verified. The company expects GTC 2020's attendance to reach 10,000. Similarly, the role of the GPU in simulation has significantly expanded over the past decade.

GPU vs. CPU

To be GPU-accelerated means to enlist the many processing cores in the GPU for general computing tasks. Therefore, tasks that are highly parallel—those that can be divided and simultaneously processed on many cores—tend to benefit more from GPU acceleration (for example, real-time rendering). By contrast, tasks that must be completed sequentially in a specific order, such as CAD geometry modeling, get more benefit from a faster CPU.

"Within the context of our integrated multiphysics solutions, our focus is on computational fluid dynamics [CFD], where for many years incremental improvements in performance have been realized by utilizing GPUs to perform the linear system solve. The real benefit comes when the full computational algorithm, including assembly of the linear system, is suitably implemented for GPUs," notes Matt Godo, senior manager, Simcenter STAR-CCM+ Technical Product Management, Siemens Digital Industries Software.

"This is quite naturally done for things like discrete element method (DEM), smoothed particle hydrodynamics and Lattice-Boltzmann technologies. It's also been shown recently that Navier-Stokes solvers, if properly architected, can obtain similar significant speedups on GPUs," he adds.

In 2016, Siemens acquired CD-adapco, maker of the STAR-CCM+ multiphysics simulation software suite, to add to its roster of simulation and test software solutions. The software is widely used by engineers to optimize and

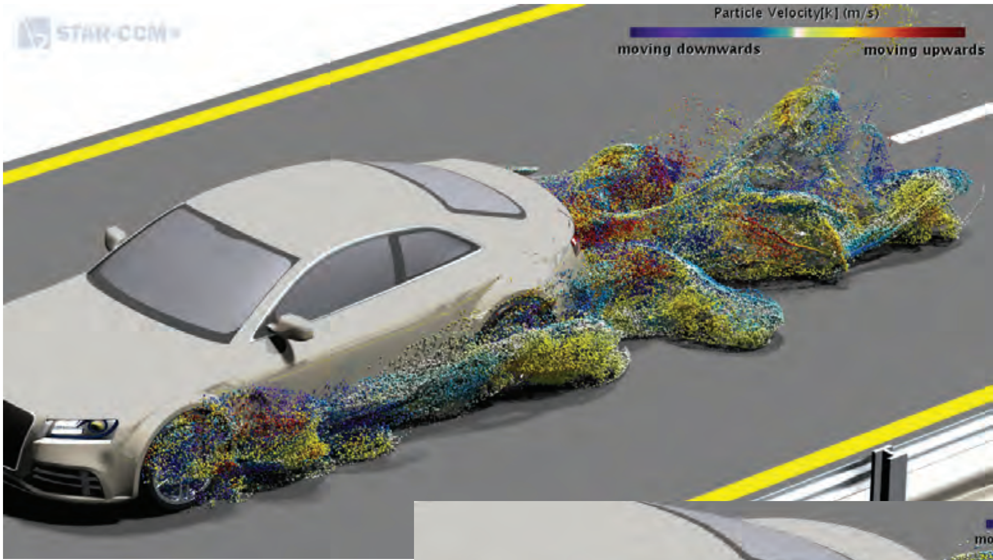

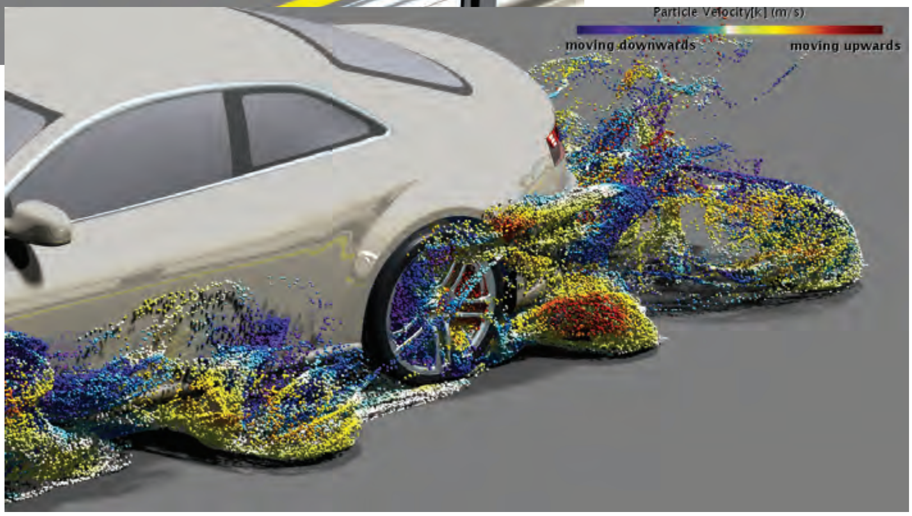

Using Siemens software to simulate how weather conditions like rain, snow or dust can adversely affect autonomous driving sensors. The software uses GPU acceleration to speed up Lagrangian Mass Particle (LMP) simulation methods. *Images courtesy of Siemens.*

test their designs across a broad range of use cases such as single-phase and multiphase flows, aero and hydrodynamics, heat transfer, electromagnetics and fluid-structure interactions.

Not Just CFD

In late 2018, ANSYS launched ANSYS Discovery Live, primarily targeting design engineers. Different from traditional simulation software targeting experts and sophisticated users, Discovery Live emphasizes speed. The software is written to compute and display simulation results in real-time or near-real-time speed. The speed comes primarily from GPU acceleration, implemented through NVIDIA's CUDA framework. The software currently runs exclusively on NVIDIA GPUs.

"We found that almost all types of simulation can be accelerated by the GPU," notes Mark Hindsbo, VP and GM, ANSYS. "Our Discovery Live, for instance, runs natively on the GPU. We have structural, thermal, fluid and electrostatic simulations running on the GPU."

Artificial intelligence training and machine learning—newer types of computation—are made possible by GPU computing. It's technically possible to run the same types of computation on CPU clusters, but the dramatic speedup possible and the large dataset involved make GPU clusters a more practical choice.

Legacy Code and the GPU

The origins of many modeling and simulation software for the engineering sector date all the way back to the 1980s and 1990s, when the yearly CPU speed increases (better known as Moore's law) was the bankable rule.

Many of the program codes, therefore, were architected to harvest the CPU's processing speed; GPU acceleration is a more recent phenomenon. Rewriting CPU-straddling codes into GPU-friendly codes is no small feat. In some cases, writing brand new programs may prove easier than revising the legacy code.

"Many simulation software vendors are still investigating how and if to switch to GPUs," observes Joe Walsh, CEO and cofounder of the ASSESS Initiative. "The challenge is, to get these performance improvements, significant architectural changes are needed in the software. Look for new applications built from scratch on GPUs rather than porting of existing large simulation applications."

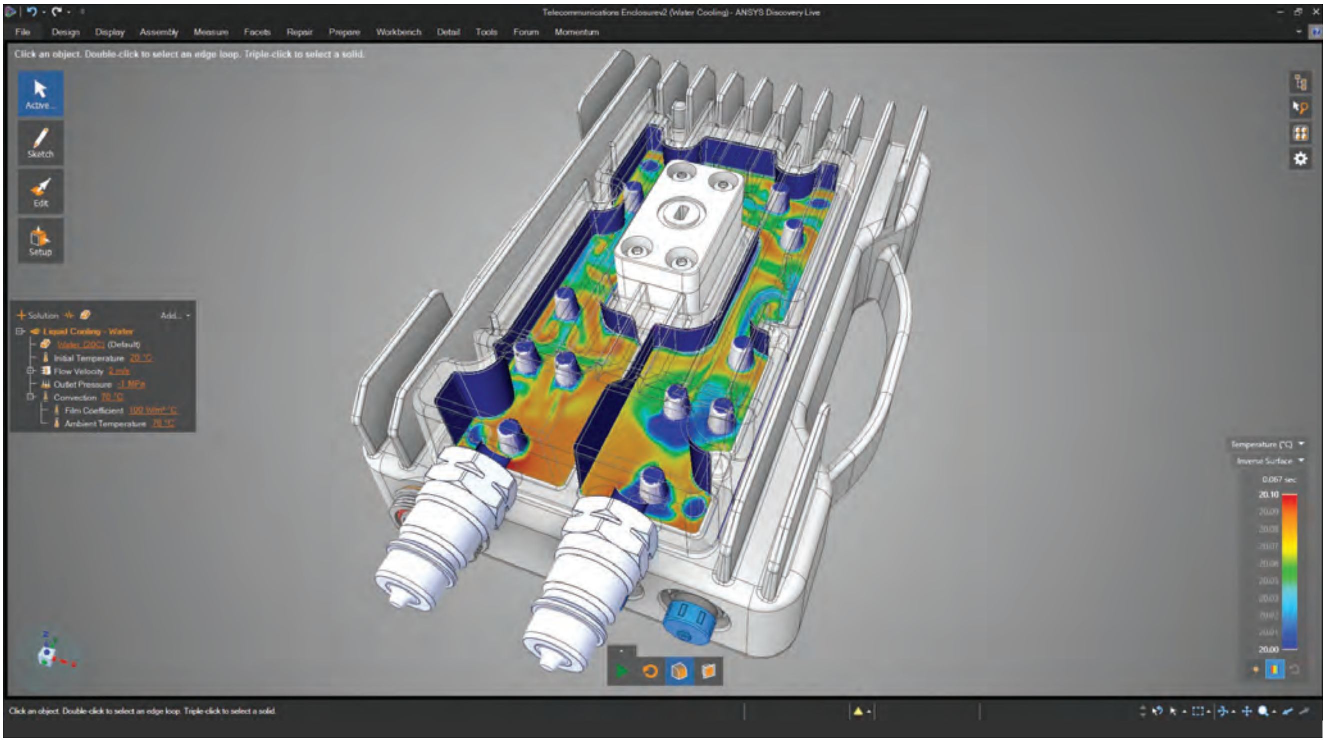

Using GPU acceleration, ANSYS Discovery offers near-real-time response to complex flow and thermal simulation. *Images courtesy of ANSYS.*

ASSESS, which stands for Analysis, Simulation & Systems Engineering Software Strategies, is a membership-based industry group. It holds annual congresses that bring together simulation software industry leaders and insiders to swap knowledge, present findings and discuss challenges. The ASSESS 2020 Congress is set for November 2-4, in Braselton, GA.

New Licensing Models Emerging

The effects of GPU acceleration may also spill into simulation software licensing. GPU acceleration encourages engineers to run design of experiments (DOE)—evaluating multiple design options simultaneously. It also plays a role in topology optimization and generative design software, which uses algorithms to compute and propose suitable design options based on user-stated requirements. But it also challenges the design software industry that traditionally charges users by seat or node (machines operating the software), not by consumption, as software-as-a-service vendors typically do.

Vendors such as ANSYS introduce high-performance computing (HPC) licensing to address the shift. “Each GPU is treated as a CPU core in terms of licensing, so users can gain higher productivity through GPU simulations,” writes ANSYS on its homepage for HPC licensing.

“The improved performance can enable near-real-time response, opening the market for significantly broader usage. To enable this broader usage, new pricing models will be explored,” says Walsh.

Boosting Performance

In benchmark results published by NVIDIA and ANSYS titled, “NVIDIA Quadro Enabling ANSYS Discovery Live,” running the same test job with NVIDIA Quadro RTX 4000 (2304 CUDA cores) provides more than 3x in performance boost over the older generation Quadro M4000 (1664 CUDA cores).

Switching to Quadro RTX 6000 (4608 CUDA cores) results in nearly 4.5x in performance boost over Quadro M4000.

These stats suggest that the more processing cores are in the GPU, the higher the speed bump. But there’s a point at which the performance boost becomes less dramatic, due to how much a job can be parallelized or how well the software is written to harvest the GPU.

“The wisdom of Amdahl’s law also comes into play here. We can reasonably expect performance gains to plateau,” warns Godo. This rule refers to a principle by computer scientist Gene Amdahl, who came up with a formula for estimating theoretical speedup in parallel computing workloads.

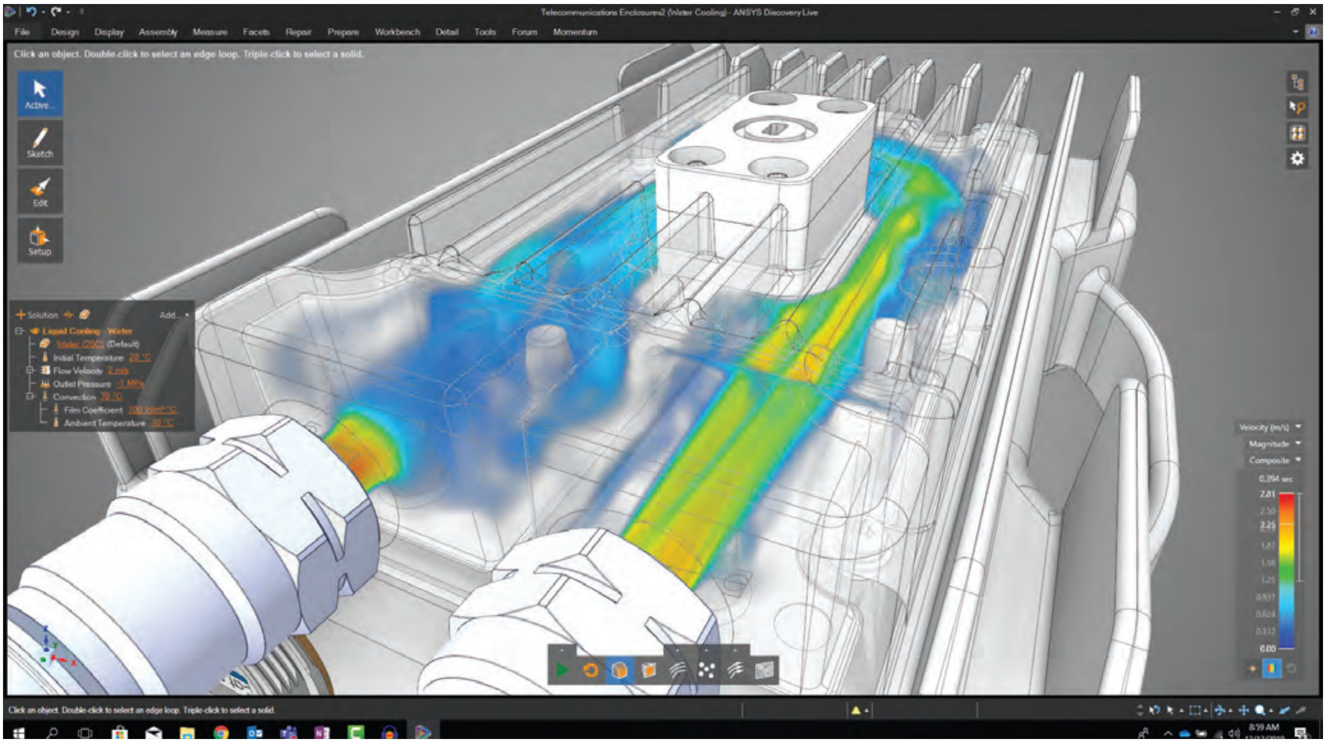

Amdahl stated that "speedup is limited by the total time needed for the sequential (serial) part of the program" (Bit.ly/2urVpzw).

"ANSYS Discovery Live can take advantage of two GPUs—one for general compute, the other for rendering visuals. The performance boost varies from application to application, but with today's software, you'll probably see [the] most significant performance improvements with one or two GPUs," says Hindsbo.

"But it's also important to look at the amount of memory," he adds. "How big a problem can you load into the GPU? If you have to go out of the GPU to load or save data, you pay a high performance penalty."

Fitting Into the GPU Memory

A GPU with more memory allows larger computing jobs to be efficiently loaded and processed. Therefore, if the size of the model or dataset exceeds the memory available on the GPU, the memory shortage undermines the GPU cores' ability to divide and conquer the job itself.

Hindsbo believes the best approach is the hybrid approach—to get the best available CPU and GPU. In his view, investing in one to the exclusion of the other is not wise.

"When planning for tomorrow's engineering compute platform, both on the desktop and in the cluster, you should be thinking hybrid GPU and CPU," he explains.

"The majority of the applications today are legacy applications written for the CPU, so you cannot realistically get away from the CPU. Even if you started from scratch, there are certain things the CPU will be better at than the GPU. Similarly, having no GPU computing will be a major deficiency both today and tomorrow," he adds. 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.*

→ MORE INFO

- **ANSYS:** ANSYS.com
- **ASSESS:** ASSESSInitiative.com
- **Siemens Digital Industries Software:** SW.Siemens.com
- **NVIDIA and ANSYS benchmarks on Discovery Live:** NVIDIA.com/en-us/data-center/gpu-accelerated-applications/ansys

For more information on this topic, visit DigitalEngineering247.com.

HPC Transforms 3D Printing

Design for additive manufacturing is pushing more companies to tap into cloud and high-performance computing resources.

BY BRIAN ALBRIGHT

Although additive manufacturing offers clear advantages in many manufacturing applications, current design and simulation resources may be challenged by the complex computational fluid dynamics simulations and shape optimization processes required. Cloud and high-performance computing resources can help companies take better advantage of new design possibilities enabled by 3D printing without requiring a costly investment in new compute resources.

3D print designs that can be simulated and optimized in an HPC environment (on-premise or, increasingly, in the cloud) will be a key enabler of digital manufacturing applications. A number of simulation software vendors, for example, have been working to allow their solutions to run simulations on the Microsoft Azure Big Compute platform.

Cloud resources are increasingly used for simulation and optimization purposes as designs become more complex, and models increase in size, requiring more multiphysics simulation work. For example, Utah-based Optisys creates 3D-printed radiofrequency (RF) antennas. The company uses ANSYS HFSS and Mechanical simulation tools on Rescale's cloud-based HPC infrastructure. Rather than tying up local computers for hours, the company can run advanced simulations in just a few minutes.

To take advantage of the possibilities of design for addi-

The Fortissimo Project

At the end of 2019, the European Commission wrapped up its Fortissimo Project, which was launched to encourage the use of advanced modeling and simulation by providing small and medium enterprises (SMEs) access to cloud-based HPC resources.

The experiments conducted in the early phases of the project resulted in 79 success stories and case studies that demonstrated how the use of cloud-based HPC could improve the value chain.

The project was primarily targeted at cloud-based simulation, a number of additive manufacturing use cases were enhanced through HPC use. For example, HSL, a high-tech product development center, estimates that using the Fortissimo HPC Cloud helped speed up development times for new air box optimizations, saving hundreds of thousands of dollars in the process.

The company is now able to provide alternative component designs for its customers that correspond to appropriate performance indicators, as well as the ability to offer shape optimization services in parallel with its rapid prototyping services. You can learn more about Fortissimo at www.fortissimo-project.eu.

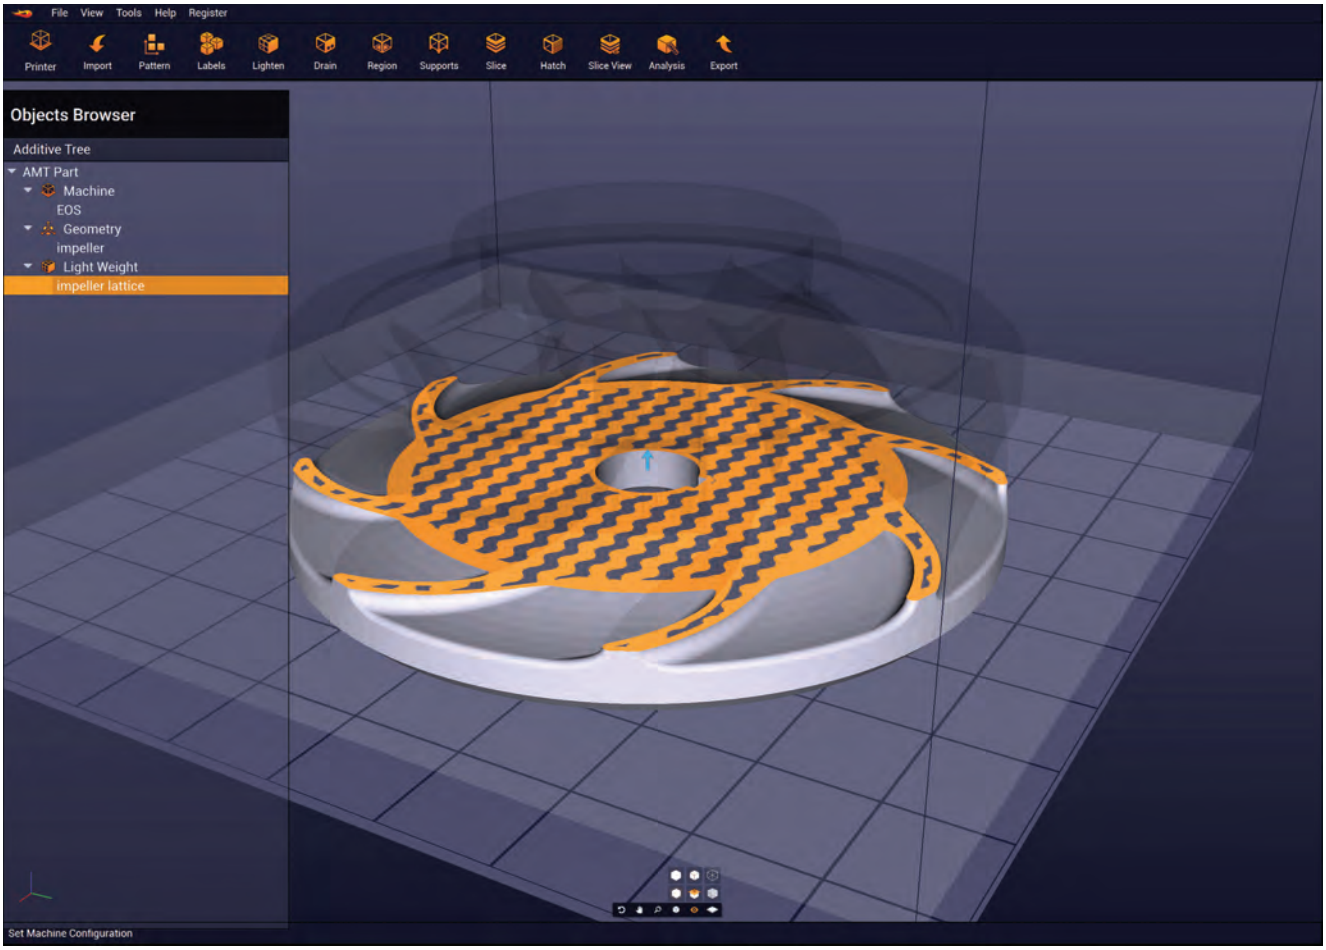

Dyndrite's GPU-based computation enables the addition of latticework and supports in seconds or minutes, compared to legacy additive manufacturing solutions that take hours and days to process the same elements. *Image courtesy of Dyndrite.*

tive manufacturing, companies increasingly need to leverage faster, graphics processing unit-based resources on their workstations and cloud-based HPC for simulation.

"We are at a unique point in time where the manufacturing hardware has outpaced the software," says Harshil Goel, founder and CEO of Dyndrite, which has developed a solution to streamline CAD-to-print processes.

Computing Options Evolve

Although concerns about IP and security have held back cloud adoption in some sectors, an increasing number of companies are more comfortable with using cloud resources for AM applications. "There's been some reluctance around CAD in the cloud because of a misperception about security," says Brent Stucker, director of additive manufacturing at ANSYS. "On the flip side, users really want to take advantage of the geometric complexity that additive has to offer, so those concerns about security are being outweighed by this advantage of running 1,000 design points or 1,000 different options on a part without having to buy a bunch of hardware. That is pushing them toward cloud and HPC resources."

The use of multiphysics in AM is pushing some of this activity. "With parts consolidation in additive, there is a huge push to create these large assemblies with fewer components, and those assemblies are responsible for achieving a lot of different physics goals. That can be computationally intensive. That ties up a lot of hardware that you don't want to buy locally, which is driving the use of cloud-based HPC resources," Stucker says.

Design, Simulation Tools Adapt

Software providers are developing new ways for CAD and simulation tools to help accelerate their performance for additive manufacturing applications. Those strategies include leveraging cloud-based access to tools and computing resources, as well as taking advantage of modern GPU capabilities. ANSYS, for example, launched a major initiative to enable all of its flagship products to solve through the cloud.

"We are also in the process of bringing all of the parts of our additive suite together in a single interface, and as we roll that out users will have the exact same option to solve to whatever hardware they want to solve on," Stucker says.

The Dyndrite solution serves as a bridge, similar to Adobe PostScript, that streamlines CAD-to-print workflows using a fast, GPU-powered geometry engine. In software demos, the product can create billions of voxels from

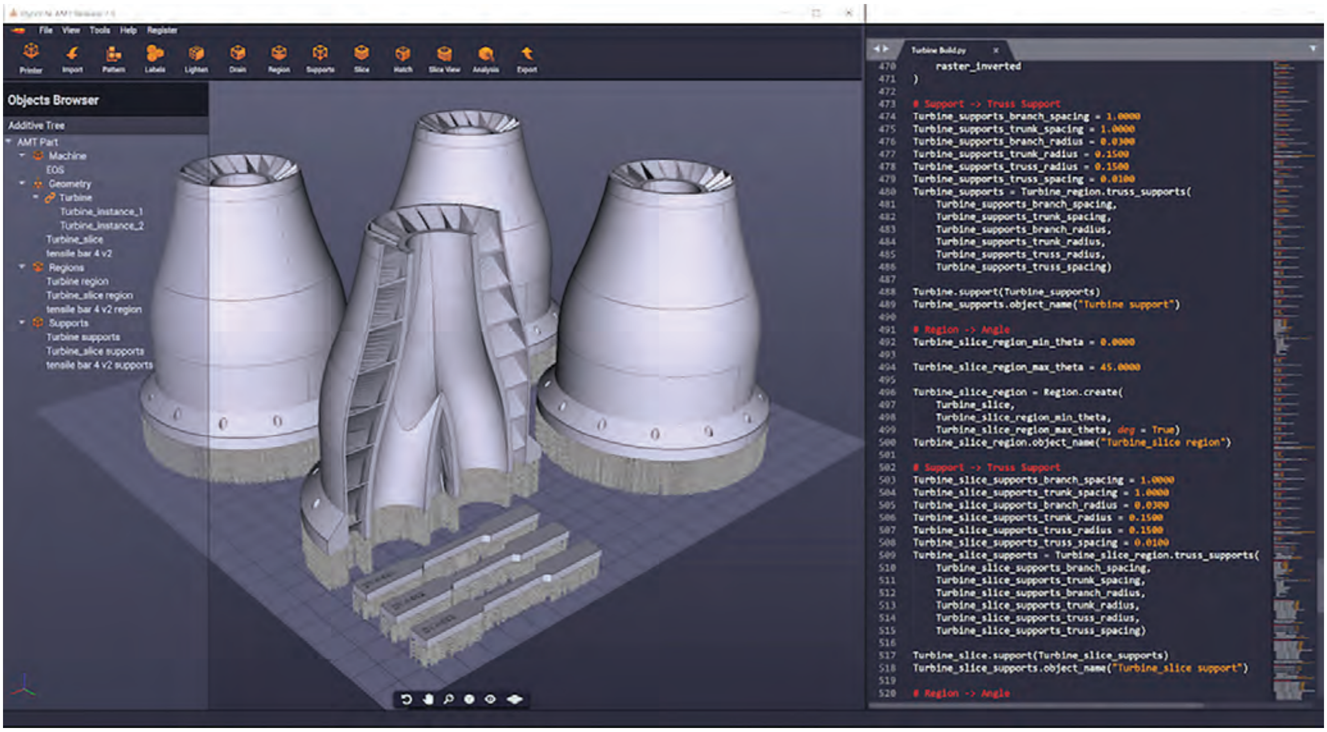

Dyndrite's additive manufacturing software delivers semi-automated python scripting routines to enable the automation of repetitive and time-consuming tasks during the preparation of 3D data for 3D printing. *Image courtesy of Dyndrite.*

a CAD file in just seconds, and slice the files just as quickly. "The only way to make that happen is using the GPU," says Dyndrite CEO Goel.

"Existing geometry kernels were not built to handle these complex geometries," Goel says. "Companies wasted a lot of money just loading files. Some customers we've spoken to used to take 60 hours to process a geometry. We help them cut that down to two. Some of this work can be completed on a laptop in less than 15 minutes. It dramatically affects iteration times and cycles."

This type of tool also eliminates steps in the traditional process, which require CAD files to be translated to STL files, then sliced and sent to the printer software. "If something goes wrong in that process, you have to start all over again," Goel says. "You aren't just losing time, you are also losing quality because you constantly sacrifice quality in the importing and exporting process."

Dyndrite is also working closely with simulation providers like ANSYS and Altair to integrate its geometry creation capabilities with their software. "The creation of the voxels is an expensive operation that we can accelerate," Goel says. "They still have to do the hard work on their side to [prompt] the GPU to simulate faster, but our goal is to provide the best geometry system so that part of the problem is taken care of."

Goel sees STL eventually going away. "It's a broken way of representing geometry," Goel says. "This delivers on the promise of AM. We were promised unlimited complexity, but it's not really unlimited because you still have a computer involved. If I want to hit the resolution of the machine, I need software and hardware to handle it."

Better compute resources are also required to handle complex additive manufacturing data management.

"Cloud and HPC resources also play a role in data management when it comes to design optimization for AM. There is all this information about machines and materials, setting up default options in experiments, simulation results from before, best practices for simulation, etc." Stucker says. "It isn't just about how to design a component, but how do you keep track of this massive amount of information that has been generated through experiments and simulation?"

In ANSYS' case the company's Granta acquisition provides a simulation data management tool called Minerva that provides distributed access to that information.

Navigating the Cloud

A variety of public, private and hybrid cloud infrastructures are available for companies that want to access design and simulation tools or tap into on-demand HPC resources. Increasingly, firms are more comfortable using public cloud services like Amazon Web Services (AWS), Microsoft Azure,

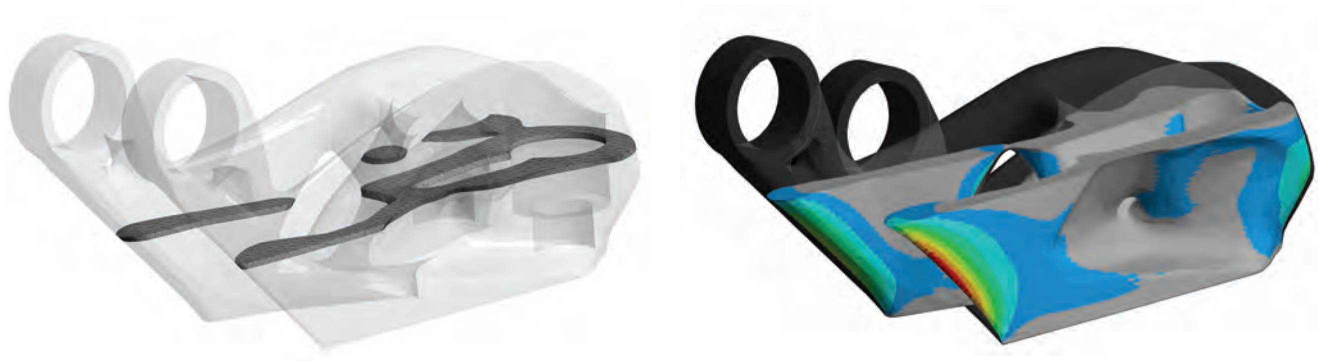

ANSYS provides a comprehensive and scalable software solution, which minimizes the risk of AM processes and ensures high quality, certifiable parts. *Image courtesy of ANSYS.*

IBM Cloud and others. Most software providers that offer cloud options also typically provide some flexibility for customers who either won't or can't (due to regulatory requirements) move to a fully public cloud infrastructure.

According to Stucker, early adopters in the cloud space have focused on private clouds, but that is slowly shifting to public cloud infrastructures. "These installed cloud architectures can become obsolete quickly," Stucker says. "Leasing compute power makes more sense for companies than buying hardware and hiring a team to manage it."

Goel, meanwhile, thinks that the next computing frontier in additive manufacturing will be pushing compute resources to the edge of the network, right down to the printers themselves. "The printers are increasing in size and in resolution, which means there is more data required to print," Goel says. "Something needs to be in place at the point of manufacture to handle that data."

As build volumes increase, the data required to create an object increases cubically. "This only gets worse as the machines grow in size," Goel says. "The amount of horsepower needed for that data is increasing rapidly. The fact that the printers are getting bigger and more process creates a problem and a challenge. People are also making more complicated things. You need a better and more sophisticated geometry engine to handle those more complicated things."

That could mean the possibility of real-time, in situ operational simulations that can help improve print quality and predictability. "If you are able to compute and simulate a part before it has been printed, the likelihood of a successful print goes up," says Shawn Hopwood, Dyndrite chief marketing officer and head of OEM relations.

"We are seeing more people putting simulation and printing software inside the machines," Goel says. "If you put a GPU in the machines, the printer is analyzing the part as it is being printed, re-simulating it and dynamically adjusting it. That's only possible when the machine itself is fast enough to compute."

These capabilities will grow in importance as more com-

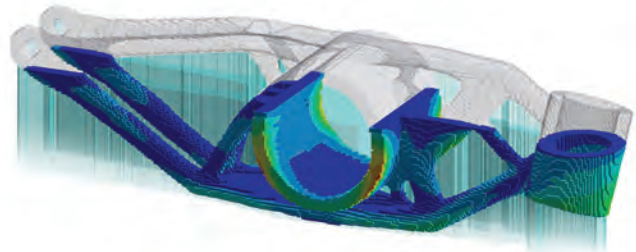

Additive solutions like those from ANSYS can reduce the number of failed builds and allow users to explore materials and optimal machine parameters. *Image courtesy of ANSYS.*

panies embrace additive manufacturing. "We're seeing a transition away from managing additive as a niche area for a company, to applying additive throughout an entire product portfolio," Stucker says. "To do that, we need everyone designing from the same data and the same assumptions. The only way to support enterprise-level sharing of resources is using the cloud and some form of HPC resources." DE

Brian Albright *is editorial director of Digital Engineering based in Cleveland, OH. Send e-mail about this article to de-editors@digitaleng.news.*

MORE INFO

- **ANSYS:** ANSYS.com
- **Dyndrite:** Dyndrite.com

For more information on this topic, visit DigitalEngineering247.com.

Making **Digital Thread** Work for You

Digital thread initiatives may require a mix of on-premise and cloud-based compute resources.

BY TOM KEVAN

The evolution of digital thread is still very much in flux. A number of essential components have yet to be clearly defined, and industry leaders and technology providers still have to agree on the best way to implement the concept to enable broad adoption. As a result, engineers are often left with as many questions as answers.

Some big issues that challenge companies interested in implementing digital thread include questions like: What is the best scale for the initial stages of a digital thread initiative? Is it practical to expect most companies to use on-premises resources to collect the relevant data and extract knowledge via big data techniques, or must they turn to cloud-based systems and services? What type of compute resources must companies bring to bear to handle the data that comes into play when implementing a digital thread, and what role does high-performance computing (HPC) play?

The answers certainly are not etched in stone. A look at today's practices and tools, however, does give engineers an idea of the current state of the technology and offer a glimpse of what may come tomorrow.

The Concept

Just one of several buzzwords to emerge from the convergence of big data and the drive toward digitization, "digital thread" refers to a communication framework that connects data flows to create a holistic view of an asset's data throughout its lifecycle (Fig. 1). The overall goal is to eliminate silos of information that so far have hindered data exchange between disparate software systems.

By aggregating all design, manufacturing, performance, supply chain and maintenance data of an asset, the digital thread aims to provide the means to deliver the right information, to the right place, at the right time. This translates into providing engineers and companies with the ability to access, integrate, analyze and transform data from all relevant systems into actionable information.

This digital mechanism promises to revolutionize design, manufacturing and supply chain processes by supporting more effective assessment of an asset's current and future capabilities; enabling early discovery of performance deficiencies; opening the door for continuous refinement of designs and models; and optimizing manufacturability.

Bottom line? Advocates contend that the digital thread will mitigate the challenges posed by the complexity of today's product development processes.

Where Does the Data Come From?

To construct a digital thread, engineers tap two groups of data. The first consists of "digital definitions data," provided by domain platforms like product lifecycle management (PLM), application lifecycle management, manufacturing operations management and enterprise resource planning (ERP) systems.

"PLM systems source most of these digital definitions," says Dave Duncan, vice president, product management, at PTC. "ERP/CRM orders, product clouds and other sources provide the lighter weight product properties, like model numbers and options. PLM retrieves configuration-specific detail from these product properties by decoding 'recipes' of overloaded design definitions, such as CAD, BOM [bill of materials] and process plans."

The second data group contains information pertaining to the physical experiences of products and processes. This data is often sourced by Internet of Things (IoT) platforms, which provide sensor-based activity data on business system activity, such as manufacturing histories, field dispatches, maintenance work orders and warranty claims.

Basic Implementation Steps

There are several phases to building a digital twin. The first step calls for the engineer to create a model, or

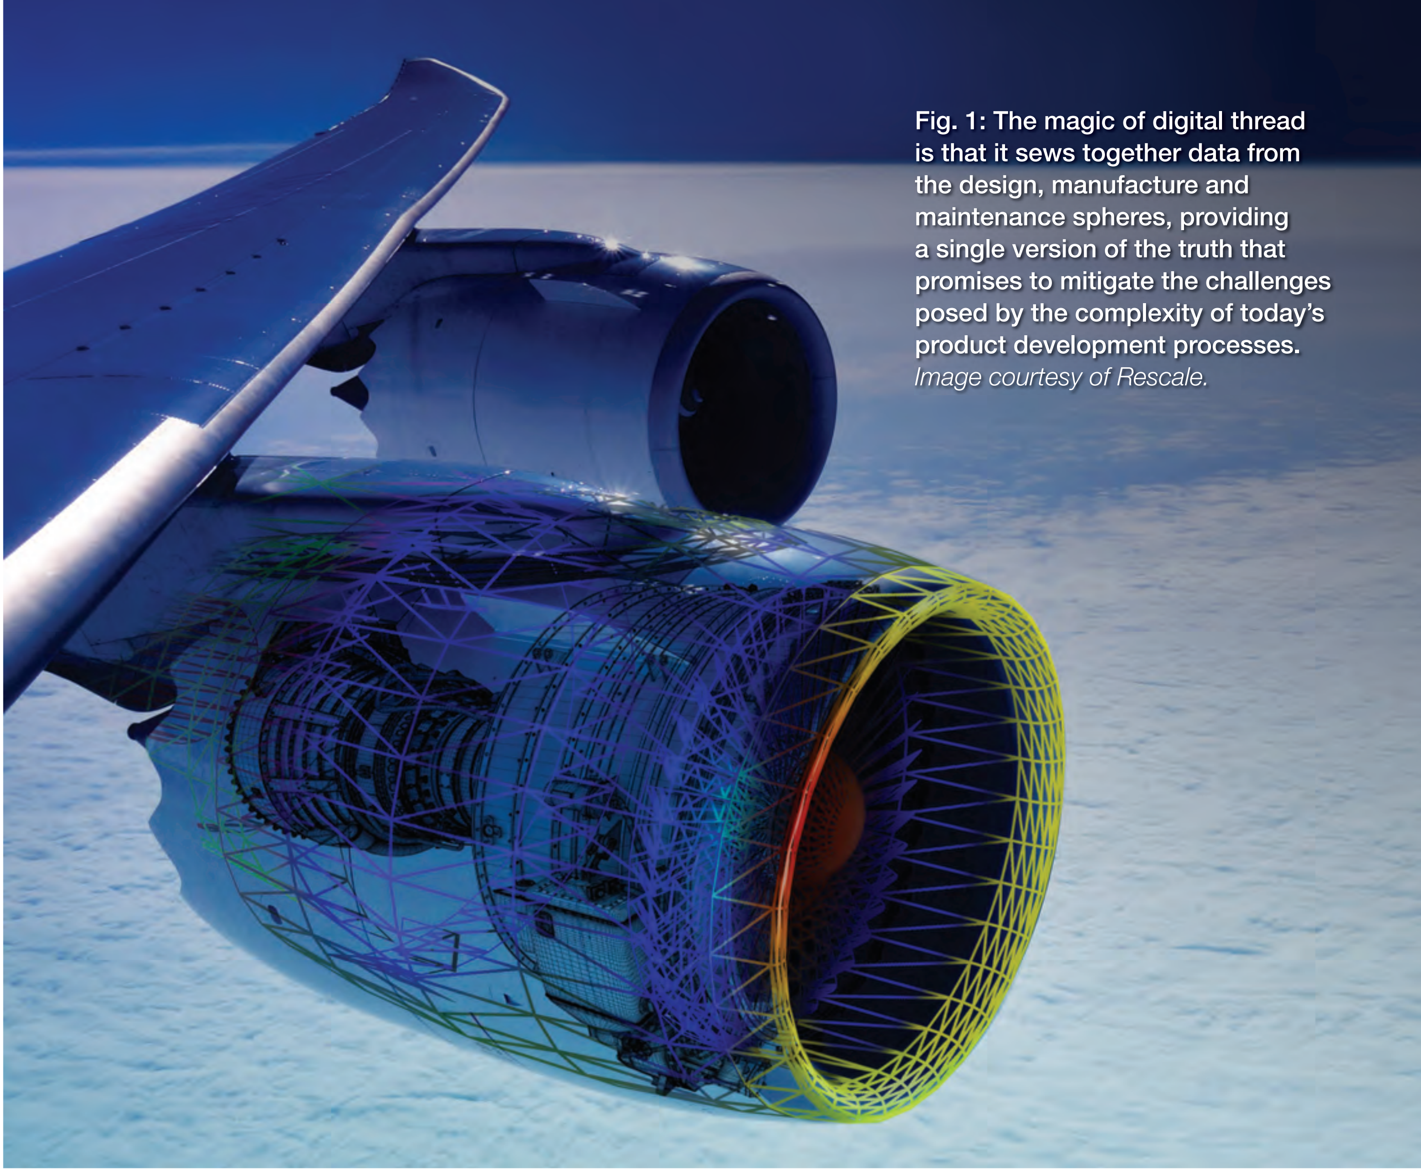

Fig. 1: The magic of digital thread is that it sews together data from the design, manufacture and maintenance spheres, providing a single version of the truth that promises to mitigate the challenges posed by the complexity of today's product development processes. *Image courtesy of Rescale.*

digital representation, of one or multiple physics that support the product or system.

The development team then gathers a set of data to feed the model, using different sources, which range from IoT-device sensor data to traditional data sources like comma-separated value (CSV) files. Everything that will enrich the model and add value is welcome.

The challenge then becomes stitching the disparate types of data together using a digital thread. This represents the glue that allows the engineer to connect the right data with the right model or models.

All of the aforementioned data must be supported by a set of tools that combines CAE, IoT and advanced analytics that run on top of the infrastructure. This, in turn, calls for a transversal layer of security for the whole chain.

These, however, are only the first steps toward reaping the full value of digital thread.

"So far, the industry has been moving in one direction, getting everything in a single digital environment," says Alvaro Everlet, senior vice president, IoT platform, Altair Engineering. "Now it's time to use that privileged point of view to turn it into value that translates into insights or actions that the analytics tools can provide."

Start Small

In general, all digital thread implementations follow these standard steps. That said, every implementation is shaped by application-specific requirements.

Foremost of these modifications are often process constraints rather than technology limitations. To avoid problems throughout the development process, these obstacles should be addressed holistically, early in the initiative.

A chief challenge involves determining the ideal scale of the initiative. This requires identifying priority use cases and determining what data is required to support them. It's best to start with a subset of potential digital twins, and stitch together only what you really need.

"We've worked with customers that initially tried to extract every piece of product and process data they could find in their systems," says Duncan. "They then link them together and tell their app developers, 'Here's everything we have. So go get whatever you need.' Not only were these budget-busting implementations, but the end result was too complex for the organization to utilize. If it requires data scientists to query a digital thread, it's a miss."

Starting small allows development teams to deploy the digital thread quickly and to achieve value early, but engi-

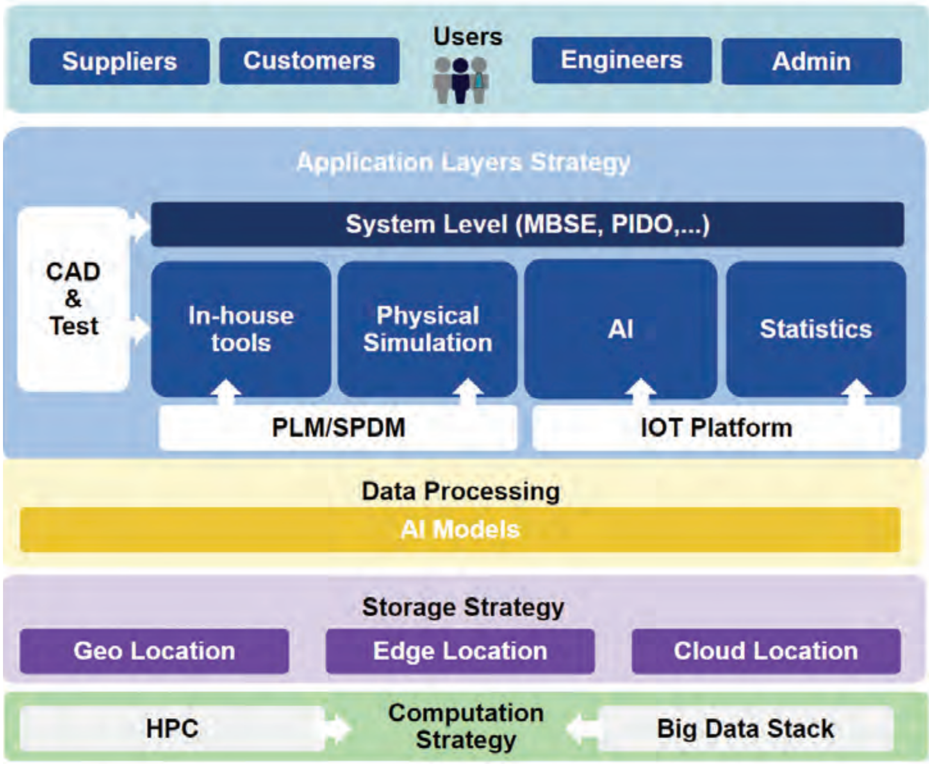

Fig. 2: Hybrid systems can help companies extend on-premises high-performance computing systems to the cloud. These platforms provide an entry point for companies to begin tapping cloud resources. *Image courtesy of Rescale.*

neers should also avoid boxing themselves into a corner. Early on, they should plan how they are going to iterate.

This requires considering two dimensions. The first is use case coverage, which involves expanding the data and query paths. Ensure these anticipated data and paths will gracefully layer onto the digital thread foundation.

The second dimension is integrating capability within the existing use cases, from descriptive to predictive to prescriptive value. Descriptive results return data fields specific to a product or process, in context to the user and task. A lightweight digital thread that keeps most of the data in place serves descriptive value well. Data in place inherently secures the data, too. Users can access only the data that they can reach directly via the source applications.

Predictive and prescriptive results involve heavier data lifting and more robust security structures to calculate similarities, analyze outcomes and apply machine learning.

Value of Local Computing Resources

Perhaps one of the biggest digital thread considerations revolves around the compute resources required to ensure a successful initiative. The question becomes: Can most companies succeed using on-premises resources to collect all the relevant data and extract knowledge, or must companies turn to cloud-based systems and services to achieve digital thread's full value?

Like other most digitalization initiatives, no single answer fits all situations. For development teams to make the right choice, they have to weigh a number of issues.

For example, the local, or on-premises, approach has great appeal when issues like regulatory compliance and latency play key roles.

"Many digital thread initiatives incorporate some level of field sensor data or OT [operational technology] factory data, so it will be necessary to ensure that investments are made in extending the enterprise IT capabilities to the edge and to recognize the cost and effort required to link IT and OT systems," says Rick Arthur, senior director of digital engineering at GE Research. "Real-time and latency-sensitive data and processing will likely reside on edge/IoT and in dedicated [on-premises] systems."

Opting for the Cloud

Although arguments can be made for the on-premises approach, the scale, dynamic nature and workloads of digital thread initiatives increasingly draw engineers and companies to the cloud. In general, most processing-intensive and highly variable workloads benefit from the elastic qualities of a cloud infrastructure.

"[Digital thread] workload demands are almost always nonlinear, and having the ability to burst capacity—while incurring a temporary, incremental cost—will meet tight business requirements with a significantly lower [total cost of ownership]," says Taylor Newill, senior manager of product management at Oracle.

"User-friendly deployment automation options, along with on-demand burst infrastructure capacity, allow for fast, easy and low-cost options for testing new software versions and performing additional process iterations on large data sets. In the cloud, it's a trivial exercise to snapshot terabytes of data on-demand, replicate a complete production compute environment, test version changes, play back logged activities during an investigation, or perform additional analytics without impacting production activities," he adds.

Cloud infrastructures also perform well in supporting digital thread initiatives that call for agile enterprise-scale access to specialized compute resources. Rather than restricting or prioritizing access to on-premises resources based on limited hardware capacity, the cloud promises to make it possible to address demand in a timelier manner.

"When it comes to ease of scalability for the digital thread, we think it is more realistic for companies to turn to cloud-based systems and services," says Duncan. "Engaging with a cloud-based system or service allows for a more flexible and agile process. The information making up the digital thread is not fixed, and the means by which it is used should not be either. Additionally, as a digital thread iterates, suppliers and other partners often participate directly, which is much easier and [more] secure to broker with cloud approaches than corporate intranets."

This type of flexibility becomes particularly important when the digital thread's development process takes on a global scale.

"Today's organizations have engineering centers and customers distributed across the world, and the complexity of the supply chain is increasing," says Fanny Treheux, director of industry solution marketing at Rescale.

A Little Bit of Both

Digital thread developers don't have to limit themselves to either on-premises computing infrastructures or cloud-based systems and services. Today, a growing number of companies are choosing to have it both ways (Fig. 2).

"Increasingly, organizations are adopting hybrid infrastructures that can leverage either local or cloud resources through technology like cloud bursting, depending on the workload requirements," says Everlet. "For example, an organization using digital thread physics-based models may work with data-driven models in the cloud to define an algorithm that can be run in a local environment for low-latency decisions."

This approach allows companies to maximize the benefits they derive from IoT assets while meeting their need to craft digital threads that operate well on a global scale.

"Depending on regulatory, security and process needs, we see companies deploying cloud and non-cloud solutions for different domains," says John Grape, Sr., director of technology and innovation at Siemens Digital Industries Software. "IoT platforms and edge solutions are being adopted by companies for gathering data from their operating assets for digitalizing their service processes. Cloud-based solutions are being created to address the global reach and scalability challenges.

"From an engineering data perspective, companies could consider leaving the work-in-progress data on-premises and storing released data in the cloud to the level that they need to support their business scenarios," Grape adds.

Is HPC Necessary?

After deciding whether cloud services are required for a digital thread implementation, the next question development teams should address is the need for high-performance computing (HPC). As with many other aspects of digital thread development, application requirements play a big role in helping engineers answer this question. Often the approach pursued includes a variety of computing platforms.

"What we see with our customers is that they often run many types of simulations, many of which are best served by different types of infrastructure," says Edward Hsu, vice president of product at Rescale. "This can include general-purpose CPUs, GPU-enabled, high-memory or high-interconnect systems—what people often think of as HPC."

In addition to the type of simulations being run, the type of compute resource required also depends on the active phase of the digital thread's lifecycle. "Typically, while authoring the digital thread, engineers operate within a working context," says Grape.

"There is not a lot of need for additional computing power to support these tasks. The issues of massive data do not really come into play here," Grape adds. "Consuming a digital thread across disciplines and domains can, however, be more challenging, depending upon the scope of the digital thread required for a business need."

For example, when the digital thread is being used to improve product design and performance, it may be necessary to gather as-operated field data from customers, as-validated performance data from test labs, as-built yield and manufacturability data from factories and suppliers, and as-designed and design trade-off options from engineering.

"Consolidating and cross-referencing these data will require significant investments in big data systems to capture and unify field, physical test, design simulation, manufacturing and supply chain data," says Eric Tucker, senior director of HPC and machine learning products at GE Research. "Depending on the complexity of your product and maturity of your design practice, HPC resources may be required."

When cost becomes the determining factor, market research increasingly indicates that the use of HPC resources quickly provides a good return on investment. A 2018 study conducted by Hyperion Research (hpcuserforum.com/ROI) on economic models linking HPC with return on investment shows that a dollar invested in HPC resources returns approximately 100 times on each dollar invested in traditional product design practices. DE

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

MORE INFO

- **Altair:** Altair.com
- **GE Research:** GE.com/research
- **Oracle:** Oracle.com
- **PTC:** PTC.com
- **Rescale:** Rescale.com
- **Siemens Digital Industries Software:** SW.Siemens.com

For more information on this topic, visit DigitalEngineering247.com.

Security in the Era of Cloud

Stolen credentials are a greater threat than stolen laptops.

BY KENNETH WONG

Gartner's Magic Quadrant is the little square that all companies want to be a part of. Published as part of Gartner's annual industry sector reports, the square identifies market leaders, challengers, niche players and visionaries. In the analyst's "Magic Quadrant for Unified Endpoint Management Tools," published August 2019, VMware, along with Citrix, Microsoft and IBM, showed up in the Leaders' corner.

"VMware shows strength in helping its customers bridge between legacy client management tools (CMTs) and UEM unified endpoint management (UEM) by building on previous common CMT functions such as patch management natively within the Workspace ONE offering," wrote Gartner. It also singled out "VMware's 2018 launch of the Workspace ONE AirLift offering" as an innovation in the space.

"We are seeing a trend toward virtualization technology, and this is a key enabler for running our software in the cloud," notes Prakash Kota, CIO, Autodesk. "Many Autodesk customers are exploring the option of moving their desktops running Autodesk software to the cloud. We are also seeing companies such as Citrix, AWS and Microsoft Azure offer cloud-based desktops."

The shift from on-premises desktop to virtual machines suggests the security risks, too, have migrated to the cloud. But human behavior, as it turns out, remains a bigger security risk than the holes in the company firewall and the cloud.

A Famous IP Theft Case

One of the most high-profile IP theft cases in the CAD industry occurred in 2002, when the source code of a prerelease version of SolidWorks was stolen. The culprit, Shekhar Verma, was a former employee of an outsourced contractor in India. He was eventually apprehended when he attempted to sell the code to an undercover federal agent from India. ("At Risk Offshore," Michael Fitzgerald, Computerworld, November 2003, ComputerWorld.com.) It was long before the rise of software-as-a-service and virtual desktop infrastructure, and ultimately an act of a disgruntled employee.

Could mission-critical CAD design files be stolen from the cloud, or via a virtual machine? It's quite probable, but Robert Thompson, senior technical instructor, VMware, worries more about poor security practices among humans than the virtual holes in the VDI.

"Have you ever worked with designers who said, 'Hey, log in and take a look at project XYZ. Here's my password'," he asked. "Never give your password out to anyone, especially if your login has access to high visibility projects with lots of IP.

SaaS vs. Virtualization

Software-as-a-service (SaaS) describes software that does not need to be installed on a local PC or workstation to run. It tends to run from a standard browser and is typically licensed via subscriptions or pay-as-you-go models.

QuickBooks online for accounting and SAP cloud apps for customer relationships and human resources management exemplify SaaS' widespread adoption in the office productivity and enterprise sectors.

Virtualization, or virtual desktop infrastructure (VDI), describes the use of remote computing resources, such as virtual workstations and virtual servers in the cloud.

The key to virtualization is the software that enables the user to interact with the virtual product as though it were a real physical piece of hardware. The widespread availability of high bandwidth connections and the upcoming 5G wireless networks both work in favor of this trend.

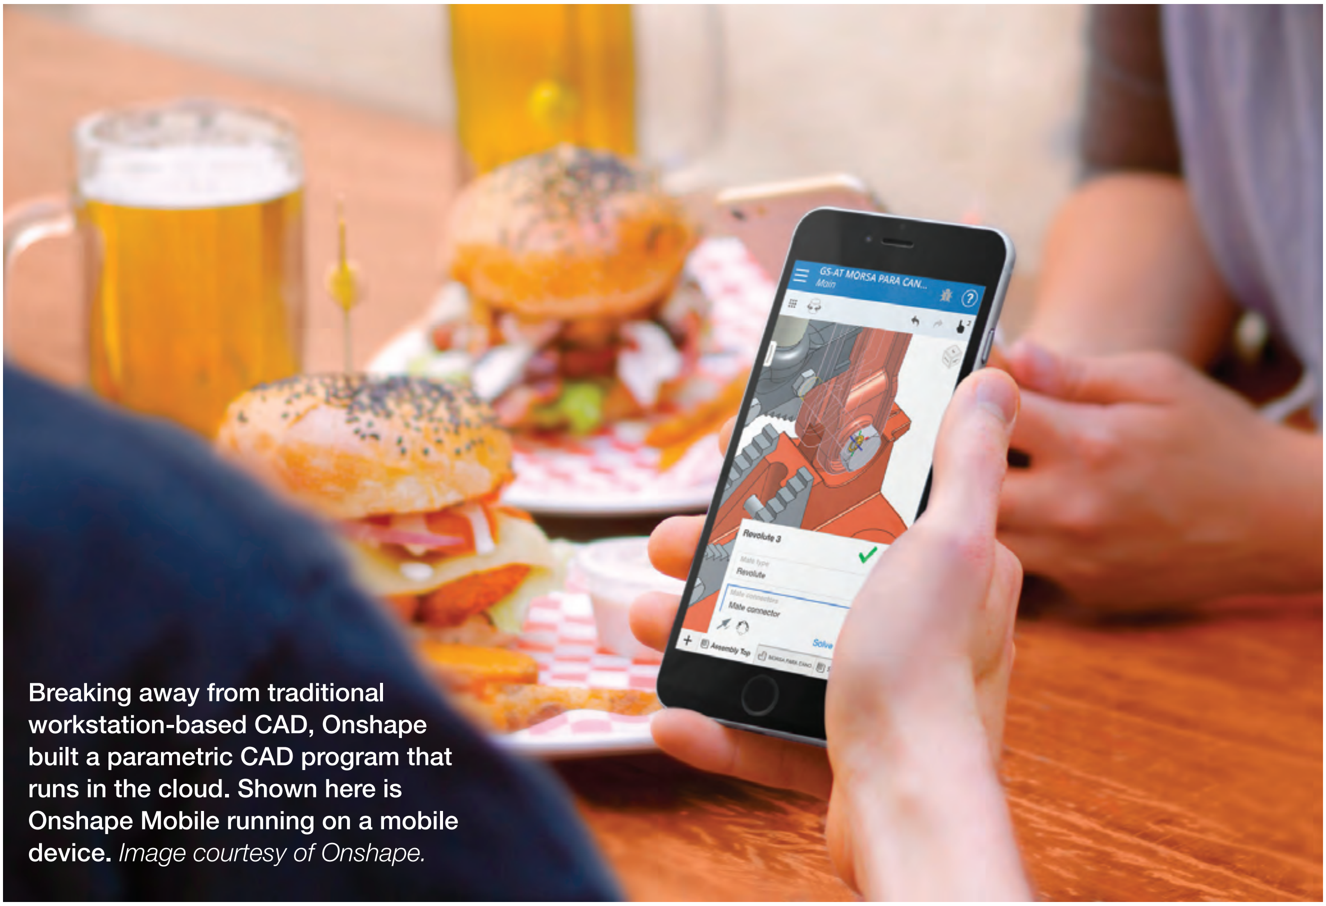

Breaking away from traditional workstation-based CAD, Onshape built a parametric CAD program that runs in the cloud. Shown here is Onshape Mobile running on a mobile device. *Image courtesy of Onshape.*

The extra step of creating [and] requesting another login for your colleague could very well prevent theft," he adds.

Tight deadlines and a human preference for efficiency may tempt some to bypass established security protocols, perhaps not out of malicious intent but purely out of the need to get things done quickly.

Password sharing is not only dangerous but it's unnecessary with good VDI. "There's a feature in VMware Horizon called session sharing that would remove the need to share passwords with applications like CAD," says Thompson.

CAD in Cloud

CAD software in its origin was written to run on personal workstations. That was the case with SolidWorks, Autodesk Inventor, PTC Pro/ENGINEER (later rebranded as PTC Creo), and Siemens NX and Solid Edge. In that sense, Onshape, launched in 2012, was a groundbreaker. It was architected from the start to run in the cloud, SaaS-style. The software was acquired by PTC in late 2019.

"SaaS is very different from virtualization. With SaaS, you have a centralized multi-tenant network of computers that provides service to people," says Jon Hirschtick, co-founder and CEO of Onshape. "As a result, you are not working in the inherently insecure Windows OS workstation environment; you are not copying and emailing files around; and you are not maintaining hundreds of different instances of the project."

In Onshape, users collaborate with others by inviting them into the same modeling session, or sharing the centralized cloud-stored file with someone via a link, where you control the recipient's reading, writing and editing privileges.

"The argument that you are more vulnerable in the cloud is a bit old-fashioned, like keeping your money under the mattress because you don't trust the bank," Hirschtick says. "We have to worry about hacking, just like any modern business these days. But I dare say, our security measures for our data center are far more secure than the steps taken by a typical corporate IT team or PC user."

"With some government and military customers, they don't even allow you to bring in your cellphone or storage devices into the facility," observes Hirschtick. "These types of places currently won't let you run SaaS software."

Compliance and Security

But it's not necessarily the end of the road for SaaS vendors who want to work with strict military or government clients. Amazon, one of the biggest cloud service providers, launched GovCloud in late 2011, creating a way for SaaS vendors to comply with the regulations.

"AWS GovCloud is a new AWS Region designed to allow U.S. government agencies and contractors to move more sensitive workloads into the cloud by addressing their specific regulatory and compliance requirements," writes Amazon.

Attendees at VMware 2019 conference in San Francisco, CA. VMware provides virtualization software that enables companies to deploy remote workstations. *Image courtesy of VMware.*

AWS data centers running the AWS GovCloud services are located in geographic regions acceptable to government and military users, both for compliance and security reasons. The stringent requirements go both ways. To use the GovCloud services, the customer also has to meet certain criteria.

"AWS GovCloud (US-East) and (US-West) regions are operated by employees who are U.S. citizens on U.S. soil. AWS GovCloud (U.S.) is only accessible to U.S. entities and root account holders who pass a screening process. Customers must confirm that they will only use a U.S. person (green card holder or citizen as defined by the U.S. Department of State) to manage and access root account keys to these regions," Amazon states.

"We will at some point put our product on AWS GovCloud. When we do, we can reach these types of clients," says Hirschtick.

Stolen Credentials

With increasing numbers of companies switching to SaaS or VDI for cost savings, stolen credentials are now a much greater risk than stolen computers.

"If you steal my computer right now, you still cannot get my Onshape CAD files. Onshape data never leaves the server. It never gets copied into your local machine," says Hirschtick. "But there's no system that can fully prevent credential misappropriation. It's just like someone getting your banking credentials and logging into your bank account."

"Most of the customers I work with do not see virtualized infrastructure as an added vulnerability," says Thompson. "Whether you are using physical or virtual machines, there is still a chance of intrusion, especially if your machines are secure but your network internet connection/firewall is wide open."

"Some common causes of data breach regardless of physical or virtual environment include weak and stolen credentials, application vulnerabilities, malicious insiders or insider error. Addressing these issues and implementing a proper security and defense strategy is essential to minimizing risk," says Kota. 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.*

→ MORE INFO

- **Autodesk:** Autodesk.com
- **AWS Govcloud:** AWS.Amazon.com/govcloud-us
- **Gartner:** Gartner.com
- **Onshape:** Onshape.com
- **VMware:** VMware.com

For more information on this topic, visit DigitalEngineering247.com.

The Changing Nature of Hackers

In a recorded webinar titled, "Enhance VDI Security for Virtualized Desktops and Apps with VMware NSX and Horizon" (available on the VMware EMEA YouTube channel) presenter Brigitte Skakkebaek, product marketing manager for end user computing company, VMware, discussed the changing nature of cyberattacks. Likely attackers now include organized crime, insiders, hacktivists and nation states.

"First and foremost, inherent in virtualization itself is the elimination of the loss of data sitting on the edge devices [for example, a poorly secured or unmonitored laptop]. With VDI, every time a user logs in, they get a brand new desktop; every time they log off, their virtual desktop is destroyed," said Skakkebaek.

The Forbidden Simulation Riff

Discover how bad simulation "riffs" can ultimately lead you down a slippery slope of bad modeling.

BY TONY ABBEY

Back in the '70s, when I was in college, everyone had a Led Zeppelin LP collection, but the ultimate street credibility came from playing the tunes on a guitar.

Guitar playing had exploded in popularity; it was such a cool instrument to play. What made it more attractive was the underlying simplicity of the riffs the great bands were playing. Even I could pick up a simple chord or two and fingerpick my way to an approximation of the iconic riffs.

Therein lies the snag—my contemporaries and I were playing crude approximations to the real thing! It sounded good to us, and probably to most of our audiences after a beer or two. But have pity on the professionals who had to suffer through this noise.

The Forbidden Riff of FEA

I remember trying out guitars in the store, and my favorite tunes were "I Was Born Under a Wandering Star"—goodness knows why—and "Stairway to Heaven." In fact, I was one of many thousands of very poor players bashing out the latter, perhaps

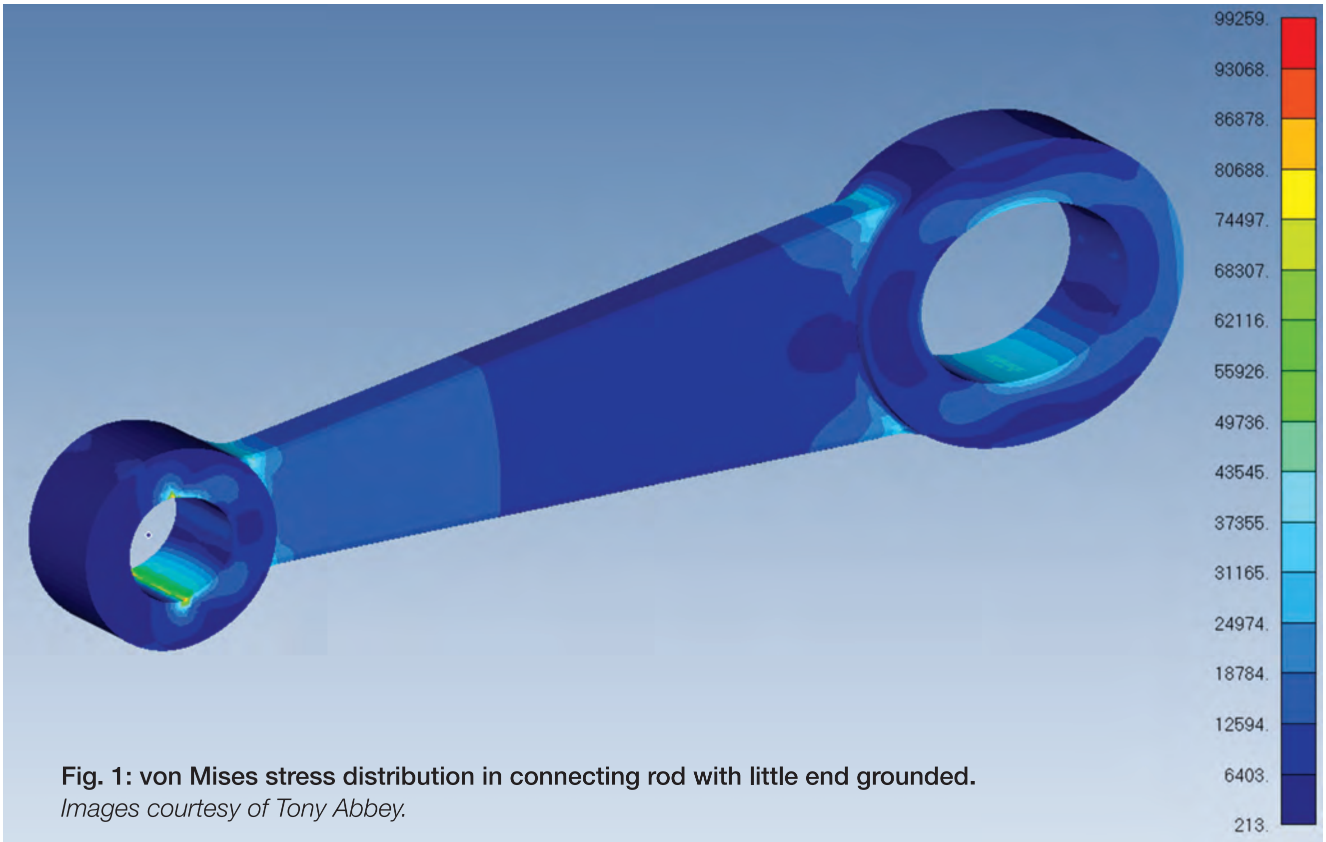

Fig. 1: von Mises stress distribution in connecting rod with little end grounded. *Images courtesy of Tony Abbey.*

one of the most iconic guitar intros ever.

The staff in the guitar stores probably had to resort to ear plugs after hearing this tune massacred so many times. An urban legend was born over the next few years that the "Stairway to Heaven" riff was forbidden to be played in the stores, and even in public. Buskers and other street performers were too embarrassed to play it!

I thought about the forbidden riff a few days ago, when I stumbled across yet another YouTube video showing "how to do finite element analysis (FEA)." The component was a connecting rod, loaded at one end and fully constrained at the other. There was no discussion as to whether this was a good, bad or indifferent technique.

Instead of providing a link, I decided repeating the analysis with a similar model for this article. The loading consists of a force applied at the big end and a constraint at the little end, to produce a tensile stress in the conrod as shown in Fig. 1.

The peak von Mises stress is 99,529 psi and appears as a sharp "blip" at the little end. This is due to the fully fixed constraint method, which is shown in Fig. 2.

The constraint is shown applied over the bearing face at the little end. This can be imagined as "supergluing" the little end to an infinitely stiff surface.

There are two side effects from this action. The first effect is that the natural stress distribution created by a compressive bearing force and tensile hoop force in the left side of the little end is completely locked out. The second effect is that the stresses at the run adjacent points (top and bottom) are a singularity.

Fig. 2: Little end fully fixed on bearing face.

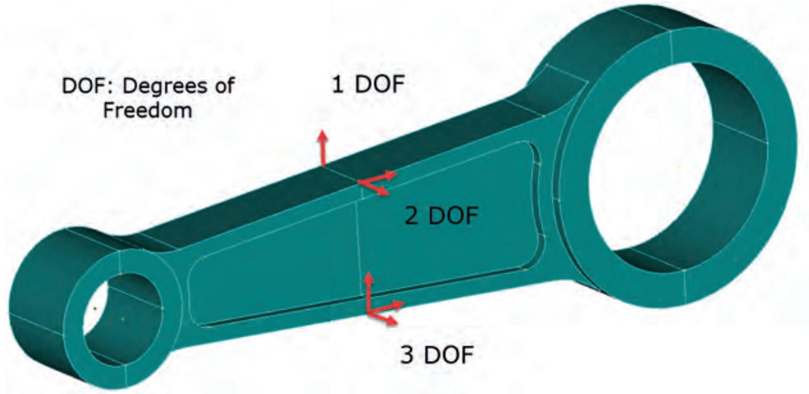

Fig. 3: The 3-2-1 method applied to the connecting rod.

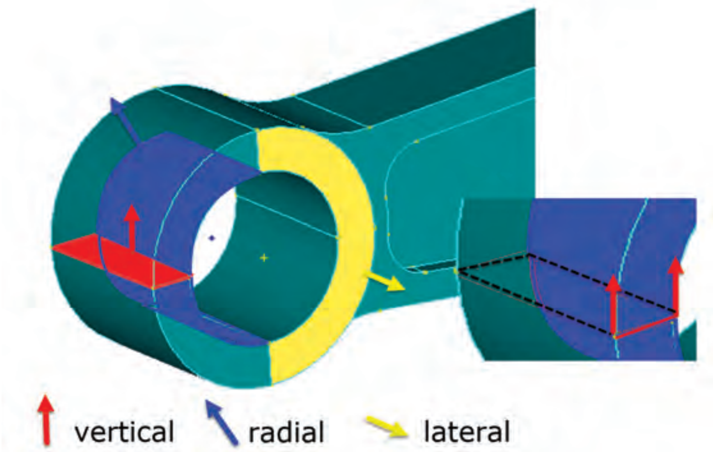

Fig. 4: Ad hoc method for constraint at little end.

If the mesh is refined, the stress will increase in value, but never converge.

In my analogy, this is the forbidden riff, played very badly! It is a quick and dirty way to set up a model to show the basics of FEA, but leads us straight down a slippery slope of bad modeling.

Playing the Riff Professionally

Various methods produce a much better result. The 3-2-1 method, described in the February 2015 issue (tinyurl.com/wb4nz6g), uses a balanced set of applied loads and the minimum constraint set of 6 Degrees of Freedom (6DOF) shown in Fig. 3.

When meshing, three points help identify three nodal locations. The recipe for constraints applied as shown in Fig. 3 ensures that all six rigid body motions are eliminated.

We then apply a set of balanced loads; in this case, the equal action and reaction through the little end and big end of 15,000 lbf. The loads are applied as approximations to a bearing distribution, using pressure over 180° applied as a sine distribution. This type of loading is supported in most FEA tools.

EDITOR'S NOTE: Tony Abbey provides live e-Learning courses, FEA consulting and mentoring. Contact tony@fetraining.com for details or visit his website at www.fetraining.net.

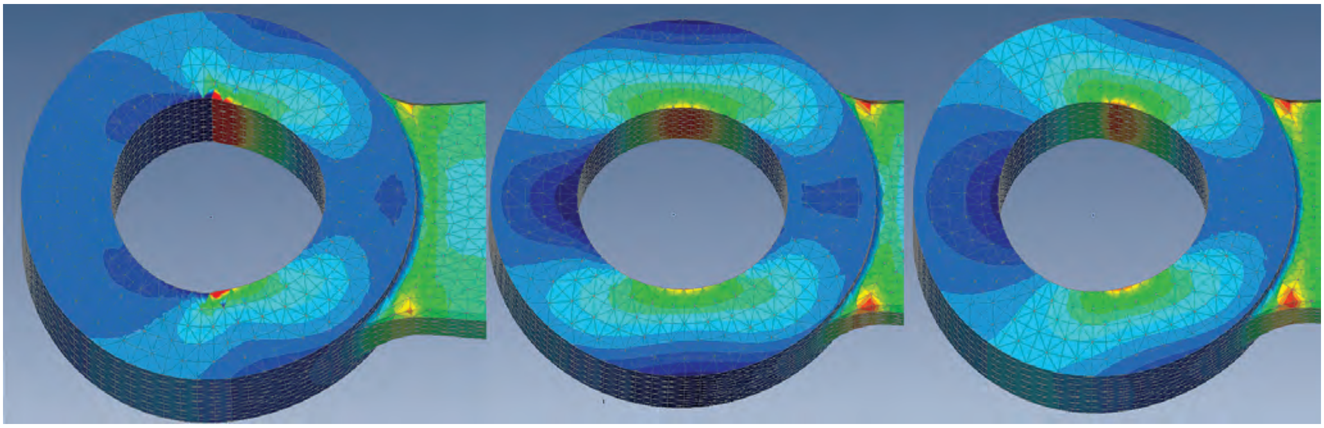

Fig. 5: Axial stress and deformation at the little end; (left) fully constrained, (middle) 3-2-1 method, (right) ad hoc constraint.

Inertia relief is a useful alternative to the 3-2-1 method. This applies balancing distributed inertia body forces to match any out of balance loads and suppresses rigid body motion. If the out-of-balance loads are zero, it effectively becomes a zero "dummy" set of inertia forces. I won't use this technique here, but you can consider it.

Another alternative to the formal 3-2-1 method is to investigate ad hoc constraint systems that achieve the required constraints for specific loading actions. Note that if the loading line of action changes, the constraint system must be changed. Fig. 4 shows a setup for the axially loaded connecting rod.

We can constrain the three translational systems of 6DOF by restricting:

• Vertical motion using vertical constraints applied as shown in red in the figure. A cut plane in the geometry model had to be created to achieve this. After meshing, the nodes are merged in this plane. Some FEA tools may not allow constraining an embedded surface like this. In that case, a pair of vertex points can be used from the cut plane intersection, positioned as shown in Fig. 4 inset with the red arrows, to provide a reaction couple about the bearing rotation axis.

• Radial motion using radial constraints in a cylindrical system about the bearing axis, as shown in blue. This is a common feature in many FEA tools.

• Lateral motion on an arbitrary face, as shown in yellow.

The corresponding moments are also constrained with this set. In each case, reacting couples are created from the translational degrees of freedom.

The vertical constraint is overkill in the translational sense, as the radial constraint has vertical and axial components. But it is needed to provide a vertical couple about the bearing spin axis. This could be done, as mentioned, more elegantly with just two points constrained in the vertical direction, but offset axially.

Fig. 5 shows the peak axial stress and deformation comparison between fully constraining, using the 3-2-1 method and the ad hoc method.

Comparison of Constraint Methods

The peak stresses at the little end seen in Fig. 5 are 69,163 psi, 55,823 psi and 45,795 psi for the fully constrained, the 3-2-1 constraint and the ad hoc constraint, respectively. The fully fixed constraint, as expected, destroys any realistic stress distribution to the left of the bearing surface.

Here, we are only looking at the axial stress component, yet the absence of a compressive load path can be clearly seen. The singularity created by this method also can be seen. With a finer mesh, the stress value at the constraint runout will increase; it just isn't a good method.

Let's now focus on the two other methods.

The middle plot in Fig. 5, of the 3-2-1 constraint shows a symmetrical ovalization of the little end bearing region. There is no representation of the pin, which has a "choking" action in practice, and would have prevented the vertical narrowing of the little end.

The right-hand plot in Fig. 5 shows a non-symmetric displacement of the little end bearing region. The radial constraint used here is preventing the inward motion of the little end, so it cannot narrow vertically.

The material can stretch around the constraint in a hoop, or circumferential, direction. The corresponding stress plots also show a variation on stress distribution. The right-hand (3-2-1 method) peak stress plot shows the peak stress moving away from the top and bottom central position seen in the middle (ad hoc method) peak stress plot.

As noted, the peak stresses are different. The lack of pin stiffness representation is sometimes a weakness of the 3-2-1 method, but at least it is providing a conservative stress estimate.

A Nonlinear Riff

In reality, the physics of the little end and pin interaction will be nonlinear. Both of the last two methods are brave attempts to use linear finite element analysis to represent this. The 3-2-1 method ignores pin interaction, and the ad hoc method simulates it with an infinitely high radial pin stiffness.

The main motivations for avoiding a nonlinear analysis are solution time, the unpredictability of the nonlinear convergence and the accuracy of the contact responses. These can be big factors in assessing risk to the project in timescale, and even perhaps technical success.

Fig. 6 shows the axial stress distribution at the little end under the same loading, but now using a nonlinear analysis with the pin modeled.

The peak stress is 48,647 psi, compared to 55,823 psi and 45,795 psi for the 3-2-1 constraint and the ad hoc constraint, respectively. The stress distribution and off-center peak stress position agree well with the ad hoc constraint method, seen in Fig. 5 left-hand. The deformation also matches the ad hoc method. The total stretch along the overall conrod also was very similar for the ad hoc constraint and nonlinear methods.

Fig. 7 shows the confirmation of the pin to bearing gapping.

The bearing contact region extends past the vertical center line. This is an interesting effect, which could be fed back into the linear ad hoc model to increase the bearing angle there beyond 180 degrees. The tricky part is that the effect is nonlinear, so the correction would not be valid with increased or decreased load levels. It is also interesting that the position of the peak stress beyond the 180-degree plane was predicted by the linear ad hoc model, so is not perhaps so sensitive to the actual bearing angle.

The nonlinear analysis included both pins (big end and little end) and a dummy piston. The model required careful setup of constraints to avoid the pins spinning about their own axes and to eliminate the rigid body motions of the conrod.

Eventually, the piston was eliminated, and the load applied directly to the little end pin. A nonzero friction value was required in the contacts to achieve analysis stability. This again highlights the additional effort required when moving to a nonlinear analysis.

Does the forbidden riff analogy work?

I think so. A poorly defined FEA simulation looks vaguely like a well-defined one. If the poor model is presented as a workflow example without any qualification, then it can be very misleading.

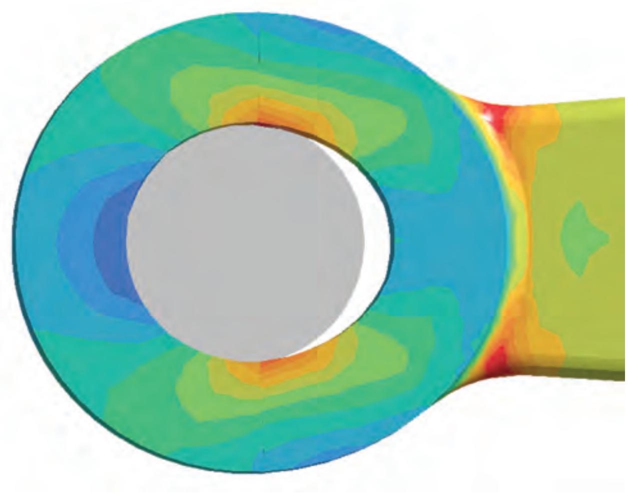

Fig. 6: Axial stress and deformation at the little end using a nonlinear analysis.

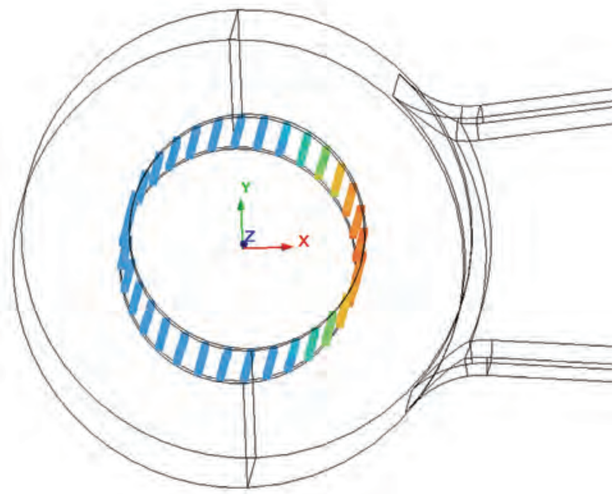

Fig. 7: Gapping at the little end bearing face; blue, no gap and red, maximum gap.

Often, YouTube and similar videos are more focused on showing the user interface setup than on good FEA practice. A similar situation can arise with verification benchmarks where the objective is to match a classical manual calculation, rather than model a physically meaningful configuration.

Bad guitar playing is self-correcting; keep on doing it and you just lose your audience, as I remember well! Bad modeling, however, might not be so easily spotted. DE

Tony Abbey *creates e-learning classes that are presented in partnership with NAFEMS. Check out the latest course and contents at* http://bit.ly/courses-Nafems

Lenovo ThinkPad P53: Power-Packed ThinkPad

Lenovo delivers the fastest 15-in. laptop we have ever tested.

BY DAVID COHN

Lenovo recently sent us its ThinkPad P53, the P51's successor and the latest addition to its P-Series mobile workstation lineup. Designed for professionals who need the highest performance in a mobile workstation, the ThinkPad P53 delivers lots of features in a 15.6-in. form factor.

Like the ThinkPad P72 we reviewed earlier this year (*DE*, August 2019; digitalengineering247.com/r/22956), the Lenovo ThinkPad P53 comes housed in a charcoal gray case with a glass-reinforced polymer top and a magnesium aluminum alloy bottom cover to dissipate heat.

The system measures 14.86x9.93x1.16-in. (WxDxH) and weighs 5.89 lbs. At 6.63x3.38x1.0-in., the 230-watt power supply is 35% smaller than the previous generation and adds just 1.95 lbs., including its cables. The dot over the "i" in the ThinkPad logo on the lid glows red when the system is powers up.

Raising that lid reveals a 15.6-in. display and 105-key backlit spill-resistant keyboard with separate numeric keypad. The consistently reliable Lenovo keyboard is once again one of the best available in any laptop, with nicely sculpted keys and a great feel. Lenovo offers a choice of four in-plane switching anti-glare displays.

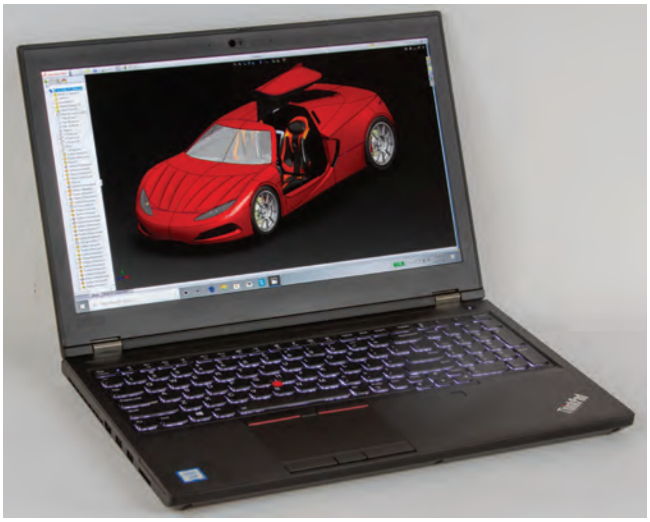

The base configuration includes an HD 1920x1080 panel with a brightness rated at 300 nits. Our evaluation unit included a brighter 500-nit display with Dolby Vision HDR 400, which added $25. Lenovo also offers ultra-high-definition (3840x2160) displays—a 500-nit version ($200 extra) or a 300-nit panel with multi-touch ($270).

A round power button is located on the upper-right corner of the numeric keypad, while a fingerprint reader is positioned to the lower left of the keyboard, just below the cursor keys. A 4x2.25-in. touchpad with three dedicated buttons is centered below the spacebar. There is also a red pointing stick nestled between the G, H and B keys with its own three buttons directly below the spacebar. A pair of stereo speakers are concealed beneath a perforated screen just above the keyboard. The caps lock and number lock keys each have their own lights, as do the function keys dedicated to the speakers and microphone as well as the ESC key, which doubles as FnLock. There are also hard drive activity and Wi-Fi lights in the hinge area below the center of the display.

Preconfigure, Then Add Options

Lenovo offers a number of preconfigured systems, including a $1,489 base model equipped with a 2.5GHz Intel Core i5-9400H 4-core CPU, a 1920x1080 15.6-in. display, 8GB of RAM, a 256GB M.2 solid-state drive (SSD), an NVIDIA Quadro T1000 graphics processing unit (GPU) and a 720p web cam. That's the starting point.

The company offers a choice of four different Intel Core processors as well as a Xeon CPU. The ThinkPad P53 can accommodate up to 128GB of memory and up to 6TB of storage. All models include a discrete NVIDIA graphics card in addition to integrated Intel graphics, and customers have a choice of five different NVIDIA Quadro GPUs—the Quadro T1000 in the base unit, the slightly

Fig. 1: The Lenovo ThinkPad P53 is a powerful 15.6-in. system with a top-of-the-line Xeon CPU and NVIDIA Quadro RTX GPU. *Image courtesy of David Cohn.*

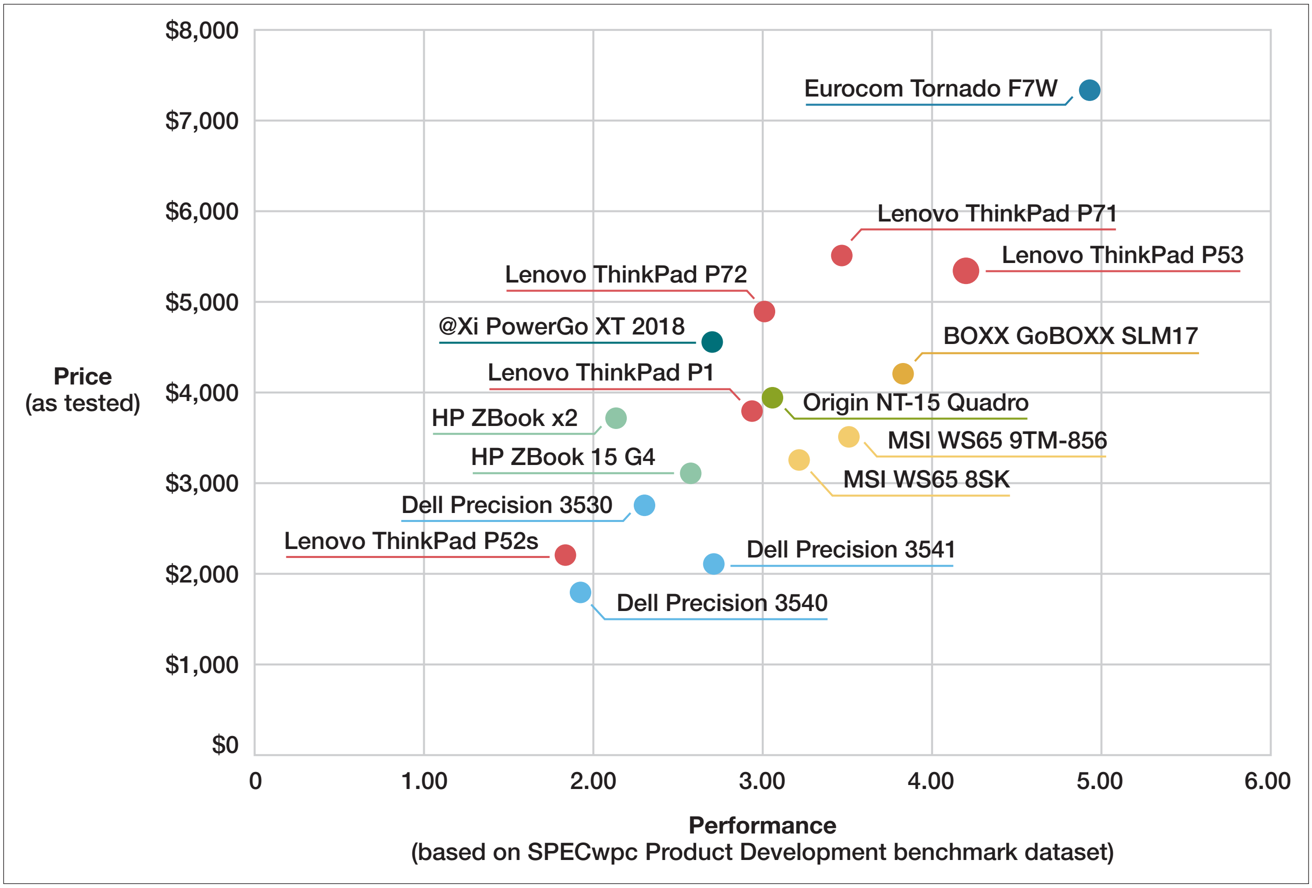

Fig. 2: Price/Performance chart for recently reviewed workstations.

more powerful T2000 or three VR-ready Quadro RTX GPUs: the RTX 3000, 4000 or 5000.

Configuration choices depend on whether you choose an Intel Core or Xeon processor. The system we received was based on a 9th-generation Coffee Lake 2.8GHz Intel Xeon E3-2276H six-core processor with up to 4.7GHz turbo boost, a 12MB cache and a thermal design power (TDP) rating of 45 watts.

That system, which includes an NVIDIA Quadro RTX 5000 graphics card, costs $4,289 after online discounts, but our system included several enhancements. For example, our evaluation unit came with 64GB of memory, installed as a pair of 32GB DDR4 2666MHz error-code correcting memory modules, which added $690 and left two memory sockets available for future expansion.

The NVIDIA Quadro RTX 5000 graphics card in our evaluation unit provides 3,072 compute unified device architecture processing cores, 48 RT cores, 384 Tensor cores and includes 16GB of GDDR6 memory, while consuming 80-110 watts. Its 256-bit memory interface provides a 448 GB/second memory bandwidth.

All ThinkPad P53 systems include two M.2 drive sockets, while some systems based on Intel Core processors also can accommodate a third 2.5-in. data drive. Our evaluation unit came with a 1TB Samsung M.2 PCI3-NVMe Opal2 drive, which added $205. A similar 2TB SSD costs $465. A second 1TB M.2 drive would have added an additional $325 and you can have the second drive configured as a RAID array. Our system also included an infrared and 720p camera ($20).

Unsurprisingly, the P53 offers lots of connectivity. The left side provides an HDMI port, a pair of USB 3.1 ports (including one always-on port to charge a USB device whenever the computer is connected to AC power) and a four-in-one media card reader. The right side includes a security lock slot, a SIM tray, a USB Type-C port and a combination microphone/headphone audio jack. The rear panel provides an RJ45 Ethernet jack, two USB Type-C Thunderbolt ports, and the connector to the external power supply.

Dual-band Wi-Fi and Bluetooth come standard and the system is powered by a six-cell 90 watt-hour battery that kept our ThinkPad P53 running for 5 hours and 30 minutes, about the same as the P72 we tested earlier this year. The Lenovo mobile workstation remained cool and quiet throughout our tests.

Great Performance

Lenovo workstations have a history of delivering great performance and the ThinkPad P53 continued to live up to expectations.

On the SPECviewperf benchmark, which focuses on graphics, the P53 scored near the top on all datasets, turning in the best results we have ever recorded on several of those tests. The 15.6-in. Lenovo also did very well on the SPECapc SolidWorks benchmark. In fact, the only systems that outper-

Mobile Workstations Compared	**Lenovo ThinkPad P53** 15.6-inch mobile workstation (2.80GHz Intel Xeon E-2276M 6-core CPU, NVIDIA Quadro RTX 5000, 64GB RAM, 1TB NVMe PCIe SSD)	**BOXX GoBOXX SLM17** 17.0-inch mobile workstation (2.30GHz Intel Core i9-9880H 8-core CPU, NVIDIA Quadro RTX 3000, 32GB RAM, 512GB NVMe PCIe SSD)	**HP ZBook 14u G6** 14.0-inch mobile workstation (1.90GHz Intel Core i7-8665U 4-core CPU, AMD Radeon Pro WX3200, 32GB RAM, 512GB NVMe PCIe SSD)	**MSI WS65 9TM-856** 15.6-inch 2.60GHz Intel Core i7-9750H 6-core CPU, NVIDIA Quadro RTX 5000, 32GB RAM, 512GB NVMe PCIe SSD	**Dell Precision 3541** 15.6-inch 2.60GHz Intel Core i7-9750H 6-core CPU, NVIDIA Quadro P620, 16GB RAM, 512GB NVMe PCIe SSD	**Dell Precision 3540** 15.6-inch 1.80GHz Intel Core i7-8565U 4-core CPU, AMD Radeon Pro WX 2100, 16GB RAM, 512GB NVMe PCIe SSD
Price as tested	$5,338.00	$4,200.00	$2,649.00	$3,499	$2,0687	$1,782
Date tested	**10/24/2019**	10/23/2019	8/8/19	7/12/19	7/3/19	7/3/19
Operating System	Windows 10 Pro 64	Windows 10 Pro 64	Windows 10 Pro 64	Windows 10 Pro 64	Windows 10 Pro 64	Windows 10 Pro 64
SPECviewperf 13.0 (higher is better)						
3dsmax-06	181.47	148.65	36.42	169.25	47.53	16.54
catia-05	269.51	200.15	35.60	213.02	53.63	31.56
creo-02	255.96	185.52	34.17	210.09	52.16	15.27
energy-02	38.63	29.94	2.61	39.87	8.25	0.42
maya-05	261.90	187.67	34.17	206.74	56.88	35.19
medical-02	85.31	63.59	9.73	80.88	12.63	8.52
showcase-02	63.79	79.50	13.93	92.57	21.87	11.20
snx-03	223.64	218.39	52.78	288.08	71.37	41.00
sw-04	88.51	123.98	46.04	123.16	60.95	49.00
SPECapc SolidWorks 2015 (higher is better)						
Graphics Composite	5.47	5.03	2.27	3.73	4.16	4.14
Shaded Graphics Sub-Composite	3.53	3.01	1.39	2.23	2.69	2.28
Shaded w/Edges Graphics Sub-Composite	4.38	3.89	2.06	2.96	3.51	3.09
Shaded using RealView Sub-Composite	4.05	3.57	1.63	2.63	3.05	2.87
Shaded w/Edges using RealView Sub-Composite	4.73	4.35	2.98	3.12	3.50	4.87
Shaded using RealView and Shadows Sub-Composite	4.59	4.11	1.45	3.04	3.48	3.02
Shaded with Edges using RealView and Shadows Graphics Sub-Composite	4.71	4.56	2.47	3.30	3.66	4.83
Shaded using RealView and Shadows and Ambient Occlusion Graphics Sub-Composite	15.06	13.54	3.01	10.06	10.13	7.75
Shaded with Edges using RealView and Shadows and Ambient Occlusion Graphics Sub-Composite	14.58	13.35	4.58	9.59	9.73	11.27
Wireframe Graphics Sub-Composite	3.92	4.34	2.76	3.50	3.69	3.50
CPU Composite	5.32	5.33	1.85	2.71	4.17	3.37
SPEC Workstation v3 (higher is better)						
Media and Entertainment	2.07	1.98	0.8	1.82	1.37	0.92
Product Development	2.24	2.07	1.04	2.01	1.57	1.21
Life Sciences	1.77	1.99	0.87	1.97	1.00	0.94
Financial Services	1.69	2.16	0.70	1.49	1.41	0.80
Energy	1.37	1.32	0.54	1.28	0.95	0.50
General Operations	1.89	1.79	1.13	1.75	1.67	1.52
GPU Compute	3.31	3.09	0.60	3.41	1.00	0.35
Time						
AutoCAD Render Test (in seconds, lower is better)	49.20	45.90	140.40	43.80	59.70	77.50
Battery Life (in hours:minutes, higher is better)	5:30	8:37	5:30	6.07	15:28	15:17

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

Fig. 3: Graphic Performance chart based on SPECviewperf 13 Geomean.

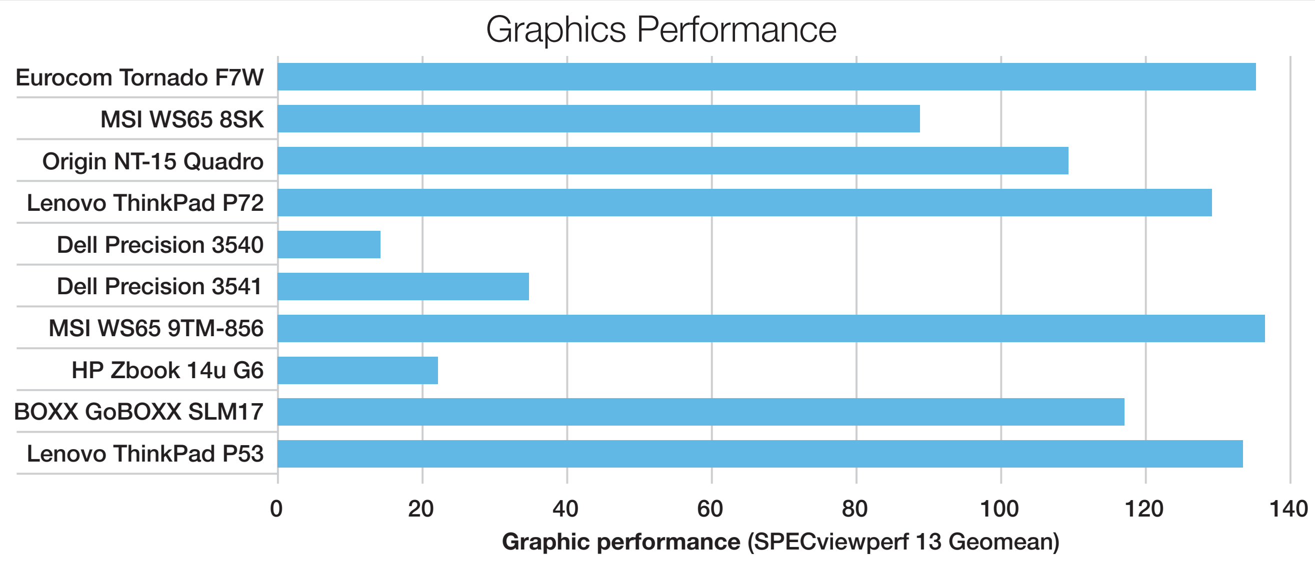

Fig. 3: Graphics performance chart based on SPECviewperf 13 Geomean.

formed the P53 were 17-in. systems with over-clocked CPUs.

On the very demanding SPECwpc workstation benchmark, the Lenovo ThinkPad P53 also delivered excellent results, again garnering top marks in several categories. Its graphics subsystem score was the best for any mobile workstation and its storage score was much better than the ThinkPad P72, thanks to Lenovo's switch back to a Samsung M.2 drive.

Although its 49.2-second average to complete our AutoCAD rendering test was a few seconds slower than some of the other mobile systems we have tested recently, those systems were equipped with faster CPUs with more CPU cores, giving them a decided edge on this multi-threaded test.

Although base configurations with Intel Core processors come with Windows 10 Home, Lenovo also offers Windows 10 Professional for just $35 more. Since the system we received included a Xeon CPU, our evaluation unit included Windows 10 Pro for Workstations.

Lenovo's standard warranty now covers its systems for just one year, with depot or carry-in service. Additional coverage is available at the time of purchase that can extend the warranty for up to five years, including accidental damage protection, on-site service and premier support. Since most other workstation-class machines come with a three-year warranty, we boosted the standard warranty to three years for pricing purposes, which added $109.

The ThinkPad P53 is certified for applications from vendors including ANSYS, Autodesk, Dassault Systèmes, PTC and Siemens. You can purchase a preconfigured system or build a custom configuration on the Lenovo website, where our P53 priced out at $5,338 after an automatic online discount, making it one of the more expensive systems we have tested recently.

But its performance elevates it to rarified air—we've tested only one other mobile workstation that outperformed the P53 and it cost several thousand dollars more.

The P53 configuration we tested certainly isn't for everyone. Less expensive components could significantly reduce the price without dramatically reducing performance. But if you require a top-of-the-line mobile CPU and GPU, the Lenovo ThinkPad P53 certainly delivers. **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 www.dscohn.com.

→ MORE INFO

Lenovo: Lenovo.com

Lenovo ThinkPad P53

- **Price:** $5,338 as tested ($1,489 base price)
- **Size:** 14.86x9.93x1.16-in. (WxHxD) notebook
- **Weight:** 5.89 lbs. (plus 1.95-lb. external power supply)
- **CPU:** Intel Xeon E3-2276M 2.80GHz 6-core w/12MB cache
- **Memory:** 64GB ECC DDR4 at 2666MHz
- **Graphics:** NVIDIA Quadro RTX 5000 w/16GB GDDR6 memory
- **LCD:** 17.3-in. HD (1920x1080) IPS
- **Hard Disk:** 1TB SSD M.2 2280 PCIe NVMe Opal2
- **Floppy:** None
- **Optical:** None
- **Audio:** Built-in speakers, headphone/microphone jack, built-in microphone array
- **Network:** Integrated Intel AX200 11ax 2x2 plus Bluetooth 5.0, one RJ45 gigabit Ethernet port
- **Modem:** None
- **Other:** Two USB-A 3.1, one USB-C, two USB-C Thunderbolt 3, HDMI, 4-in-1 media card reader, SIM tray, IR camera and 720 web cam
- **Keyboard:** Integrated 105-key full-size backlit keyboard with numeric keypad
- **Pointing Device:** integrated touchpad with three buttons, pointing stick with three buttons, fingerprint reader

For more information on this topic, visit DigitalEngineering247.com.

EDITOR'S PICKS

Each week, DE's editors comb through dozens of new products to bring you the ones we think will help you do your job better, smarter and faster. Here are our most recent musings about the products that have really grabbed our attention.

Engineering GPU with World-First Specs

AMD Radeon Pro W5700 is a professional PC workstation graphics card.

The Radeon Pro W5700 professional PC workstation graphics card is designed for 3D designers and engineers using 7nm technology. It combines GDDR6 memory and PCIe 4.0 bandwidth. Together, these make it possible to work fast with larger datasets. If you already create photorealistic product renderings or virtual reality environments for design review, this GPU will do it faster. One more industry first for the AMD Radeon Pro W5700 is the addition of a USB-C connector.
MORE → digitalengineering247.com/r/23428

Bridging Gap Between Design and Additive

MSC built Apex Generative Design to boost productivity workflow.

MSC Apex Generative Design combines elements of MSC Simufact (metal) and MSC Digimat (polymer) for build process simulation. In use, a designer only needs to specify boundary conditions and design objectives. The software can then produce multiple lightweight design candidates.

Apex Generative Design combines manufacturing knowledge with automated generative design. It can evaluate thousands of alternatives before suggesting options to the user. **MORE →** digitalengineering247.com/r/23396

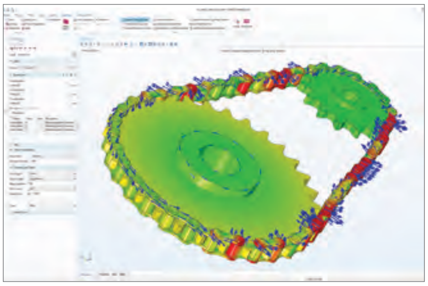

Update to COMSOL Multiphysics Is Now Available

Release plays well with others.

The big draw for COMSOL Multiphysics has always been its ability to account for coupled or multiphysics phenomena. Also, it interfaces well with other CAD and CAE products.

The Design Module gains a new parametric sketching tool. The user can assign dimensions and constraints to planar drawings for 2D models and 3D work planes. This new sketching tool extends the capabilities of the Model Builder, making it easier to use model parameters to drive the simulation.
MORE → digitalengineering247.com/r/23355

Automate New Design Space Workflows

nTop Platform 2.0 designed to help engineers design performance parts faster.

nTopology's nTop platform is designed for importing relevant engineering data from various sources into a single workspace where the data becomes a single model. The geometry engine uses implicit modeling, where data drives shape instead of shape driving data.

nTopology says it is not creating a product to replace CAD, CAE, CAM or PLM, but to augment them. nTop Platform can bring in data from all these formats and work with them in a unified environment.
MORE → digitalengineering247.com/r/23327

Next-Gen Engineers

Student Design Competition Profile: FlexFactor's Final Pitch Event

Shark Tank-Style Program Helps Students Kickstart STEM Careers

BY JIM ROMEO

FlexFactor's Final Pitch Event is a one-day competition held at the end of FlexFactor's 4-week technology and entrepreneurship program. This multi-week program challenges students to work in teams to identify a real-world problem, conceptualize an advanced hardware solution and build a business model around it that they present at the final pitch event. Brynt Parmeter is the director of workforce development, education and training at NextFlex. We spoke to Parmeter to learn more about their competition.

FlexFactor focuses on design and marketing new solutions. *Image courtesy of NextFlex.*

Digital Engineering: **Can you provide an overview of your competition, how it came to be and its intent?**

Brynt Parmeter: The program was started in the Silicon Valley in 2016 and within only two years, the program has expanded from eight students to over 3,500 students in three geographic regions: California, Ohio and Alabama. The core intent of this program is to inform, attract and recruit young people from unengaged to engaged talent and motivate them to pursue education and career pathways in technology and advanced manufacturing.

FlexFactor inspires middle school, high school and higher education students from all over the country to go into fields of work they might not have known of without completing the program while also encouraging female and minority student involvement in tech and advanced manufacturing.

DE: **Can you tell us about some of the designs that are part of the event and how they came to be?**

Parmeter: The FlexFactor program gives students the opportunity to see how advanced materials and products can combine to create innovative solutions to a diverse range of problems. It also helps kickstart their advanced manufacturing career paths by familiarizing them with the technology and the vast range of professional opportunities in the field.

Throughout the program, students gain hands-on experience applying cutting-edge flexible electronics to real-world problems, and acquiring industry exposure to further explore their career paths by visiting advanced manufacturing facilities, and by seeing how advanced technologies are being used today and what it's like to work in the technical fields.

At the end of the program, teams pitch their product ideas "Shark Tank" style with guidance from industry mentors and educators. After completion of the program, students receive further information from participating colleges regarding pursuing STEM-related majors for future careers in technology and advanced manufacturing.

DE: **Can you provide some examples of what the event has produced or what you expect it to produce?**

Parmeter: Last month, NextFlex co-hosted FlexFactor Final Pitch Event: Silicon Valley, in partnership with Evergreen Valley College. During the finals, top student teams from across Silicon Valley delivered a hardware device "pitch" on their stated problem area, their solution for that problem and a business model to a panel of industry experts.

DE: **Does your organization, or the consortium of organizers of this competition, have a particular stance on adopting an innovation that is linked to the program?**

Parmeter: At the core of the FlexFactor program is flexible hybrid electronics (FHE). NextFlex focuses on FHE technology because … conceptually speaking, it is a fairly easy technology to understand both in how it is created and how it will change the world around us. It provides an excellent "innovation vehicle" for student teams looking to explore how advanced technologies will be able to address the most series challenges of the future.

DE: **Is there anything else you'd like to tell us about the event?**

Parmeter: The FlexFactor program taps into the potential in youth whose circumstances sometimes have isolated them from the experiences, networks, pathways and opportunities that so frequently determine future economic success and social mobility.

Through FlexFactor, the talent pipeline is diversified as students who previously hadn't considered higher education or STEM-based careers muster the courage to realize their potential. **DE**

Jim Romeo *is a freelance writer based in Chesapeake, VA. Send e-mail about this article to de-editors@digitaleng.news.*

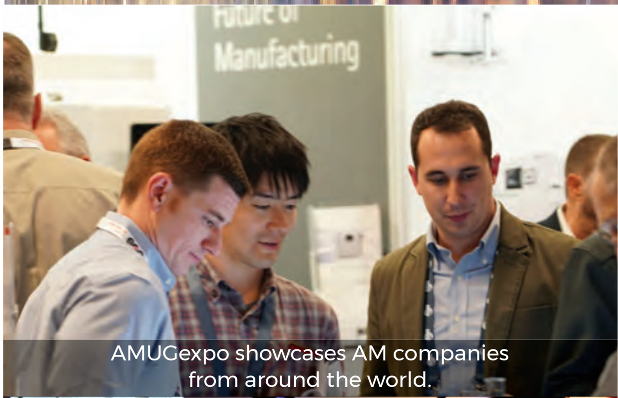
AMUGexpo showcases AM companies from around the world.


2019 Innovator Recipient Winner Gideon Levy.


Diamond and Platinum sponsor sessions all week.


Technical Competition highlights advanced concepts and finishing.

REGISTER TODAY!
WWW.AMUG.COM

AMUG

2020 AMUG CONFERENCE
HILTON CHICAGO
CHICAGO, ILLINOIS
MARCH 22 - MARCH 26, 2020

ADDITIVE MANUFACTURING USERS GROUP
32ND ANNUAL EDUCATION & TRAINING CONFERENCE

- Technical Presentations and Panel Discussions
- AMUG Tracks include:
  - Aerospace, Defense, Military and Transportation
  - AM Technologies (Non-Metal and Metal)
  - Casting
  - Education and Training (Design/Software, Post Processing, Finishing, Tips and Tricks)
  - Materials (Plastic and Metal)
  - Medical and Dental
  - Software
  - Scanning and Metrology
- Daily Diamond and Platinum Sponsor Sessions
- Hands-on Workshops for Plastics and Metals
- AMUGexpo Showcases New Developments in AM
- Networking with AM Industry Experts and OEMs

2020 AMUG DIAMOND SPONSORS

For a complete list of sponsors and exhibitors visit www.amug.com

3D SYSTEMS Additive Manufacturing Solutions
3YOURMIND
BASF We create chemistry FORWARD AM
Carbon®
COMPUTER AIDED TECHNOLOGY
DSM BRIGHT SCIENCE. BRIGHTER LIVING.
Dyndrite
eos
ESSENTIUM
ExOne
formlabs
hp
RENISHAW apply innovation™
SLM SOLUTIONS
stratasys
VELO 3D