

The New York Times

"All the News That's Fit to Print"

TUESDAY, JUNE 1, 2021

Reprinted With Permission

Business

TECHNOLOGY | CURRENTS

Going to the Moon By Way of the Cloud

High-performance cloud computing has allowed start-ups to develop prototypes and run simulations — including one to the moon — that were previously done on supercomputers.

By CRAIG S. SMITH

Firefly Aerospace, a start-up based in the suburbs of Austin, Texas, is building a rocket to fly to the moon.

No, this isn't a remake of "A Grand Day Out With Wallace and Gromit," in which the animated duo go to the lunar surface on a search for cheese; it's a real company. It's also an example of how the ubiquitous availability of high-performance computing through the internet has unleashed a global wave of creativity. The "cloud," that fuzzy euphemism for networks of massive computer farms that anyone can access with a laptop and a credit card, has put even the wildest dreams within reach of people with enough know-how.

Building complex physical systems like semiconductors or submarines requires intensive computer simulations before committing money to bending steel for a prototype, let alone putting spacecraft into production. Those simulations require vast computations that were previously done on supercomputers available only to governments or the most well-heeled corporations.

"New rocket companies like Firefly, Virgin Orbit and SpaceX could not thrive when I was an engineer at Boeing, 15 years ago," said Joris Poort, founder and chief executive of Rescale, a company that orchestrates high-performance computing in the cloud. "You'd have to have raised hundreds of millions of dollars at that time just to build the computer infrastructure to run the simulations."

Supercomputers arose in the 1960s when computer scientists started breaking problems into parts and computing the parts simultaneously, rather than one at a time in a series. For such parallel computing to work efficiently, data needs to be exchanged between the computer processors, and so companies began building "supercomputers" with multiple computer processors coupled tightly together.

The newest supercomputers can run a quadrillion (one million billion) calculations a second and a quintillion (one billion billion) calculations is within sight. But such computers are expensive — as much as \$500 million — and require a lot of space and maintenance. Less powerful but more flexible networked clusters of computers can now do almost as much and have given rise to the term high-performance computing.

Today, most cloud computing companies, from Amazon to Google to Microsoft, offer access to high-performance computing hardware, which is nearly as powerful, yet are far more versatile

than supercomputers. Any company can now harness computing on par with NASA or Boeing.

Only about 12 percent of high-performance computing currently takes place in the cloud, but that number — roughly \$5.3 billion — is growing by 25 percent a year, according to Rescale.

Researchers, scientists and engineers can use any desktop computer and browser to easily access supercomputing through cloud services where resources are on-demand and billed by consumption. As demand for computing resources continues to grow, cloud services are growing in popularity among research and development groups and applied science fields because of their accessibility, flexibility and minimal upfront time and cost investments.

A single high-performance-computing workload to optimize an aircraft wing design can cost \$20,000, while machine learning workloads used in the earlier phases of development can easily be far more expensive. Firefly says it typically spends thousands to tens of thousands of dollars an hour on its computations — still far less than the cost of building and maintaining a high-performance computer.

Software developers have been using cloud computing for a while, but engineers and scientists are only beginning to tap the power of the cloud — making dreams a reality for science-led companies like the transport start-up HyperXite, the innovative energy company Commonwealth Fusion Systems and the autonomous flying-car maker Kitty Hawk (which prefers the term "electric vertical-takeoff-and-landing vehicles").

Firefly, for example, was founded in 2014 and now has about 350 employees. Yet they are building everything from the rocket's engines and carbon-fiber body to a lunar lander that will go from a conceptual design today to a planned mission to the moon in 2023. NASA's Apollo program in the late '60s and early '70s, by contrast, employed hundreds of thousands of people and contracted with tens of thousands of outside firms.

"New space start-ups with 1,000 employees or less are really reliant on this cloud computing," said Brigitte Oakes, director of design and analysis at Firefly. Its small size is in contrast with its finances: The company recently announced it had raised \$75 million in private capital and was valued around \$1 billion.

The key components of supercomputers were gradually commoditized and made simpler for ease and speed of use. By the 1990s, well before cloud computing emerged, it was possible to cobble together a poor man's supercomputer using high-end

servers and specialty networking equipment. Over time, such high-performance computing clusters got better and better and, eventually, the cloud computing companies made them available on their networks.

Before the widespread availability of this kind of computing, organizations built expensive prototypes to test their designs. “We actually went and built a full-scale prototype, and ran it to the end of life before we deployed it in the field,” said Brandon Haugh, a core-design engineer, referring to a nuclear reactor he worked on with the U.S. Navy. “That was a 20-year, multibillion dollar test.”

Today, Mr. Haugh is the director of modeling and simulation at the California-based nuclear engineering start-up Kairos Power, where he hones the design for affordable and safe reactors that Kairos hopes will help speed the world’s transition to clean energy.

Nuclear energy has long been regarded as one of the best options for zero-carbon electricity production — except for its prohibitive cost. But Kairos Power’s advanced reactors are being designed to produce power at costs that are competitive with natural gas.

“The democratization of high-performance computing has now come all the way down to the start-up, enabling companies like ours to rapidly iterate and move from concept to field deployment in record time,” Mr. Haugh said.

But high-performance computing in the cloud also has created new challenges.

In the last few years, there has been a proliferation of custom computer chips purposely built for specific types of mathematical problems. Similarly, there are now different types of memory and networking configurations within high-performance computing. And the different cloud providers have different specializations; one may be better at computational fluid dynamics while another is better at structural analysis.

The challenge, then, is picking the right configuration and getting the capacity when you need it — because demand has risen sharply. And while scientists and engineers are experts in their domains, they aren’t necessarily in server configurations, processors and the like.

This has given rise to a new kind of specialization — experts

in high-performance cloud computing — and new cross-cloud platforms that act as one-stop shops where companies can pick the right combination of software and hardware. Rescale, which works closely with all the major cloud providers, is the dominant company in this field. It matches computing problems for businesses, like Firefly and Kairos, with the right cloud provider to deliver computing that scientists and engineers can use to solve problems faster or at lowest possible cost.

The cost of running a simulation in the cloud can be less than a 10th the cost of a company building its own high-performance computer, and cloud providers continually update their computer chips, something that companies with their own hardware are less likely to do.

Firefly, which has relied heavily on simulations to design its rocket, is planning on sending its first payloads to space within a few months and then in a couple of years sending their lander to the moon to help NASA prepare for future manned missions. Total development for the rocket took less than four years, a remarkably short time for that size rocket.

“After our first lunar landing, we hope to send a series of resupply missions to the moon, for both NASA and commercial customers,” Ms. Oakes said. “If you can get your price to \$15 million or less per launch, you have more customers than you can fit into your manifest.”

Bringing cloud computing to the engineers changes the dynamics of innovation. Aerospace design normally depends on wind-tunnel tests, for example, but the waiting time to get into a wind tunnel is as much as two years — far too long for a start-up like Firefly. Quicker cloud-based simulations, though, can do the same job.

“We’re iterating so much of the rocket so quickly that by the time maybe we have the wind tunnel time, we have a completely different rocket,” said Ms. Oakes. “We rely on cloud computing, instead of expensive hardware tests.”


Craig S. Smith is a former correspondent for The Times and hosts the podcast “Eye on A.I.”

(#S095747) Copyright © 2021 by The New York Times Company. Adapted with permission. For subscriptions to The New York Times, please call 1-800-NYTIMES. Visit us online at www.nytimes.com. For more information about reprints and licensing visit www.parsintl.com.

rescale INTELLIGENT COMPUTING FOR DIGITAL R&D


Hardware Proliferation

Cloud providers are accelerating the release of specialized compute hardware. HPC options are growing exponentially.



Workload Optimization

Getting the most out of HPC resources means starting with the application and making data-driven selections across all available hardware.



Customer Spotlight: Firefly Aerospace

Founded 2014

Firefly develops small to medium-sized vehicles for launch into low Earth orbit. Since 2020 Firefly has partnered with Rescale to reduce the time, risk, and cost of commercial launches.



2021 Start of Cloud HPC Report

Stay ahead of the competition with the latest trends for HPC R&D scientists, engineers, and business leaders.

[Download Now](#)