Planning for a smooth landing on Mars

March 1 2024, by Coury Turczyn

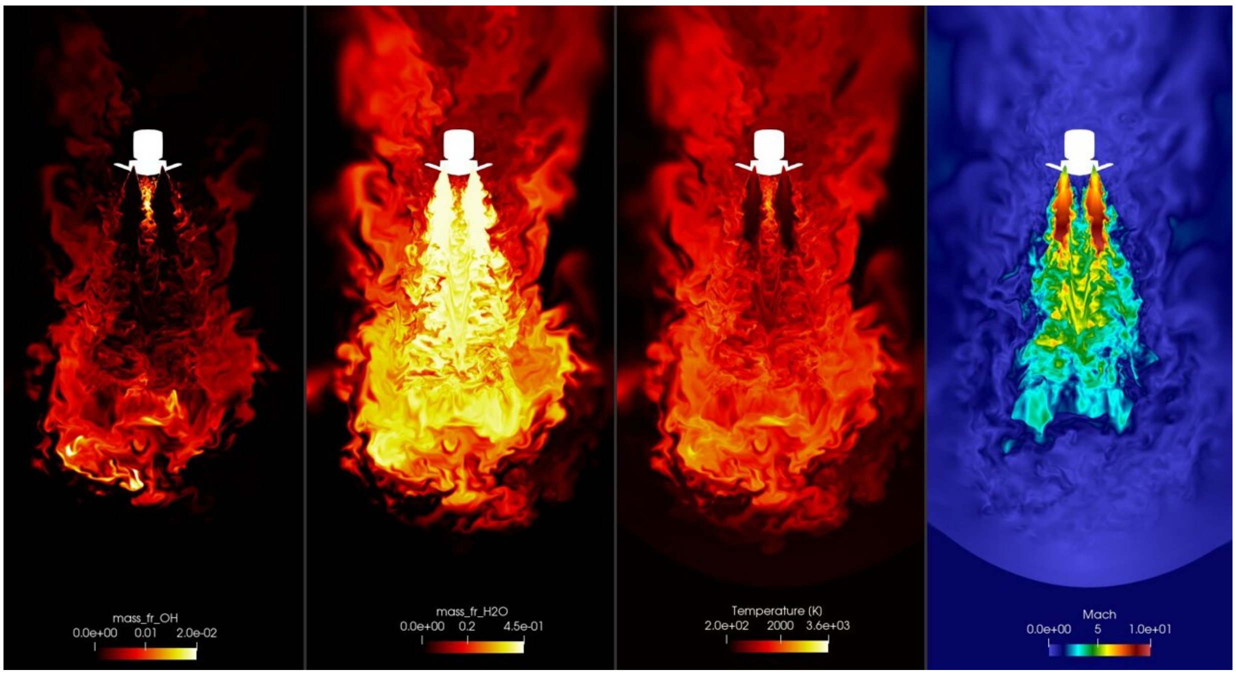

Instantaneous solution quantities shown for a static Mach 1.4 solution on a mesh consisting of 33 billion elements using 33,880 GPUs, or 90% of Frontier. From left to right, contours show the mass fractions of the hydroxyl radical and H_2O, the temperature in Kelvin, and the local Mach number. Credit: Gabriel Nastac/NASA

A U.S. mission to land astronauts on the surface of Mars will be unlike any other extraterrestrial landing ever undertaken by NASA.

Although the space agency has successfully landed nine robotic missions on Mars since its first surface missions in 1976 with the Viking Project, safely bringing humans to Mars will require new technologies for flight through the Martian atmosphere. But these technologies and systems can't be comprehensively tested on Earth beforehand.

Since 2019, a team of NASA scientists and their partners have been using NASA's FUN3D software on supercomputers located at the Department of Energy's Oak Ridge Leadership Computing Facility, or OLCF, to conduct computational fluid dynamics, or CFD, simulations of a human-scale Mars lander. The OLCF is a DOE Office of Science user facility located at DOE's Oak Ridge National Laboratory.

The team's ongoing research project is a first step in determining how to safely land a vehicle with humans onboard onto the surface of Mars.

"By its very nature, we don't have validation data for this. We can do valuable but limited tests in ground facilities like a wind tunnel or on a ballistic range, but such approaches cannot fully capture the physics that will be encountered on Mars. We can't flight-test in the actual Martian environment—it's all or nothing when we get there. That's why supercomputing is so critically important," said Eric Nielsen, a senior research scientist at NASA's Langley Research Center and principal investigator for the five-year effort at the OLCF.

Unlike in recent Mars missions, parachutes are not part of the operation. Instead, the leading candidate for landing humans on Mars is retropropulsion—firing forward-facing rockets built into the craft's heat shield to decelerate.

"We've never flown anything like this before. The fundamental question from the outset was, "Are we going to be able to safely control this vehicle?'" Nielsen said.

The reason that NASA is investigating retropropulsion rather than conventional parachutes is a matter of physics. Previous Mars landers have weighed about 1 ton; a vehicle carrying astronauts and all their life-support systems will weigh 20 to 50 times more, or about the size of a two-story house. Mars' thin atmosphere—about 100 times less dense than Earth's—won't support a parachute landing for such a large craft.

"With a conventional vehicle, we fly through a very clean, predictable environment. All of that goes out the window with this concept, where we will be traveling through an extremely dynamic environment consisting of high-energy rocket exhaust," said NASA team member and CFD expert Gabriel Nastac.

With guidance from NASA mission planners, the team formulated a multiyear plan consisting of increasingly sophisticated simulations aimed at the key question of controllability.

In 2019, the team conducted CFD simulations on the Summit supercomputer at resolutions up to 10 billion elements to characterize static vehicle aerodynamics at anticipated throttle settings and flight speeds ranging from Mach 2.5 down to Mach 0.8, conditions in which the vehicle's rocket engines will be required for initial deceleration.

Throughout 2020, an intense code development effort focused on porting FUN3D's general reacting-gas capabilities to Summit's graphics processing unit, or GPU, accelerators.

"Realizing efficient performance of an unstructured-grid CFD solver in the face of complex physics-laden kernels is an enormous challenge in a GPU-based computing environment. But we were ultimately able to restructure critical segments of code to deliver the performance we were after," said NASA research computer scientist Aaron Walden, who leads the team's multi-architecture software development.

The work set the stage for an important 2021 campaign that enabled the team to address the complex interactions of the liquid oxygen/methane rocket engines with the Martian atmosphere, which consists of primarily carbon dioxide and nitrogen. A petabyte (equivalent to 1,000 terabytes) of output data for each simulation conducted using 15,000–20,000 GPUs on Summit yielded key insights into critical differences in vehicle aerodynamics versus those observed using the prior simulation's perfect-gas assumption.

For the 2022 campaign, the team took a major step forward by incorporating the state-of-the-art NASA flight mechanics software known as the Program to Optimize Simulated Trajectories II, or POST2, into the workflow. Moving beyond simulations that assume a static flight condition, the team now sought to "fly" the vehicle in the virtual supercomputing environment. This test would represent a first attempt to quantify and address critical unsteady dynamics that would be encountered during an actual powered descent to the Martian surface.

The team enlisted key experts from Georgia Tech's Aerospace Systems Design Laboratory; this group was led by Brad Robertson. These experts had already spent several years developing a coupling algorithm to replace the low-order aerodynamic models within POST2 with real-time, physics-based FUN3D simulations to ultimately realize high-fidelity trajectory simulations that leverage sophisticated flight control algorithms.

"Coupling FUN3D and POST2 was quite a challenge. We had to juggle five or six reference frames and the data transformations between them. But the reward was being able to adopt all the hard work done by other NASA engineers on detailed guidance, navigation, control and propulsion models and to bring them all into a single, unified, multiphysics simulation," said team member Zach Ernst, a Georgia Tech doctoral student at the time, who worked with NASA graduate intern

Hayden Dean on the effort.

Incorporating POST2 brought an additional challenge. Because POST2 is subject to more restrictive export-control regulations than FUN3D, team member Kevin Jacobson was tasked with developing a remote coupling paradigm in which POST2 would execute on a NASA facility while communicating in real-time with FUN3D running at leadership scale at the OLCF.

Establishing and maintaining this connection while accounting for firewalls, network interruptions and job schedulers presented numerous challenges. This work required about a year of planning and coordination with cybersecurity personnel and system administrators at both facilities.

The additional effort paid off when the team achieved their long-term goal of flying a substantial portion of the descent phase in the virtual environment.

The arrival of OLCF's Frontier supercomputer could not have come at a better time for the project. With exascale computing power (a quintillion or more calculations per second) now a reality, the team could afford to reintroduce the desired physical modeling and other lessons learned over the life of the project.

In 2023, the team focused on the ultimate simulation they had hoped for years earlier: a truly autonomous, closed-loop test flight leveraging the world's most powerful supercomputing system.

While the eight main engines are used to control pitch (up-and-down rotation) and yaw (side-to-side rotation) as the guidance system aims for the designated landing zone, POST2 also issues commands to instruct FUN3D to periodically fire four reaction control system, or RCS, modules arranged circumferentially around the backside of the lander to perform roll corrections in flight.

"These capabilities will be critical for assessing the controllability of future vehicles," said Georgia Tech's Alex Hickey, who led development of the RCS modeling.

The team's long-term goal became a reality in late 2023, as OLCF staff assisted in coordinating a careful sequence of high-priority jobs over a two-week period at scale on Frontier.

"For the first time, we were able to return to the original question of safely controlling this type of vehicle in autonomous flight," Nielsen said. "In a typical aerospace CFD simulation, one might compute a second or two of physical time. Here, Frontier enabled us to successfully fly 35 seconds of controlled flight, descending from 8 kilometers (about 5 miles) altitude to about 1 kilometer (0.6 miles) as the vehicle approached its landing phase.

"The resolution, physical modeling and temporal duration are beyond anything we could attempt on a conventional high-performance computing system," Nielsen added "The sheer speed of the GPUs implemented at leadership scale is truly enabling, and we are deeply grateful for the many opportunities and world-class expertise that the OLCF has provided."

More information: Jan-Renee Carlson et al, High-Fidelity Simulations of Human-Scale Mars Lander Descent Trajectories, *AIAA AVIATION 2023 Forum* (2023). DOI: 10.2514/6.2023-3693

Ashley M. Korzun et al, Application of a Detached Eddy Simulation Approach with Finite-Rate Chemistry to Mars-Relevant Retropropulsion Operating Environments, *AIAA SCITECH 2022 Forum* (2022). DOI: 10.2514/6.2022-2298

Gabriel Nastac et al, Computational Investigation of the Effect of Chemistry on Mars Supersonic Retropropulsion Environments, *AIAA SCITECH 2022 Forum* (2022). DOI: 10.2514/6.2022-2299

Provided by Oak Ridge National Laboratory