Case Study

High Performance Cloud Computing
2nd Generation Intel® Xeon® Scalable Processors
Intel® software tools

OnScale and Google Cloud Make Complex Digital Prototyping Accessible and Affordable

2nd Generation Intel Xeon Scalable processors, along with Intel software and AI technologies, power OnScale solutions.

Solution Summary:

- OnScale Multiphysics solvers for modeling and simulation
- Digital prototyping for next-generation biomedical, MEMS, semiconductor, and other new technologies
- Google Cloud large-memory (up to 12 TB) M2 instances
- Google Cloud compute-optimized C2 instances built on 2nd Gen Intel Xeon Scalable processors

Executive Summary

Digital prototyping using modeling and simulation is becoming standard practice in today's engineering processes. But for large and complex device designs, digital prototyping involves 3D Multiphysics solvers that run on expensive supercomputers. OnScale enables digital prototyping using their long-trusted and established Multiphysics solvers running on Google Cloud clusters, which feature 2nd Gen Intel® Xeon® Scalable processors. OnScale's Intel-and Google Cloud-powered service helps companies accelerate products to market and cut R&D costs by offsetting costly physical prototypes with highly accurate digital prototypes.

Challenge

Digital prototyping is a critical part of design and engineering today. Compared to classic physical prototyping, test, and verification processes, digital prototyping saves money, accelerates R&D timelines, and reduces risk in new technology development.

But digital prototyping for complex designs requires solving extremely large Multiphysics problems—electrical, thermal, mechanical, material, and others—simultaneously. The more complicated the problem and the higher accuracy required, the greater demand for higher and higher levels of computing. Today's workstations, with as many as 48 or more CPU cores, help engineers understand a portion of the design space. But without the compute capabilities of supercomputers, engineers must make intelligent approximations about more complex design issues to fit simulation workloads.

For large and complex product designs, engineers need the power of supercomputers. In these engineering applications, many physical phenomena are simulated concurrently in 3D using Multiphysics solvers to understand the entire design. And highly detailed digital prototyping studies with multi-parametric sweeps across design options help engineers optimize new, complex designs quickly.

Large computers and the engineering software licenses to run on them are expensive, even for enterprises. Fixed in-house simulation capacity also limits how many engineers can access simulation resources because of the large size of simulation workloads, limited local computing capacity, and availability of fixed per-user simulation software licenses. And scaling in-house machines by adding more licenses and computing capacity add complexity for the IT department to deploy and manage both the hardware and software.

The hardware and software limitations of traditional on-premise or desktop computing create barriers for engineers, which constrains their innovative potential.

Solution

OnScale, a new Software-as-a-Service (SaaS) cloud engineering simulation platform, runs their long-established and trusted Multiphysics solvers on-demand on Google Cloud. Their solution provides comprehensive digital prototyping capabilities powered by 2nd Generation Intel® Xeon® Scalable processors to engineers in organizations of all sizes around the world.

"Our Multiphysics solvers describe the physical universe digitally," explained Ian Campbell, CEO of OnScale. "Our solvers have been used for decades by civil engineers, semiconductor designers, mechanical engineers, and engineers in government agencies and enterprises. They need the combined resources of complex Multiphysics solvers running on supercomputers."

Multiphysics on Scalable Google Cloud Simplifies Digital Prototyping

Originally developed by engineering firm Weidlinger Associates (now part of Thornton Tomasetti), OnScale adapted the solvers for harnessing the power of the cloud. The solvers were designed to run on massively parallel computing systems using the Intel Message Passing Interface (MPI) library. OnScale optimized them to scale across thousands of nodes on Google Cloud. This innovation makes the solvers ideal for flexible, on-demand deployments in a cloud environment.

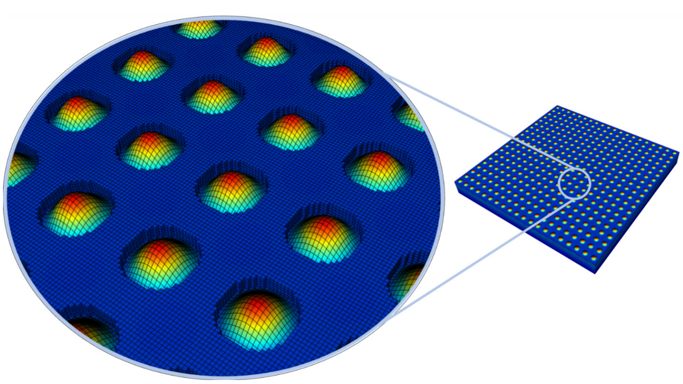

Figure 1. Parametric full 3D piezoelectric, structural, and acoustic simulation of a MEMS PMUT for ultrasonic fingerprint sensor applications. Simulation results courtesy of OnScale.

A SaaS business model allows companies of all sizes to take advantage of powerful simulation capabilities without the complexities and costs of software and hardware in-house. And, with advances in scaling out Google Cloud instances in recent years, quickly adapting compute capacity in cloud environments to workload demands is simpler and faster. The result is gains in speedup far outpace the cost of additional cores to achieve those gains. Engineers can take advantage of faster completion times without dramatic increase in costs (Figure 2).

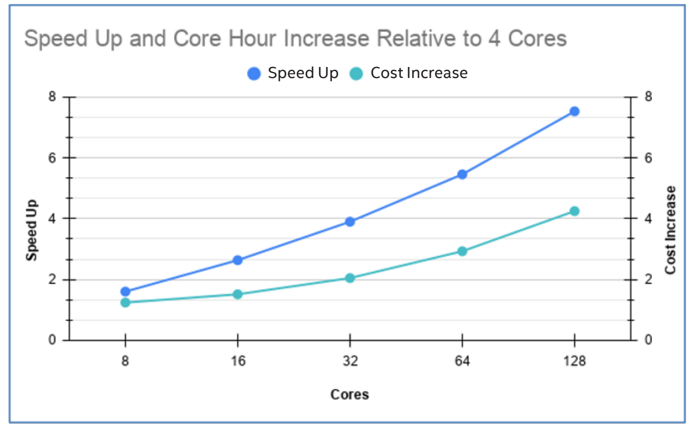

Figure 2 illustrates an OnScale mechanical simulation with 2 million degrees of freedom (DoFs) executed on cloud hardware with a varying number of physical cores. For an approximate 4X increase in computation cost (based on core hours), the job ran about 8X faster, reducing runtime from about 11 minutes to 1.5 minutes.

AI Drives Cluster Configuration

"Our simulation platform mainly runs on compute-optimized C2 and memory-optimized M2 Google Cloud instances," added Campbell. "We've benchmarked many of these Google Cloud configurations, so we know how our simulation software scales on cloud hardware. We've created a machine learning engine and trained our models on one-half million simulations. That means when an engineer sets up a new simulation study, our AI can create the best cloud configuration based on his or her needs, optimizing for accuracy, cost, and runtime."

Engineers define their simulation, upload 3D models, and choose the priorities of their study—accuracy, budget, and time to solution. OnScale's AI then provides an estimate of accuracy, cost, and runtime to complete the simulation. When the engineer executes the job, OnScale's Cloud Orchestrator builds a custom state-of-the-art cloud supercomputer configuration with Intel Xeon Scalable processors for the workloads and selected priorities. For example, it uses 2nd Gen Intel Xeon Scalable processors for fast solutions with high accuracy, if that is the priority. With Google's Kubernetes Management Engine, engineers experience very little wait time to begin building their digital prototypes with OnScale.

OnScale used Intel software tools to compile and optimize their solvers for Intel architecture to achieve high performance and code efficiency, including:

- Intel MPI Library
- Intel Math Kernel Library
- Intel Fortran Compiler

"Over the years, we've added new capabilities and features to our solvers and platform, such as how we perform advanced domain composition and dynamic load balancing," commented Campbell. "For example, during a simulation, if the workload is reaching the limits of the cluster configuration, we can pause the simulation. Then we build a new cluster with more resources and restart the simulation from where it was paused—all transparently to the engineer."

The OnScale platform monitors its own performance and processes. After a simulation study completes, metadata about platform performance is fed back to the machine learning engine to further tune OnScale's ability to accurately predict the best cluster configurations for future simulations.

"We're looking at new Intel technologies closely," explained Campbell. "Google Cloud M2 instances with large memory are interesting. Large memory will allow us to move seamlessly from running a massive simulation on a large cluster to visualization of simulation results on a much smaller machine without having to move data from memory to storage and then reload it back to memory in a new container. That can save a significant amount of time and reduce costs."

OnScale is also exposing AI and ML capabilities for their end customers to use in combination with Multiphysics simulation.

AI and ML are integral to our service offerings for our customers," said Campbell. "For example, engineers can use simulated datasets to train embedded AI algorithms, saving an enormous amount of time and cost in creating physical datasets from physical prototypes. We're looking at incorporating TensorFlow and other machine learning methods directly into our offerings for our customers to leverage. Technologies like Intel Deep Learning Boost are very interesting to power our solutions."

Result

The accuracy of digital prototyping is best proven by correlating simulation results with experimentation. Polytec, a world leader in optical measurement solutions, develops, produces, and distributes instruments and provides non-destructive testing (NDT) services for industry and research. Their optical measurement solutions capture the direct mechanical response with high accuracy and sensitivity, allowing engineers to evaluate physical prototypes—properties of materials, structures, and devices used in aerospace, automotive, medical, nanotechnology and other industries. After engineering and digital prototyping using OnScale to optimize a design, engineers turn to Polytec equipment to validate physical prototypes with experimentation.

Polytec showed the accuracy of OnScale tools by comparing the simulation results of a digital prototype to measurements from a physical prototype. They ran experiments on a physical stainless steel test block and compared measurements to a simulation. In this experiment, a transducer mounted on the block injected high-frequency elastic waves into the material and captured reflections from the holes drilled into the block. Correlating simulation results with physical measurements

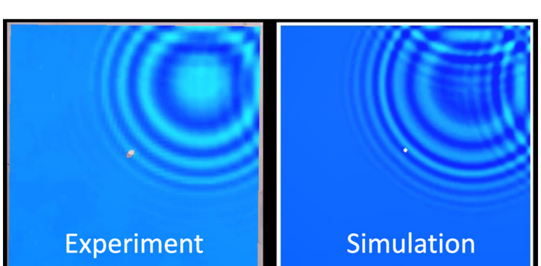

Figure 3. Polytec Laser Doppler Vibrometer Experiment Results versus OnScale Simulation Results (Courtesy of OnScale)

is shown in Figure 3. The simulation results from the digital prototype accurately predicted the experimental data of the physical prototype captured by Polytec.

"Ultimately, what the OnScale, Intel, and the Google Cloud partnership has produced is a full-stack engineering simulation solution that is magnitudes faster than comparable desktop solutions."

– Ian Campbell, CEO, OnScale

Researchers at the University of Washington are using OnScale to design breakthrough eye imaging technology for optical coherence elastography (OCE). This non-contact, noninvasive method allows clinicians to quantify and detect changes in corneal elasticity and intraocular pressure to better identify and monitor opthalmic conditions (e.g. keratoconus and glaucoma), and guide therapeutic interventions (e.g. LASIK, photorefractive keratectomy, and corneal crosslinking). Their technology noninvasively measures mechanical properties of the cornea using propagating mechanical waves over the cornea excited with an air-coupled acoustic transducer through air. Using optical coherence tomography (OCT), they can image propagating mechanical waves which allows mapping corneal elasticity.

Of particular importance to OCT technology is propagation of shear waves. The research team uses OnScale to model excitation and propagation of shear waves in simulated biological tissue—a digital prototype of the human eye—of which the cornea presents unique challenges. With OnScale, researchers were able to construct a reliable two-dimensional finite element model that closely mirrors their experimental OCE system and measurements. Their ongoing research is described in "Acoustic micro-tapping for non-contact 4D imaging of tissue elasticity" and "Nearly-incompressible transverse isotropy (NITI) of cornea elasticity: model and experiments with acoustic micro-tapping OCE".

Vestas designs and manufactures wind turbine solutions for the energy industry. A major component is the very large blade (up to 80 meters in length) made from multiple composite materials—plastics, carbon fiber, resin, and others. The blades are subject to many stresses across their length and especially at the mount. Vestas had to design various tools, such as ultrasound testing, to validate the design and quality of manufacturing and look for anomalies. But designing ultrasound testing for composites is not as well understood as it is for other materials.

"There are many complicated phenomena when using ultrasound on composites," explained Jason Hawkins, a testing engineer at Vestas. "Depending on which direction the wave is moving, you have different velocities, and that can be very complicated to read on your screen. You can improve the detection process if you can actually simulate it. You can see what processes, such as reflected primary and secondary waves, are resulting in what you're measuring. We couldn't do that before."

Having the capability to run numerical simulations gives engineers new insight into how the material is responding to ultrasound excitation for NDT. With this new knowledge, they can create inspections that were not possible before and improve design and manufacturing.

Solution Summary

Digital prototyping accelerates development timelines while saving costs by reducing reliance on physical prototypes. But digital prototypes can require compute-intensive calculations possible only with supercomputers and advanced Multiphysics solvers, which are expensive to acquire and operate in-house. OnScale provides trusted Multiphysics solvers that run on-demand on the Google Cloud on Intel Xeon Scalable processors, making these powerful capabilities accessible to more organizations.

"Ultimately, what the OnScale, Intel, and the Google Cloud partnership has produced is a full-stack engineering simulation solution that is magnitudes faster than comparable desktop solutions," concluded Campbell. "This means engineers can be more productive and bring us future technology with reduced risk, cost, and time-to-market".

Where to Get More Information

Learn more about OnScale cloud simulation solutions.

Find out more about Google Cloud.

Explore the capabilities of the 2nd Generation Intel Xeon Scalable processors with integrated Intel® Deep Learning Boost capabilities for accelerated AI inferencing.

Solution Ingredients

- OnScale Multiphysics solvers and SaaS cloud engineering simulation
- Digital prototyping solutions for industrial engineering and research
- Google Cloud M2 instances
- Google Cloud C2 instances built on 2nd Gen Intel Xeon Scalable processors