CommetCo | Hewlett Packard Enterprise

Cognitive Powers — AI and Supercomputing Are Driving a New Era of Discovery and Innovation

Can machines think?

To someone using a pocket calculator, it certainly seems that way. The first one came on the market in the 1970s. That's not so long ago. The Beatles had only recently broken up, most cars drove with rear wheels, and Neil Armstrong had already taken his famous "small step" on the lunar surface. Anyone reading this who remembers a time before the pocket calculator knows that math problems or even anything beyond simple arithmetic required a sharp brain, and usually a sharp pencil. In that seemingly neolithic era, school children actually had to learn how to solve math problems—and even memorize multiplication tables.

The advent of the pocket calculator seemed like a small miracle. Suddenly we had palm-sized devices that even kids could afford yet seemed to have more brain power than humans—at least when it came to crunching numbers. However, anyone with a basic understanding of computer science knows machines like a pocket calculator, the first personal computers, or today's remarkable digital devices like Amazon's new domestic robot only appear to have cognitive powers. Their apparent intelligence is a mirage.

What actually happens at the level of software and chips in these smart machines is mind-numbingly simple. It just seems like intelligence because it happens at lightning speed. Let's have a look.

Unlike the Hindu-Arabic number system we learned as children, which uses ten numerals—the digits 0, 1, 2, 3, 4, 5, 6, 7, 8, 9—calculators and computers are based on the Binary number system. It uses only two digits, zero and one. While the inner workings of semiconductor chips and software that make digital machines work is incredibly complicated, they rely on a simple invention—the transistor. A transistor is an electrical machine that controls the flow of electrons through gates that are either open or closed. Think of a transistor gate like a light switch that can have only two states, On or Off.

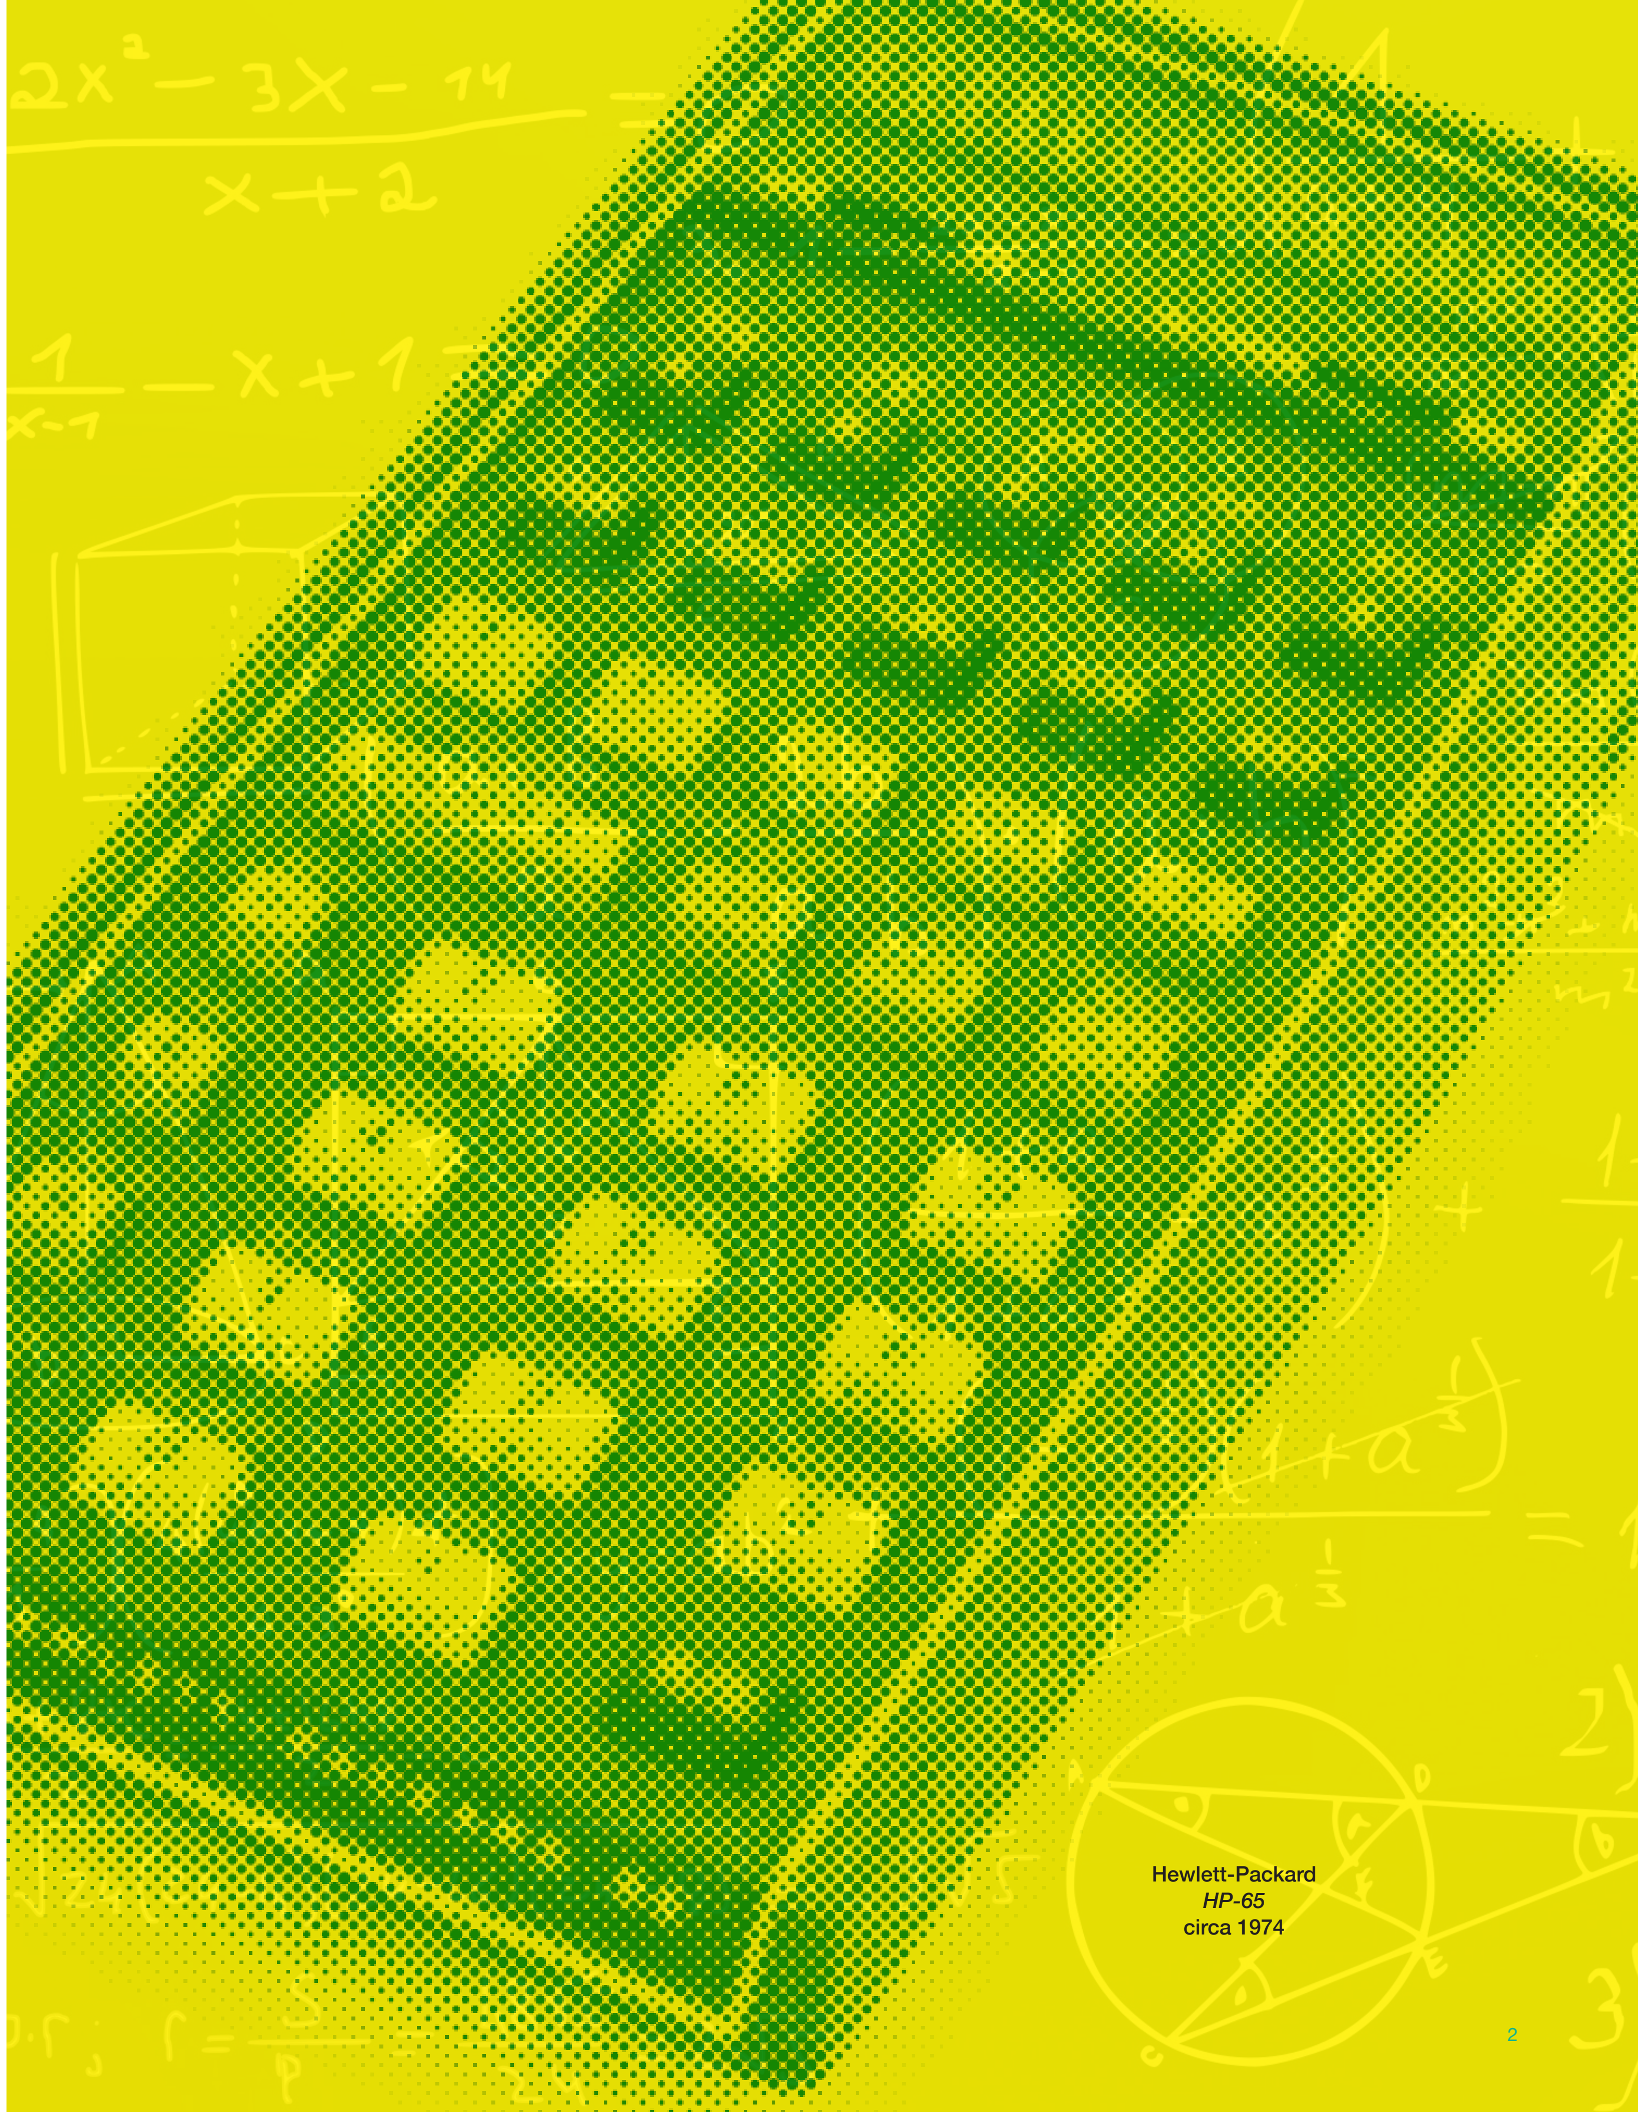

Hewlett-Packard
HP-65
circa 1974

Print enough microscopic transistors on a small chip of silicon and you have an integrated circuit (IC). Before an IC can become useful, however, the flow of electrons through the transistors' gates needs to be programmed in a pattern. Coders do this at chip hardware level with "machine language," which uses the Binary number system—using only 0 and 1—to turn the transistors' gates On or Off. A 1 in the code flips the switch on; a 0 flips it off. This programming can enable the chip to do things such as solving a simple problem like two plus two equals four.

Except in the case of very simple devices like pocket calculators, you need a vast number of transistors on a chip for much more than solving rudimentary arithmetic and math problems. Fortunately, it wasn't long before the chipmaking art had shrunk transistors to a point where you could have hundreds of thousands or even millions of transistors on a single piece of silicon the size of a postage stamp. An example of these massively integrated circuits is the central processing unit (CPU) that powered the first personal computers. And by creating programming languages that worked above the chip level, people could program these machines to do amazing things like running a spreadsheet application.

And today, as scientists have miniaturized transistors down to sizes measured in nanometers—billionths of a meter—we have massively integrated circuits that comprise whole computers on a chip. These system-on-a-chip designs (SoCs) power seemingly miraculous Lilliputian compute devices like your smartphone.

However, most computers and smartphones are not necessarily smart. It is the astonishingly fast rate of transistor gates opening and closing and the speed of electrons flowing over the silicon and along tiny wire interconnects that makes the remarkable abilities of digital devices seem like intelligence. In fact, without the help of humans they are dumb as a brick. You still need smart coders to program their problem-solving, picture-taking, or other useful abilities.

What in the world is Artificial Intelligence?

At the risk of boring computer scientist readers to tears, this elementary explanation of how computers work was necessary to help mere mortals understand one of today's hottest technology buzz phrases—Artificial Intelligence, or AI. Simply put, it involves humans programming computers and digital devices to make decisions and do things on their own. A key subset of AI, Machine Learning, involves machines using data to become "intelligent"—which means making better choices.

"Data is the source of a machine learning to be artificially intelligent," According Dr. Eng Lim Goh, Senior Vice President & Chief Technology Officer, AI, Hewlett Packard Enterprise. "In this one sentence, I'ved linked data, machine learning, and AI. The key is the phrase in the middle, machine learning. A machine is learning from data, and ultimately through that learning, it becomes smart enough to make an informed choice. Additionally, those choices sometimes reflect human choices, because the digital device or computer is trained on the data humans choose to give the machine. This explains why we say the machine is artificially intelligent."

> "Data is the source of a machine learning to be artificially intelligent."
>
> Dr. Eng Lim Goh
> Senior Vice President & Chief Technology Officer, AI, Hewlett Packard Enterprise

The concept of artificial intelligence and machines thinking for us is not new. Almost three-hundred years ago, Jonathan Swift described "The Engine" in his satirical novel Gulliver's Travels: "...whereas, by using this contrivance, the most ignorant person, at a reasonable charge, and with a little bodily labour, might write books in philosophy, poetry, politics, laws, mathematics, and theology, without the least assistance from genius or study." And as if Swift could see the future, the book's original sketch of The Engine looks remarkably like an IC.

In the modern computing world, the conversation about artificial intelligence began in 1950 with the publication of Alan Turing's "Computing Machinery and Intelligence." In that paper, Turing—who is often referred to as the "father of computer science"—posed the question: "Can machines think?" To answer the question he offered the now famous "Turing Test." It is remarkably simple: If a machine can engage in a conversation with a human expert without the human detecting it as a machine, the machine has demonstrated artificial intelligence.

How HPC gave AI a leg up

Trouble is, back in the 1950s computers did not really have enough horsepower to run AI applications in any meaningful sense. But that changed thanks to Moore's Law, which most of us know is the prediction by Gordon Moore, co-founder and chairman emeritus of Intel Corporation, that the number of transistors in a dense IC would double approximately every two years. For decades, Gordon Moore's prediction held true (Fig. 1). And the resulting growth of computing power is astonishing. While the first handheld calculator had a chip with approximately 2,000 transistors, we now have ICs packed with billions of transistors or more. In late 2020, Cerebras announced a Processing Unit (GPU) that has 1.2 trillion transistors.

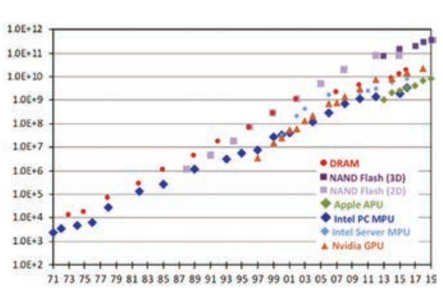

(Fig. 1) Moore's law predicted the doubling of transistors in a dense IC approximately every two years.

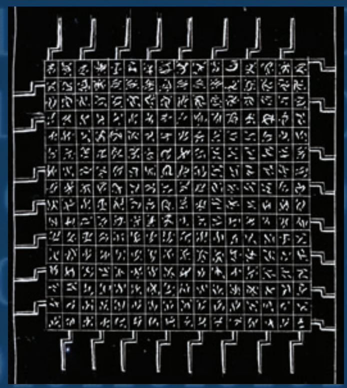

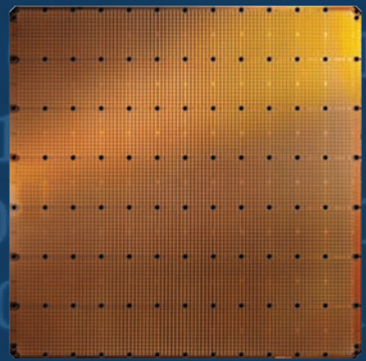

Johnathan Swift's
The Engine
circa 1726

Cerebras
GPU
circa 2020

Wisconsin becomes a new superpower

The potential of integrating many of these powerful chips was not lost on some of the world's top scientists. One was Seymour Cray who had a passion for the art of designing computers and loved the challenge of coaxing more and more speed out of circuits and wires. In 1972, he formed Cray Research in Chippewa Falls, Wisconsin with the vision of building the world's fastest computing systems. Cray succeeded with the Cray-1. A masterpiece of engineering, it measured 8½ feet (2.60 meters) wide by 6½ feet (2 meters) high, tipped the scales at 10,500 pounds (4.76 metric tons), and incorporated 60 miles (96 kilometers) of wires.

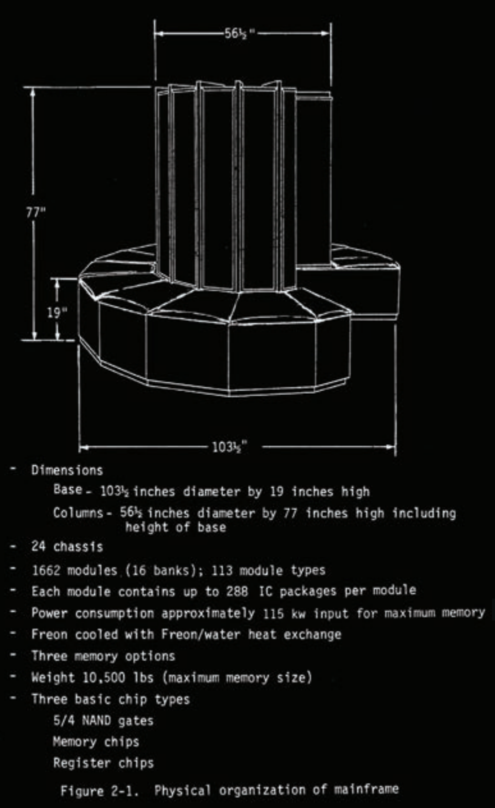

The Cray-1

Powered by newly designed ICs, the Cray-1 boasted more memory (one megabyte) and more speed (80 million computations per second) than any other computer in the world. And it retained the rank as the world's fastest supercomputer from 1976 to 1982. More importantly, it firmly established High Performance Computing (HPC) as the way to solve some of the world's toughest scientific and engineering challenges.

Hewlett Packard Enterprise (HPE) acquired Cray Inc. in 2019, and since then has made major investments in Chippewa Falls, sometimes called the supercomputing capital of the world. Today, HPE Cray EX super-computers continue to push the boundaries of HPC. The NERSC Perlmutter system at the Lawrence Berkeley National Laboratory is among the 10 ten fastest computers in the world. While not the first supercomputer, the Cray-1 helped to spawn a whole new field of high performance computing and a race for faster and faster speeds. As of mid 2021, the fastest supercomputer on the planet was Fugaku at the RIKEN Center for Computational Science in Kobe, Japan. With 7,630,848 cores, the Arm-based supercomputer achieves an impressive speed of 442 petaflops.

When you combine the enormous processing power of HPC with the potential of AI you get a new field of science. Today, this combination is ushering in a new age of insights that can benefit everything from engineering innovation to new medicines to helping meet basic human needs to breakthroughs like clean fusion energy.

HPC and AI are really taking off

Enabling this new era does not require gargantuan hardware. For example, Hewlett Packard Enterprise (HPE) tested the power of HPC in space by sending a mini supercomputer to the International Space Station (ISS) in 2017. Since the first crew arrived on ISS more than 20 years ago, scientist-astronauts have conducted countless hours of research for government agencies, academia, and commercial enterprises. And much of that requires computing power, which inspired a mission to send the first supercomputer beyond Earth – The Spaceborne Computer.

According to Hewlett Packard Enterprise and NASA, "The HPE Spaceborne Computer mission objective is to execute a long-duration operation of a Commercial Off-The-Shelf (COTS) high-performance computing (HPC) and Artificial Intelligence/Machine Learning (AI/ML) systems in the low Earth orbit internal ISS environment to determine conditions and parameters for continuance of correct computer operations."

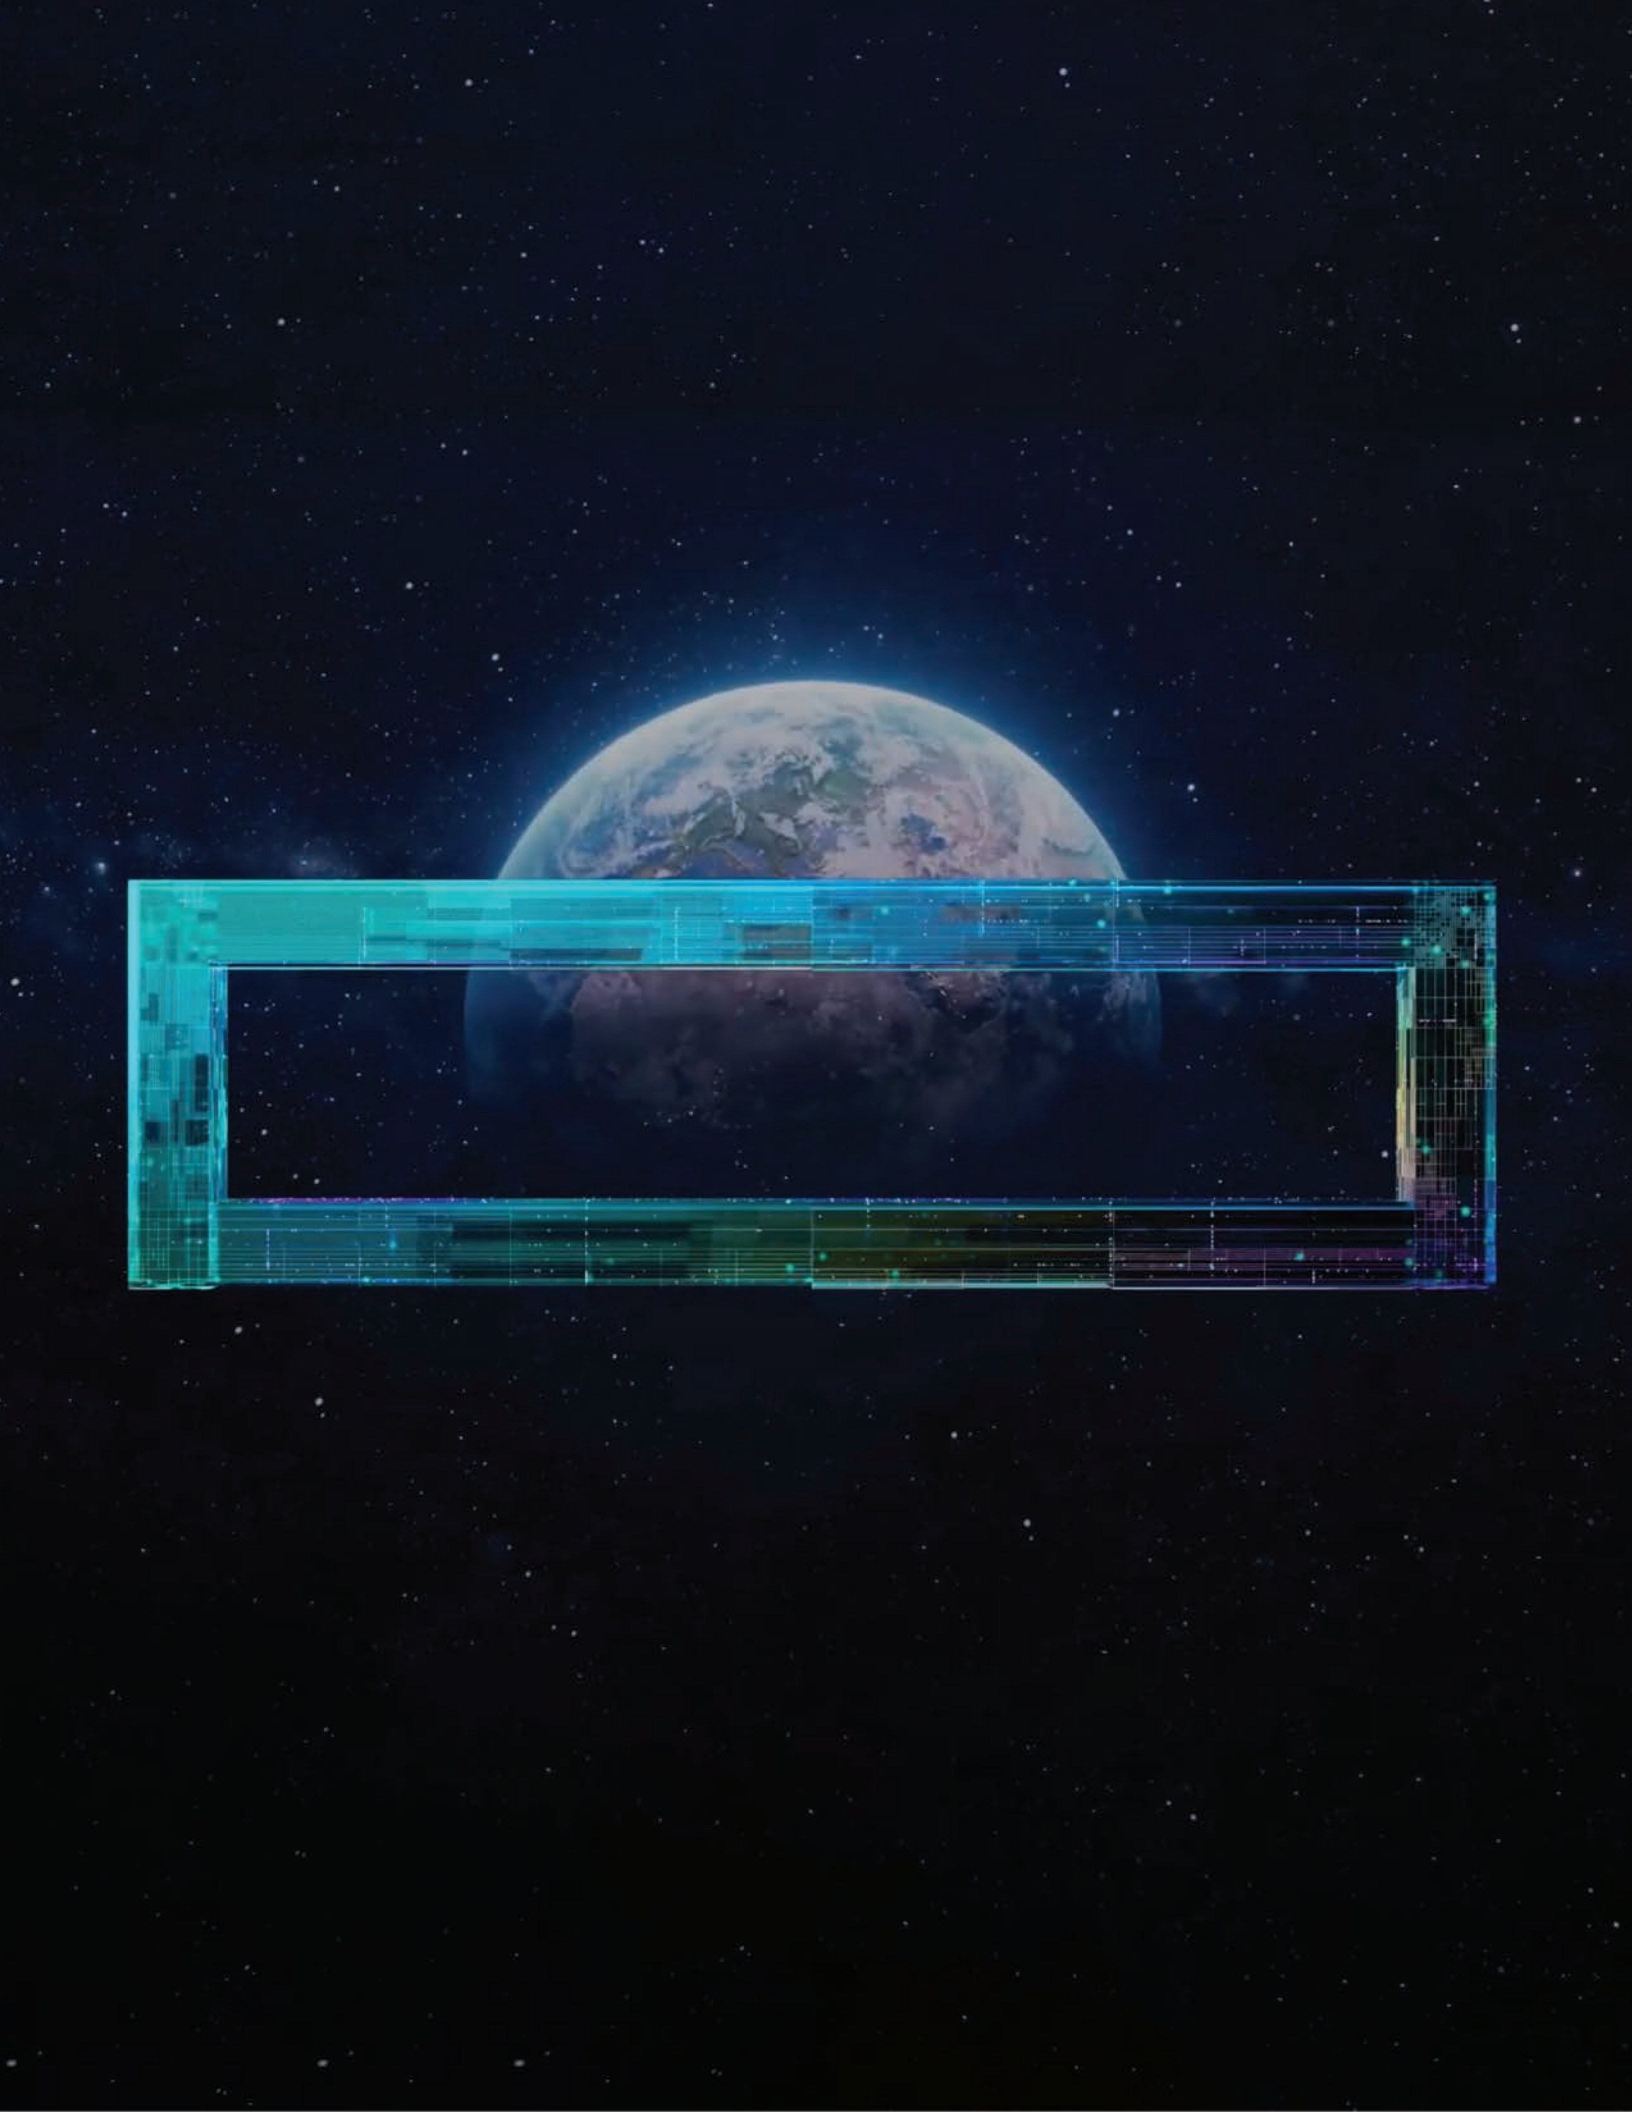

"We had many naysayers who claimed the supercomputing technology could not withstand the hazards of space travel," explained Dr. Mark Fernandez, Principal Investigator on the Spaceborne Computer-2 project. "They fretted that everything from G-forces during the rocket launch to cosmic radiation would destroy the hardware." In fact, prior to launch, Fortune magazine called it the computer that must survive 'the shake, rattle, and roll' of a space launch.

HPE further felt the success of the mission hinged on sending standard hardware to the ISS. "Computing in space has many challenges, such as cosmic radiation, which can cause computing bits in the computer's memory to flip from zero to one, or vice versa," according to Dr. Goh. "Many worried that we needed to harden the Spaceborne computer to withstand the much higher levels of radiation in space to prevent bit flip and other potential problems.

However, hardening or modifying such a complex system in any way could take years or even a decade when you include NASA's extensive qualification process. "Much of the compute hardware on the ISS is twenty years old and we wanted the Spaceborne Computer technology to be on the bleeding edge," remarked Dr. Goh. "That explains why the Spaceborne team recognized the mission's success hinged on using pure COTS hardware. Even modifying the computer to run on DC current, which is used on the space station, would risk not putting the state-of-the-art technology on the ISS."

In 2017, HPE successfully put a standard, off-the-shelf teraflop system on the ISS. "We made no modifications whatsoever to the COTS hardware," continued Dr. Goh. "For example, since engineers generally use software to minimize bit flip due to background radiation on Earth, we thought why not in space. The Spaceborne Computer not only uses similar software methods but also slows its compute speeds during episodes of high radiation to ensure proper operation. The DC power problem was solved by making a special DC/AC converter, which allows the standard AC hardware to run on DC power generated by the ISS's solar panels."

It all worked, and now the HPE supercomputer running on the ISS today proves it's possible to do a trillion calculations a second in space. "The Spaceborne Computer was not only the smartest and fastest computer in space, but thanks to its ability to use energy from solar power and cooling courtesy of the freezing temperatures in space it was also the most energy-efficient computer in existence," added Dr. Goh.

The Spaceborne Computer represented a huge leap forward, and its state-of-the-art speed continues to shorten time to discoveries that can benefit all of humanity. Nonetheless, proponents of deeper space exploration understood that computers in space, and now artificial intelligence, would be critical for these endeavors. For example, you need only travel 254 miles from the Earth's surface to reach the ISS, and you would need to travel approximately a thousand times that distance 238,900 miles (383,000 km) to reach the moon. But at its greatest distance in orbit, Mars lies 238 million miles (378,000,000 km) from the moon. That's deep space.

> "We had many naysayers who claimed the supercomputing technology could not withstand the hazards of space travel. They fretted that everything from G-forces during the rocket launch to cosmic radiation would destroy the hardware."
>
> Dr. Mark Fernandez
> Principal Investigator,
> Spaceborne Computer-2

Accelerating space exploration with state-of-the-art edge computing and AI capabilities

It can take as long as 40 minutes for radio waves to make the round trip from Mars to Earth. This results in extreme latency of data signals, which means that computers on human Mars missions or colonies would need to think for themselves—without direction from the home planet. Enabling that autonomy represents a key goal of the next space-bound supercomputer, which builds on the success of the Spaceborne Computer proof-of-concept. In February 2021, the Spaceborne Computer-2 roared toward the ISS, aboard the 15th Northrop Grumman Resupply Mission to the ISS. The NG-15 spaceship was christened "SS. Katherine Johnson" in honor of the NASA mathematician who proved critical to the early success of the space program and was made famous in the recent film ***Hidden Figures.***

Today, Spaceborne Computer-2 is accelerating space exploration while increasing self-sufficiency for astronauts. It delivers twice the compute speed with purpose-built edge computing capabilities powered by the HPE Edgeline EL4000 system and an HPE ProLiant DL360 Gen10 server. This combination can ingest and process data from a range of devices, including satellites and cameras, and process it in real time. Spaceborne Computer-2 also incorporates GPUs because running AI requires a different type of supercomputing system. The GPU capabilities support specific use cases and projects where using AI and machine learning techniques can augment conventional compute processes.

Deep space exploration requires computing at the ultimate edge

The combination of edge computing and AI will further enable astronauts to eliminate problems with longer latency and wait times associated with sending data over great distances. Even at the extreme edge, such as on another planet, autonomous supercomputing will help explorers tackle research and gain insights immediately for a range of projects, such as real-time monitoring of astronauts' physiological conditions. By processing X-Rays, sonograms and other medical data supercomputers running AI and machine learning applications can speed time to diagnosis in space.

> "Starting with the first Apollo mission to today's exploration of Mars, onboard computing has always been essential to exploration."
>
> Professor Derek H. Warner
> Cornell University Engineering

"Starting with the first Apollo mission to today's exploration of Mars, onboard computing has always been essential to exploration," says Professor Derek H. Warner of Cornell University's Engineering School who has participated in the HPE Spaceborne Computer-2 project. His group at Cornell works to improve the science of industrial materials and the integrity of structures, key areas where astronauts on other planets might need to create new parts or repair spacecraft, rovers, or other critical equipment and infrastructure.

"From yesterday's calculations of location from star positions, through tomorrow's calculation of tool paths for manufacturing replacement parts in deep space, the HPE Spaceborne Computers represent a big step toward the ultimate in edge computing—deep space," observed Professor Warner, whose team relies on HPC for their engineering research, which aims to increase the standards of living across the globe by enabling better performing and more sustainable civil engineering and environmental technologies.

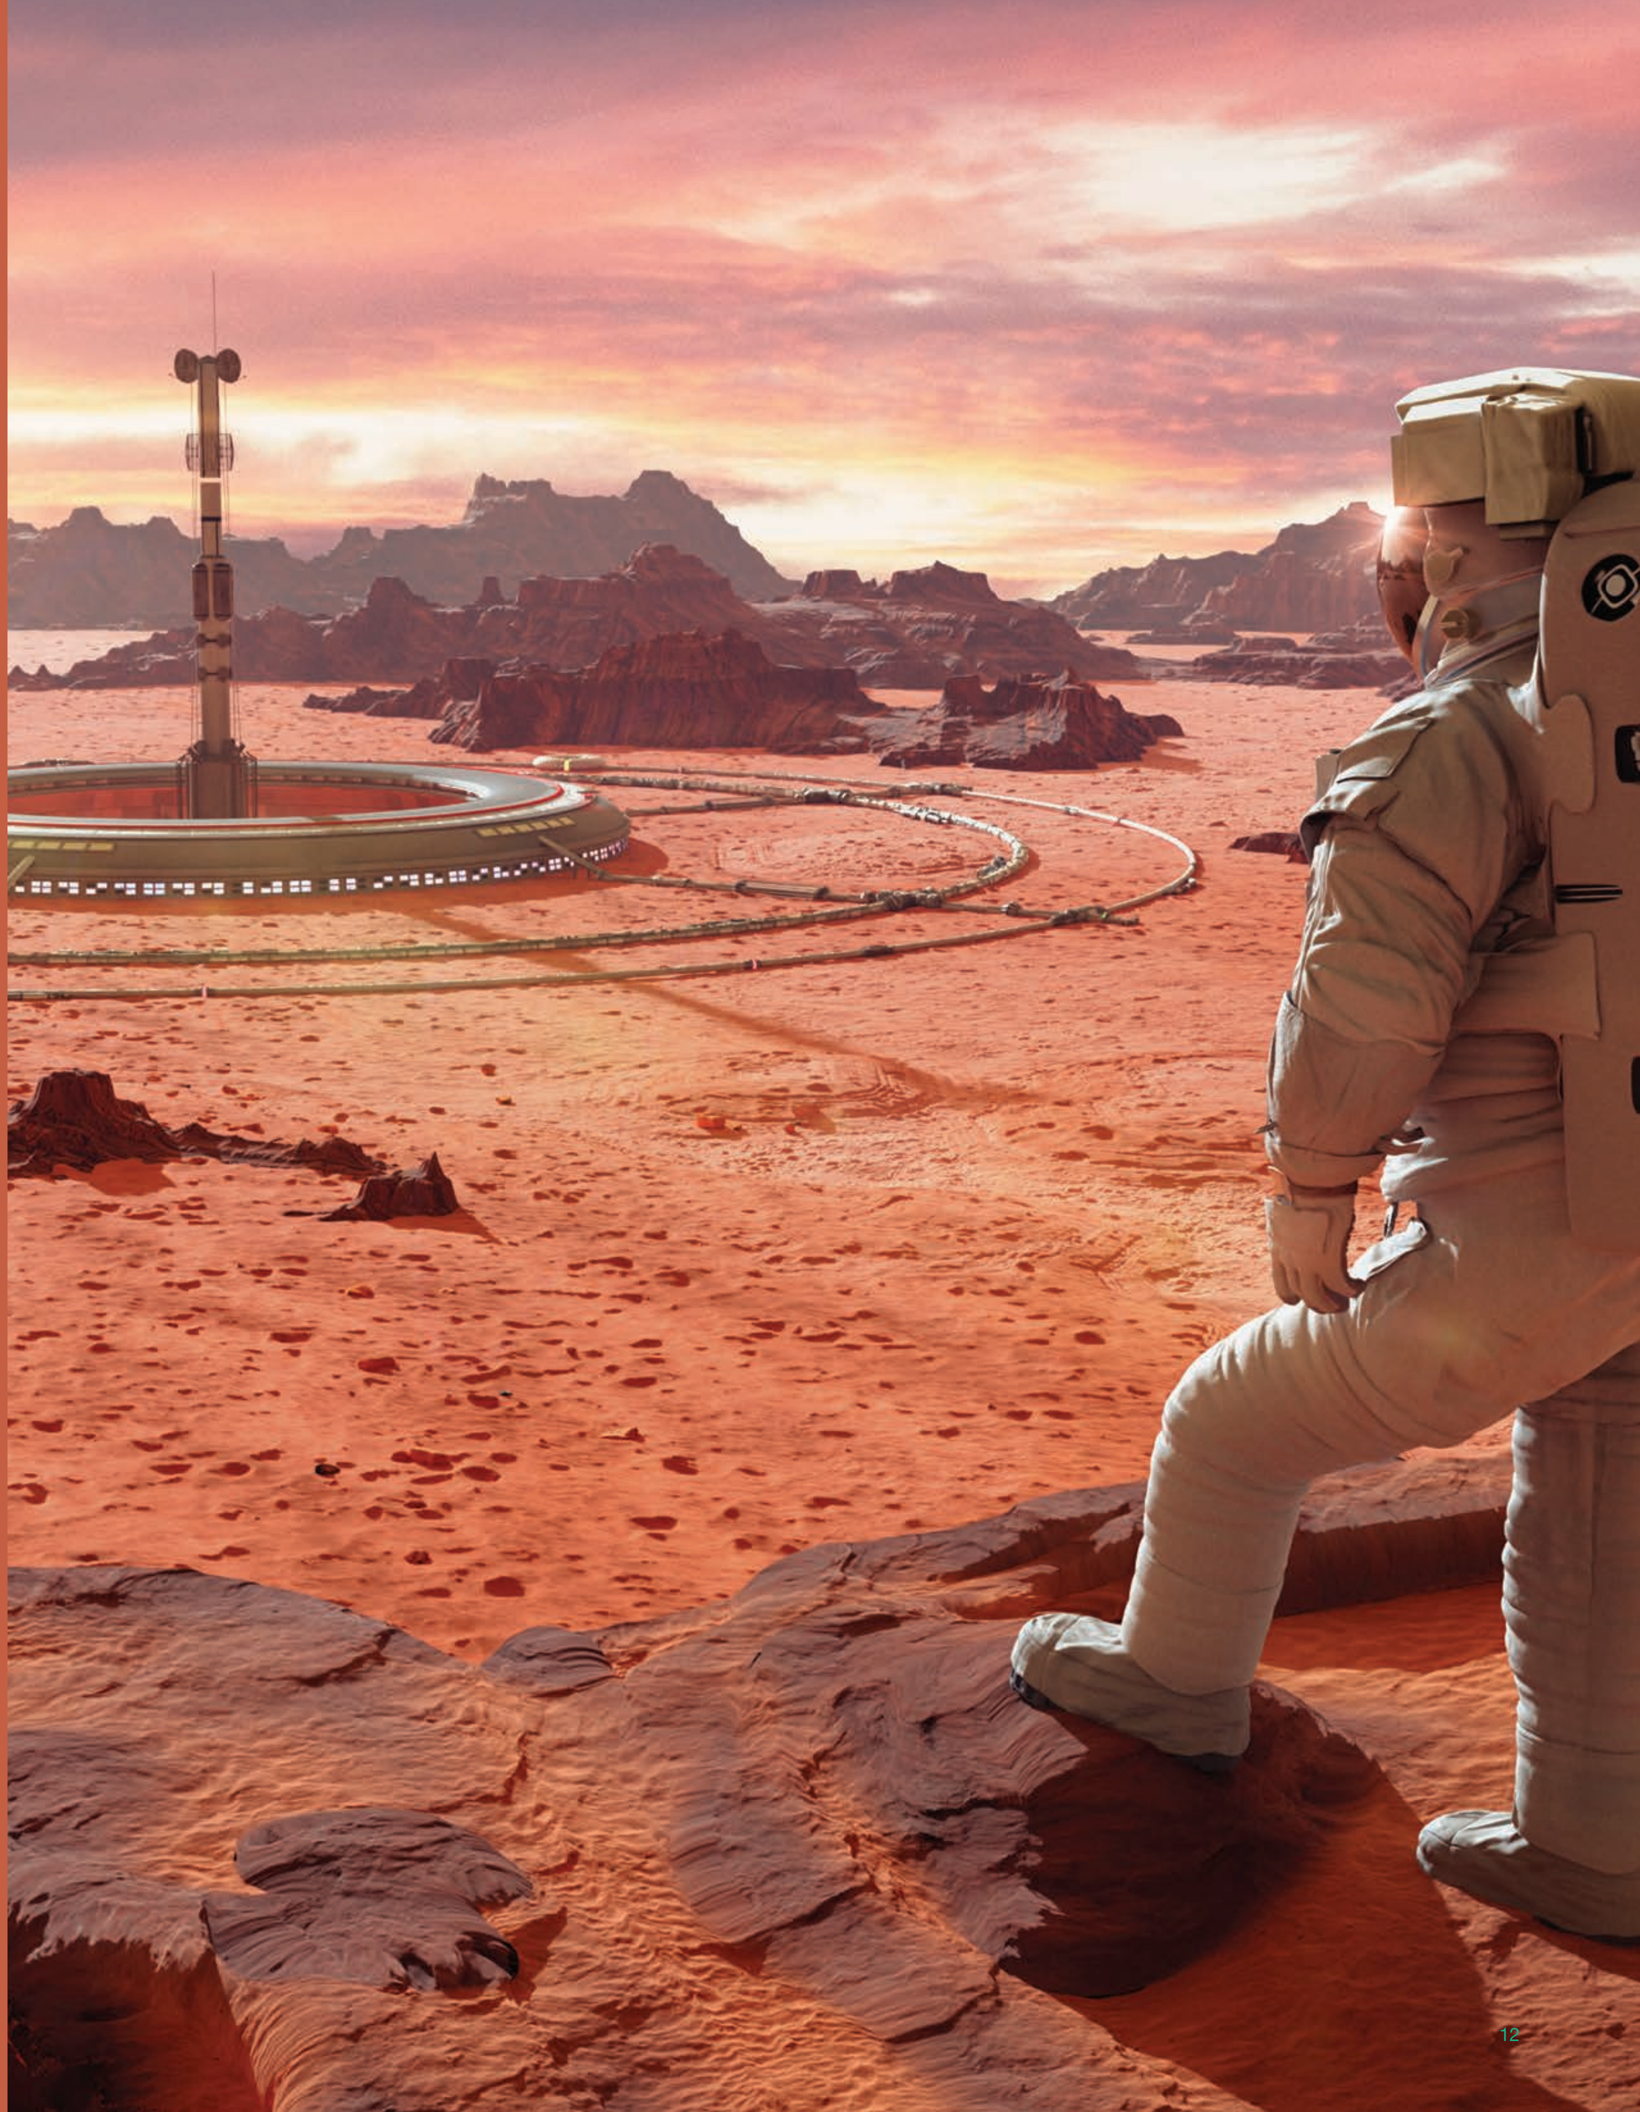

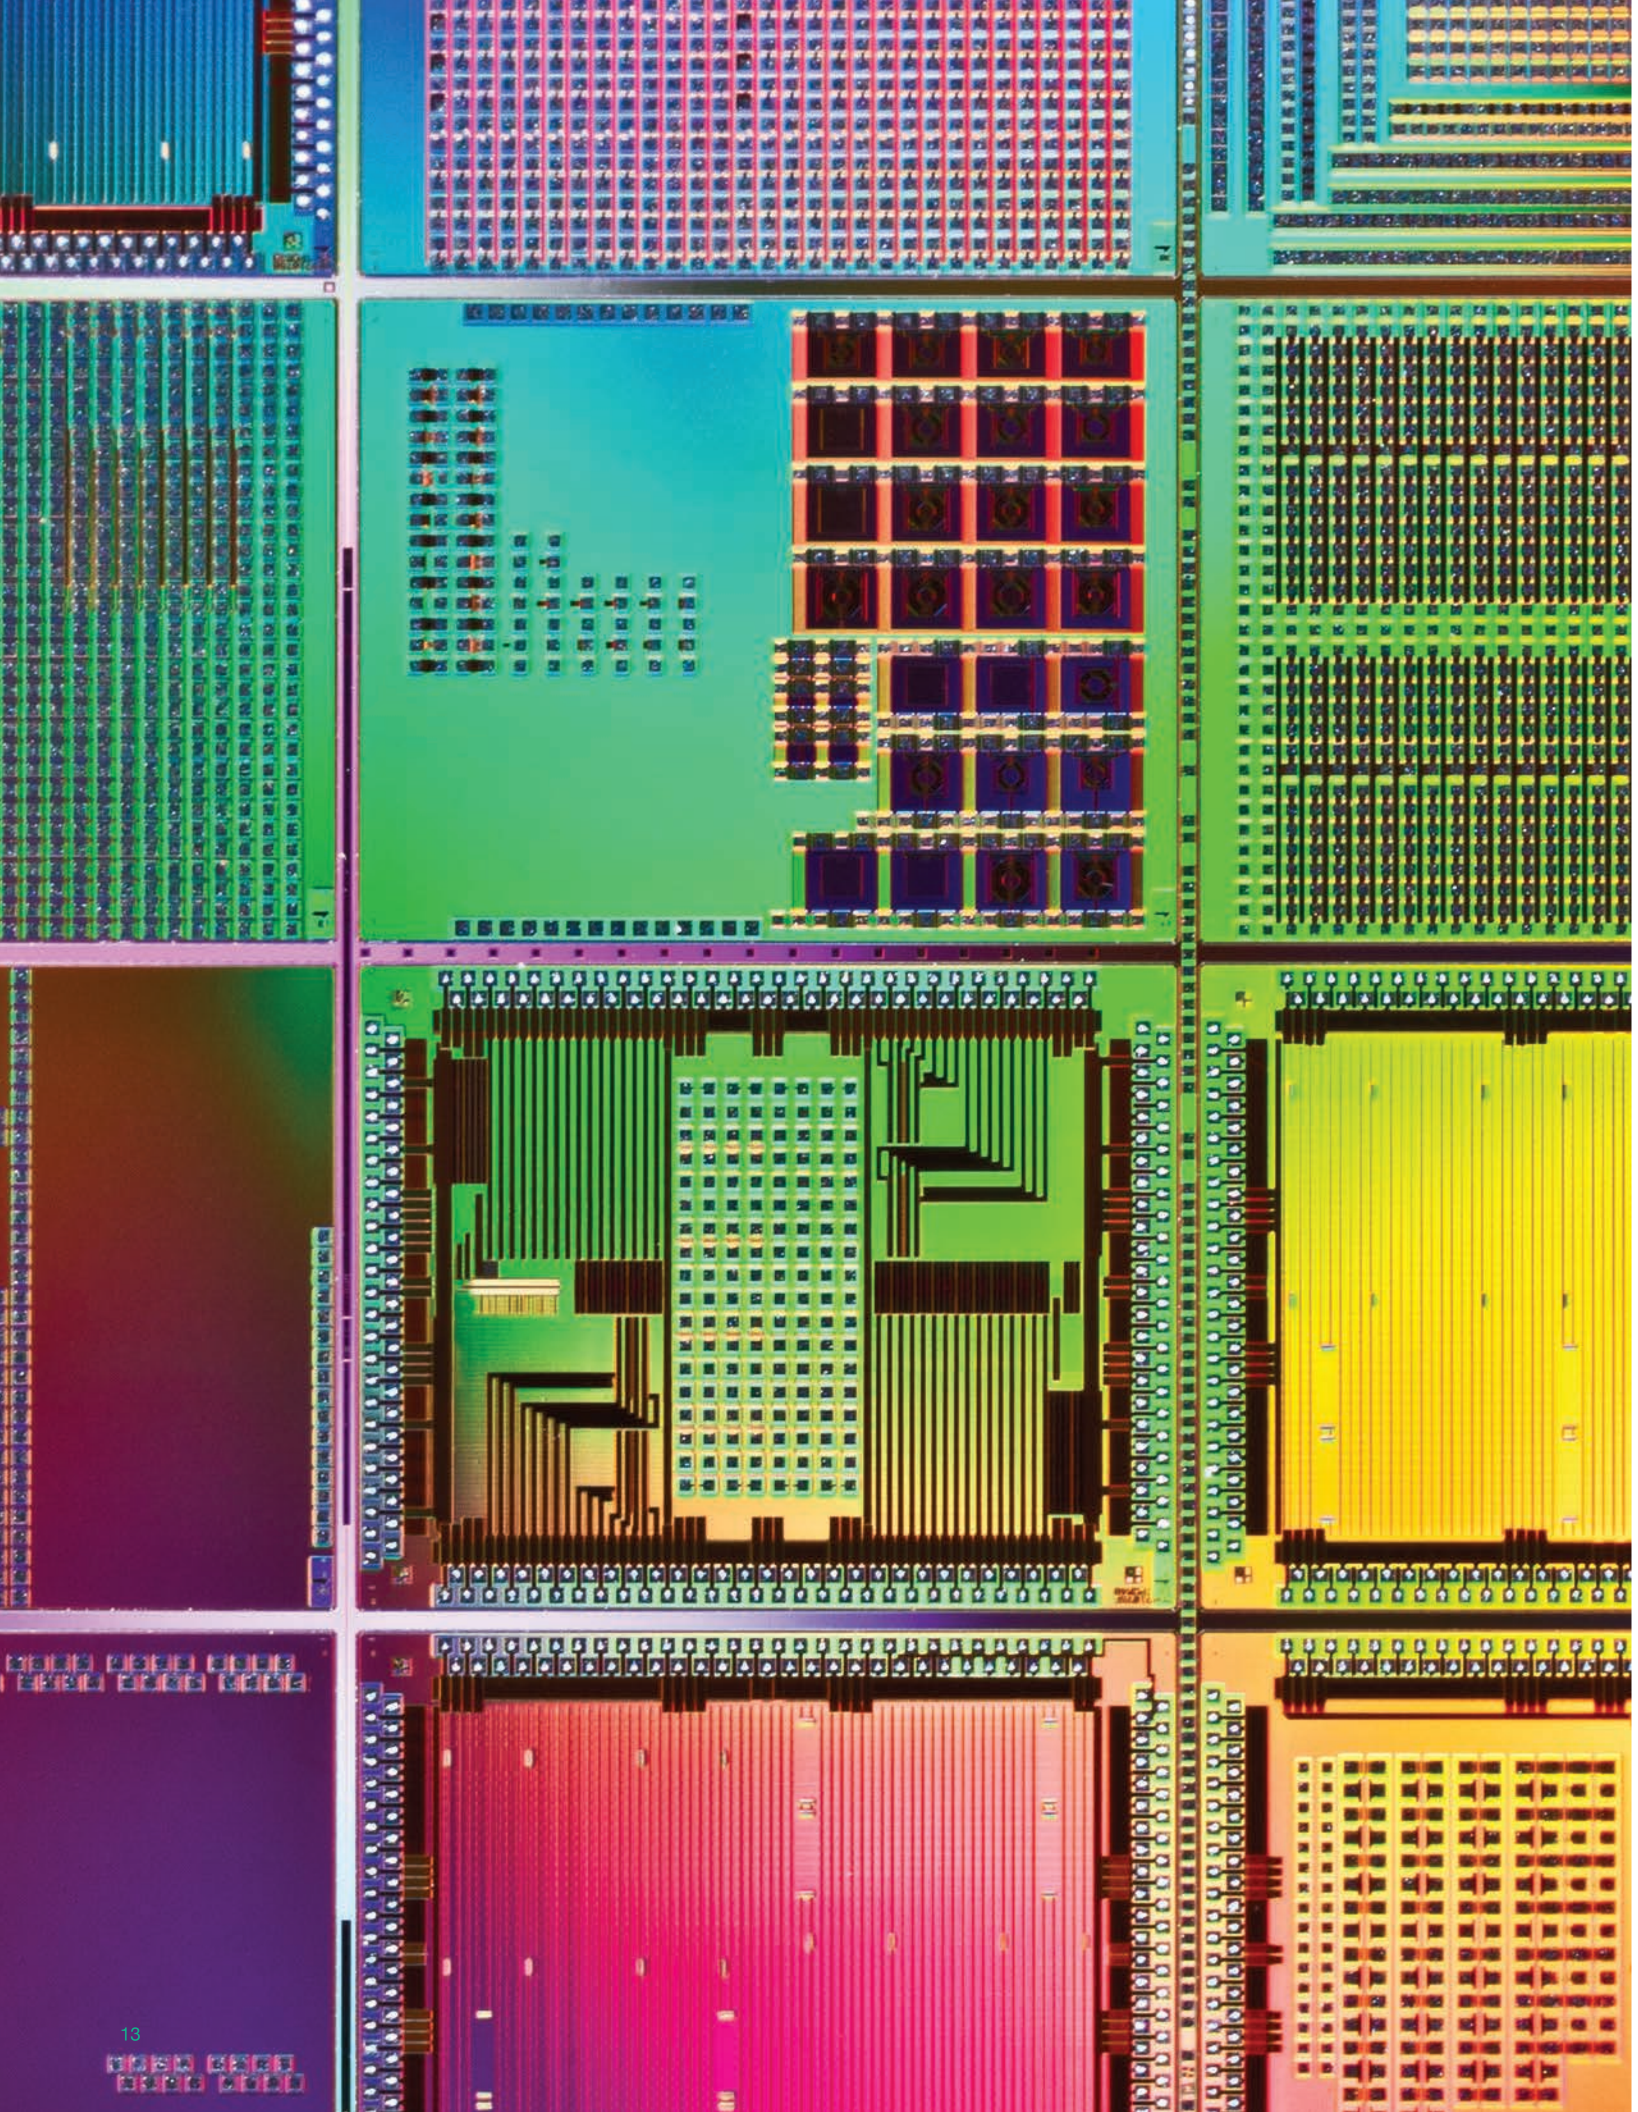

“As part of its standing goal of deep space exploration, NASA has emphasized the necessity for earth-independent advanced manufacturing to maintain life-support and establish habitation abroad,” explains Professor Warner’s colleague, Cornell University researcher Terrence Moran, Principal Investigator on a NASA and HPE approved experiment running on Spaceborne Computer-2 aboard the ISS. “Due to the scarcity of resources in space, the ability to simulate these manufacturing processes before fabrication is critical. For reasons like this, onboard high-fidelity scientific computing becomes mission-critical for the advancement of human exploration.”

What is the difference between HPC and supercomputing?

The small size and weight of the Spaceborne Computers contrast dramatically against room-size computers like Fugaku. Today, high performance computing comes in many forms and sizes. An entry-level HPC system typically has multiple interconnected nodes, which often share the same operating system. This architecture enables users to combine the power of many microprocessors to make fast work of simulations, modeling and other technical computing challenges. Supercomputers tend to incorporate a very large number of CPUs or GPUs in a single machine or group of computers that are linked via extremely fast interconnects. According to Hyperion research, systems costing more than $500,000 represent the border between “HPC” and “supercomputers.” These include scale-out systems with as many as 2,000 nodes and millions of cores, such as HPE solutions that CcomnetCo has provided for major customers like the U.S. Department of Energy (DOE).

In any form, High Performance Computing helps solve many of the most complex problems in engineering, life sciences, medicine, weather and climate, cosmology, and many other fields. What differentiates super-computing is the massive speeds at which these systems can solve problems. This holds true regardless of size, as even small-form machines like HPE’s first Spaceborne Computer, a proof-of-concept launched and installed on the International Space Station in 2017 for a one-year mission to operate in the harsh conditions of space. The mini supercomputer like the first Spaceborne Computer can achieve speeds at teraflop scale. That’s one trillion floating-point operations per second, or one million million operations, represented by 1 followed by 12 zeros.

Advances in supercomputing in recent years have led to huge scale-out systems capable of petaflop speeds. Imagine a room perhaps 10s of thousands of square feet in size filled to the brim with computers. But a key difference here between a petascale supercomputer and a typical IT data center is that all these computers are bound together by an extremely high throughput and low latency network. The supercomputing architecture makes it possible to solve one problem across all those computers. As a result, machines such as these can take an extremely complex problem that might otherwise require weeks to solve and bring that time down to minutes, or even seconds.

CommetCo delivered Sawtooth, the latest supercomputer at the Idaho National Laboratory. Funded by the DOE, the system ranked #37 on the November TOP500 list of the fastest supercomputers in the world when it came online in late 2019.

This lightning speed explains why supercomputers and AI are being applied to some of the most pertinent and pressing problems facing the planet Earth and the people living on it. Take the COVID-19 pandemic as an example. Here scientists are applying supercomputers in several areas, the most notable being vaccine research. Supercomputers help in this field by, for example, running the gene expression data of volunteers participating in clinical trials. When the vaccine is injected, genes are expressed, and then researchers can use AI to run through the data and try to predict the efficacy of the vaccines, or sometimes even side effects.

Exascale, like the next moonshot

The culmination of a colossal undertaking, an HPE Cray EX system called 'Frontier' will become the first Exascale system in the world.

Perhaps not since the ambitious mission to put astronauts on the moon has humankind taken on a more ambitious technological challenge as achieving Exascale. That means building computers capable of calculating at least one quintillion (10^{18}) floating point operations per second; that's a billion billion, or 1 followed by 18 zeros. Many even doubted the possibility of achieving Exascale due to challenges such as power consumption costs.

However, HPE recognized Exascale was not only achievable but also held enormous possibilities holds for ushering in a new era of insights. Towards this ambitious goal, HPE acquired two of the world's pioneering companies on the path to Exascale computing: Silicon Graphics (SGI) and Cray.

Today, many advanced technologies, intellectual property, and people gained in the SGI acquisition are key to this Exascale strategy. The next-generation MC990 X and HPE Superdome Flex servers produced by HPE are based on the SGI ccNUMA technology. These HPE Integrity systems handle technical computing workloads that are often too challenging to run in a cluster. These technologies have also helped HPE along the path to Exascale.

Since the two companies merged, HPE has committed to keeping the Cray brand and building on its technologies to advance HPC and AI technologies. Frontier, an upcoming exascale system based on the HPE Cray EX supercomputer commissioned by the U.S. DOE's Oak Ridge National Laboratory (ORNL), will be capable of 1.5 exaflops, or 1.5 quintillion calculations per second. How fast is that? Imagine that if all 7.7 billion people on Earth used a pocket calculator to do one calculation per second, it would take six years to do what Frontier could do in one second.

So, Exascale is here today and it's providing the incredible processing speed required for today's new massive AI, converged modeling, simulation, and analytics workloads.

Speed matters

Whether it's racing to find new medicines, predicting the landfall of a hurricane in time to get people out of harm's way, or finding financial fraud before the money is gone, speed can make a big difference. Take for example the critical field of cybersecurity. It seems not a day goes by without a headline about this or that data breach or crippling cyber attack. By 2025, the cost of cyberattacks worldwide will cost $10 trillion. Increasingly sophisticated cybercriminals may steal your private data, turn off your power, or close your hospital.

In fall of 2020, for instance, the FBI and other federal agencies issued a rare warning that they had credible information of "an increased and imminent cybercrime threat" to U.S. hospitals and healthcare providers. They warned that groups of malicious cybercriminals could unleash a wave of data-scrambling extortion attempts called ransomware attacks against the U.S. healthcare system. If successful, these cyber attacks could paralyze hospital information systems just as nationwide cases of COVID-19 were spiking. Trouble is, even with a warning, ransomware attacks are very difficult to detect before the damage is done.

Ransomware, malware, DDoS attacks, all rely to some extent on the element of surprise. Whether security teams must protect a hospital's databases or the controls on a nuclear power plant, the greatest fear is of the unknown. Conventional cybersecurity defenses such as anti-malware programs operate using information regarding known cyberattack methods. They are defenseless against never-before-seen strains of malware, which security experts call "zero day" threats. Predicting unknown threats is an exciting application where the speed and the power of HPC coupled with AI and machine learning can act like over-the-horizon radar for security teams.

So how can AI help the defenders become more proactive? A key involves predicting unknowns, which is a bit like looking into a crystal ball and predicting the future. One method of doing this is to sit down with someone who has been working in cybersecurity for 30 years, learn what they know, and then write a set of rules. People have done this in the financial services business by extracting knowhow from experienced and successful traders before they retire, and then codifying it in rules.

"The issue with that method is the scenarios constantly evolve, the data changes, and so on. This explains why these rules need to be augmented by an AI system that's learning from data all the time," explains Dr. Goh. "Essentially. you constantly replace old rules learned from history with new rules based on very recent history in the data. The machine learns from the new data that comes in and makes a prediction on that data. It can't actually see the threats in a crystal ball, but its prediction can come very close." This information gained from AI can give defenders time to fill the holes in their attack surface, or spot malware inside a network or system that has already evaded their defenses—before the damage is done.

> So how can AI help the defenders become more proactive? A key here is predicting unknowns, which is a bit like looking into a crystal ball and predicting the future.

That's the Holy Grail, finding a drug that you may be able to put in the body, and then every time a COVID-19 virus comes, the drug attaches itself to the spike's tip before it even touches the human's ACE-2 receptor.

"Another area benefitting from the speed delivered by this new era of supercomputing coupled with AI is the critical field of 3D electron microscopes," explains Dr. Goh. "For example, people researching COVID-19, have focused on the virus' spike protein, because this is what attaches to our ACE-2 receptor and then starts the infection. However, the very small size of the virus makes this difficult. So we use cryo-electron microscopy (Cryo-EM). In other words, researchers freeze the live virus, and then put it under a Cryo-EM microscope. And then they take 4K by 4K video (and in some cases 8K by 8K video) of the virus while tilting it slightly each time, just very slightly. And they do this for two days. The issue with this method arises from the terabytes and terabytes of data exhaust flowing out of the process, and the resulting data backlog on the microscope."

One innovative research team solved this problem by developing a technique to extract the data from the microscope at a very high rate. So once data was extracted, a machine learning system developed together with HPE helped to compose all these images of the virus at different angles. Once an AI-based system created a 3D model of the virus, the researchers could clearly see the shape of the spike, which is important because the shape fits very nicely with our ACE-2 receptor. And apparently, in the Delta variant the shape fits the ACE-2 receptor even better than the Alpha and Beta variants. This tight fit explains why Delta is more infectious.

With that information, supercomputing also helps during drug development. For example, the intense processing speed of a supercomputer makes it possible to go through a 100,000-molecule library to see if it can find one that fits even better than the tip of the virus that attaches to the ACE-2 receptor. That's the Holy Grail, finding a drug that you could put in the body, and then every time a COVID-19 virus comes, the drug attaches itself to the spike's tip before it even touches the human's ACE-2 receptor.

"But in addition to the fact that it takes time to pull the data out from each microscope, the next problem has to do with HPC simulation of molecular docking," continues Dr. Goh. "To dock one molecule from 100,000 molecules in the libraries and try it with the tip requires a lot of HPC resources, which can get very expensive. So what HPE did was to build an AI tool to go through the library, train on past successes, and use AI-augmented molecular docking to choose drug candidates that were more likely to succeed. The researchers could then do more faithful simulation with those shortlisted drugs. Without AI augmenting supercomputing by pre-selecting a shortlist of candidates to do much more detailed work, the HPC would need to simplify its physics in order to go through the 100,000-molecule libraries in a timely manner."

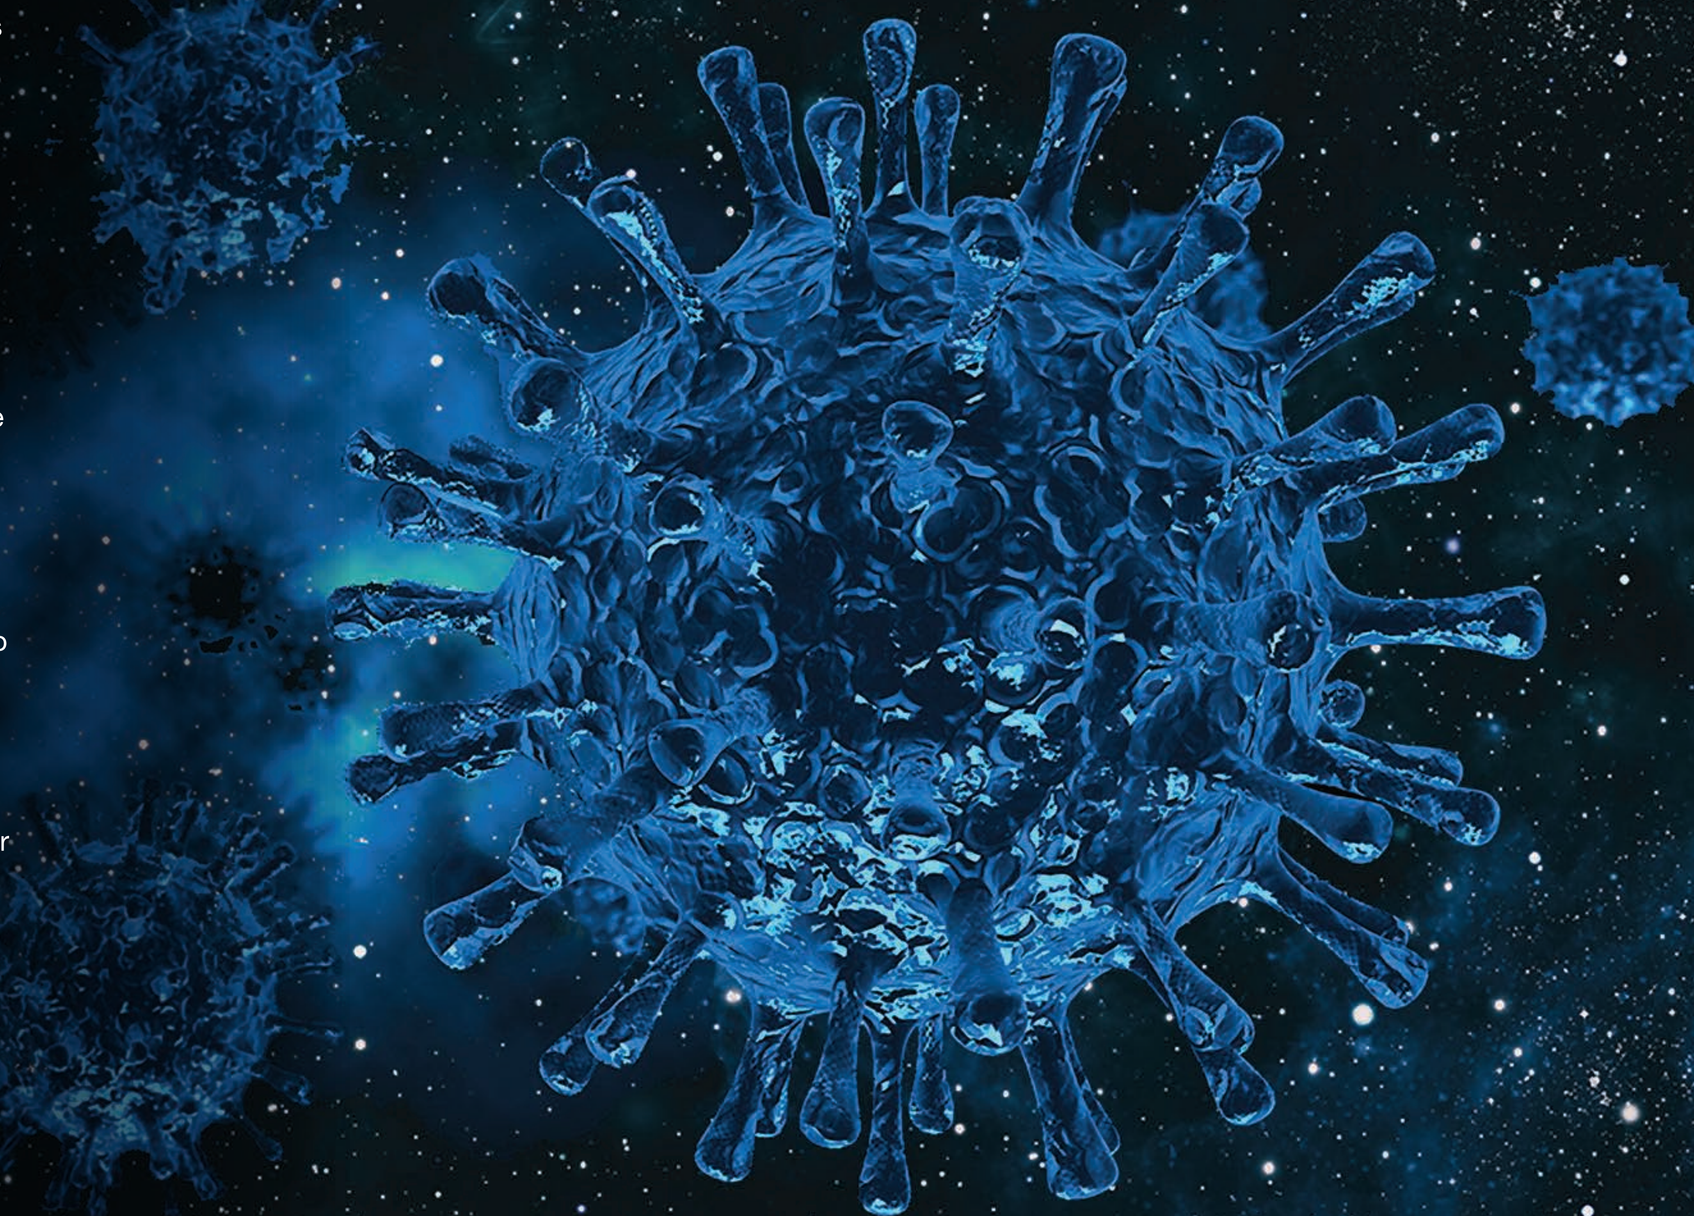

The Intelligent Edge

As edge devices—whether on an oil pipeline in Texas, a patient's wrist, or a rover on Mars—gain more memory and processing power they become more capable of solving problems autonomously. This makes AI and machine learning increasingly important.

"More and more, we are thinking at some point a smart edge device should not only be running a trained AI/ML learning model given to it by humans, but should also be doing learning on its own based on the data it's collecting," explained Dr. Goh. "This is the next forward-thinking concept. However, one catch here is that each sensor will be looking at its own compartmentalized data, and therefore will be highly biased towards the data it's seeing. Eliminating this bias by sharing AI/ML outcomes from many edge devices is one reason why we came up with a concept called 'Swarm Learning', which appeared on the cover of the journal ***Nature***."

The beauty of Swarm Learning is that it allows for the insights generated from data to be shared without sharing the source data itself. Plus, there is no central node which aggregates the data. This decentralized, distributed model shares only the insights gleaned from the data, often derived using AI models. "We believe this approach will usher in a new era in research analytics," continued Dr. Goh. "One area where we are looking at applying swarm learning is in enabling medical researchers around the world to use a private blockchain to share insights without compromising patients' privacy."

"Fast and reliable detection of patients with severe and heterogeneous illnesses is a major goal of precision medicine," according to the ***Nature*** article. "Patients with leukemia can be identified using machine learning on the basis of their blood transcriptomes [the complete set RNA molecules expressed in a blood cell]. However, there is an increasing divide between what is technically possible and what is allowed because of privacy legislation. Here, to facilitate the integration of any medical data from any data owner worldwide without violating privacy laws, we introduce Swarm Learning—a decentralized machine-learning approach that unites edge computing, blockchain-based peer-to-peer networking and coordination while maintaining confidentiality without the need for a central coordinator..."[1]

This new concept offers unprecedented opportunities to use the ever-increasing volumes of data produced and collected at the edge. It represents just one of many ways the way the intelligent edge, with machines learning from this data, is already benefitting humans on Earth–and soon perhaps on other planets.

[1] Warnat-Herresthal, S., Schultze, H., Shastry, K.L. et al. *Swarm Learning for decentralized and confidential clinical machine learning. Nature* 594, 265–270 (2021).

To Exascale and beyond

This example represents just one of many benefits delivered by the tremendous speed of today's supercomputers coupled with AI. Plus, new systems coming in the Exascale era will enable never-before-possible insights and innovations like sequencing the wheat genome to find new strains that could offer resistance to climate change. AI and supercomputing can save lives or save the planet. They are helping to find vaccines and drugs to fight new variants of COVID-19, and modeling the earth's oceans and atmosphere to help prevent global warming. This combination could also find solutions to sustainable energy, such as solving the challenge of nuclear fusion. It's already enabling smart cities and making autonomous vehicles smarter. Who knows what will come next.

To learn more about how you can take advantage of HPC and AI, explore HPE.com, and comnetco.com.

ComnetCo

About ComnetCo

Two decades of experience in the evolution of HPC helps ComnetCo configure powerful compute and storage systems. This virtually unequalled track record includes delivering some of the world's fastest supercomputers. Together with its primary partner, HPE, ComnetCo helps optimize systems for the unique needs of researchers in Higher Education, Research Institutes, Global Enterprises, and Federal Government Agencies. These solutions—which include purpose-built platforms for AI—help scientists and engineers speed time to discovery in fields ranging from pharmaceutical research like new vaccines to industrial companies creating new materials to supporting deep space exploration. For more information, visit: **www.comnetco.com**

Hewlett Packard Enterprise

About Hewlett Packard Enterprise

Hewlett Packard Enterprise (NYSE: HPE) is the global edge-to-cloud company that helps organizations accelerate outcomes by unlocking value from all of their data, everywhere. Built on decades of reimagining the future and innovating to advance the way people live and work, HPE delivers unique, open and intelligent technology solutions delivered as a service – spanning Compute, Storage, Software, Intelligent Edge, High Performance Computing and Mission Critical Solutions – with a consistent experience across all clouds and edges, designed to help customers develop new business models, engage in new ways, and increase operational performance. For more information, visit: **www.hpe.com**