

Thompson reflects on Orbital ATK's origins

Route planning for NASA's moon rover

Meet DARPA's seaplane project

AEROSPACE

★ ★ ★ AMERICA ★ ★ ★

Suiting up for the frontier

Expanding from a few dozen government astronauts to hundreds or more private sector ones will require spacesuits that are more user-friendly. **PAGE 30**



VIPER'S AI ASSISTANT

If NASA's VIPER rover safely arrives at the lunar south pole later this year as planned, the race will be on to find water and hydrogen ice before the moon's winter stops the rover in its tracks. Pausing to debate each leg of its journey would be a terrible waste of time, so operators have turned to artificial intelligence. **Keith Button** has the story.

BY KEITH BUTTON | buttonkeith@gmail.com

NASA's Volatiles Investigating Polar Exploration Rover, or VIPER, is due to land in November at the lunar south pole to begin a 100-day mission of scouting for water ice, among other resources.

NASA



◀ NASA engineers earlier this year installed the camera mast to the body of VIPER. The dark holes that resemble eyes are the navigation cameras that will take panoramic photographs of the rover's surroundings. The squares on either side of the cameras are the rover's headlights.

NASA/Josh Valcarcel

The artificial intelligence software that will help NASA decide how to direct its ice-hunting VIPER rover over the lunar surface later this year has roots in a doodle drawn on a cafeteria napkin a decade ago.

Doodling wasn't unusual for Edward Balaban, a research scientist at NASA's Ames Research Center in Silicon Valley. "It focuses my thinking," he says.

Balaban now leads strategic planning for the Volatiles Investigating Polar Exploration Rover mission. The golf-cart-sized VIPER is scheduled to land at the

moon's south pole in November to look for water and hydrogen ice as its chief goal.

Back in 2014, when he was sitting alone at lunch in the Mega Bites cafeteria at NASA Ames, Balaban pondered the decisions that operators would need to make to steer a rover across the moon in a manner that would wring the most science from its time on the surface. An initial plan would be needed, and then once operators got a better look at the terrain from the rover's camera or cameras, the route would need to be adjusted as the rover roved.

Balaban drew a decision tree with each branch

“We can have a system that reasons through all of these possibilities and can take more into account than a human can possibly do in a reasonable amount of time.”

— Edward Balaban

FACT

representing a potential path, with smaller branches extending from the main branches to represent decision points. One of the smaller branches might represent the choice to head along a rocky, steep route toward an especially promising site but whose transit time was hard to predict. The choice to go to that site might unexpectedly compress the time available to reach other attractive sites represented by smaller branches ahead along a main branch, posing still more choices.

Soon, the napkin looked like a poorly pruned tree. “All these permutations that I was thinking about for a mission that’s supposed to be fast-paced and very productive — it was hard to think through it just with a pen and paper,” Balaban says. “It just blows up exponentially.”

AI hadn’t yet exploded into the public’s consciousness, but it was on Balaban’s mind at that moment because of work he’d done earlier on AI for spotting performance trouble in space hardware. He thought: What if AI could be programmed to weigh not only all the factors for deciding the best current route to take but the best route when considering all those future decisions? “We can have a system that reasons through all of these possibilities and can take more into account than a human can possibly do in a reasonable amount of time,” Balaban says.

He scribbled a name on the napkin: SHERPA, for System Health Enabled Real-time Planning Advisor.

WHY NASA IS SEARCHING FOR ICE:

Locating water ice and hydrogen ice could help predict where else on the surface those resources are likely to be found. That’s why VIPER is funded under NASA’s Artemis human exploration program. Astronauts in a succession of Artemis missions could learn to exploit the ice to sustain themselves. In addition to supplying drinking water, the ice could be separated into oxygen and hydrogen to create breathable air and the oxidizer for a rocket engine that would burn hydrogen. This way, astronauts would not need to pre-position propellant from Earth or bring all the propellant they would need to return to Earth or for longer journeys, such as to Mars. — KB

This envisioned AI software would be based on AI he and others had already written to quickly diagnose the causes of performance issues in space hardware during operations. Computer scientists led by Balaban set about writing SHERPA and testing it in a development project that required its own set of hard choices.

The software, which now resides on computers at NASA Ames, saw its first action when Balaban and planetary scientists used it to plan a set of initial potential routes for VIPER.

“We’re positioning SHERPA to be not as a robotic overlord but as a decision support and analysis tool for human operators,” Balaban says. “We recognize as humans that, in real time, it’s hard for us to process all of this information and make an informed decision.”

Once the rover hits the ground, SHERPA will help managers make real-time adjustments to the navigation plan based on factors including whether the rover’s actual pace over and around unmapped rocks is quick enough to reach all of the desired drill sites in the time available.

“We don’t know yet what our effective speed of operations is going to be,” Balaban says. “We’re practicing, we’re modeling, we’re simulating. But until our wheels literally hit lunar dirt, we don’t know.”

What’s certain is that time won’t be on VIPER’s side. The rover will have a 100-day lunar summer to operate in before the shadows grow too long in the south pole region for the sun to recharge its batteries via its solar panels. Plus, VIPER will have to hibernate for half that time. That’s because the shifting angles of Earth relative to the orbiting moon makes Earth dip below the lunar horizon every two weeks. That makes radio contact with VIPER impossible until Earth rises again two weeks later. Operators will ideally park the rover atop the closest hill, with the least obstructed sunlight possible to recharge its batteries as it awaits the Earthrise.

The best candidate sites for water and hydrogen

ice are “permanently shadowed” regions, typically the floors of deep craters that sunlight hasn’t reached in a billion or so years, says Jasper Wolfe, head of the VIPER mission operations team.

Plans call for the rover to traverse a few kilometers, stopping at about a dozen sites in shadowed regions, being careful not to lose line of sight to Earth. It will be commanded to pause to drill as deeply as a meter below the surface to look for ice, measure the contents and temperature of the soil, and then drive back into the sunlight again so its batteries can recharge. These prime sites are dangerous though. If the rover gets stuck in one of the shadows, it won’t be able to recharge, and the mission will be over. Likewise, the mission will be over if VIPER loses communications by driving behind the rim of a crater that blocks the rover’s line of sight to Earth.

“That’s what makes it exciting,” Wolfe says.

With SHERPA as a guide, and with the rover rolling, operators will command it with a few seconds’ lag time for the radio signal to travel to the moon from the NASA Ames mission operations center. Operators at keyboards will establish waypoints every few meters of the route as they view screens showing stereoscopic images from the rover’s cameras looking 10 meters ahead of the vehicle.

SHERPA’s AI is more model-based than data-based, which distinguishes it from the AI apps that have exploded in popularity among consumers in the last 18 months. Large-language AI, such as those apps, is programmed to accurately predict word sequences in the sentences they produce based on their training on massively large datasets, such as the entire contents

FACT

THE SHERPA, OR SYSTEM HEALTH ENABLED REAL-TIME PLANNING ADVISOR, SOFTWARE WAS BRIEFLY ORPHANED, SO TO SPEAK, ABOUT FIVE YEARS AGO.

Computer scientists at NASA’s Ames Research Center in California developed it with the proposed Resource Prospector mission in mind, but that lunar mining proposal was abandoned in 2018. The Volatiles Investigating Polar Exploration Rover was born out of the Prospector proposal a year later, and SHERPA was incorporated into the plan. It grew from research scientist Edward Balaban’s earlier development of artificial intelligence for a health management system for space hardware. The idea was to identify a problem and how to solve it even when sensor readings pointed to several potential issues or the readings conflicted with one another, as sometimes happens with stuck engine valves in or pressure readings in a propellant tank. This early AI was incorporated into the X-34 reusable spaceplane, a planned uncrewed hypersonic test bed whose development was canceled in 2001, and three years later into a proposed automated drill for a Mars lander. — KB

of the internet. SHERPA, on the other hand, has been adapted to the fundamental difficulty of planning for an unknown environment without much data to go on. For VIPER, the lunar terrain has been mapped by lunar orbiters but only at about a 1-meter resolution, leaving room for plenty of rocks and craters that could hamper the rover's travels, Wolfe says. In the permanently shadowed regions, the mapping is even more fuzzy, relying only on laser altimeter measurements taken by orbiting spacecraft.

SHERPA will define each state of the rover with 12 to 15 variables, including the time, its location, battery charge, how much time before it loses sunlight at the current location, the quality of terrain and the predicted depth of ice at the current location.

Balaban and his team needed a starting point for the VIPER version of SHERPA, so they built it from an AI model intended for Resource Prospector, a proposed lunar mining mission canceled in 2018, one year before VIPER began.


For Resource Prospector, they had initially tried to develop a look-up table model, in which each decision branch and its outcome would be predetermined. They built large decision trees, with predictions of what could happen at each branch based on probability distributions for many of the unknown factors, learning the best decisions to make based on the best results on average. But that fine-grained approach to building the AI model proved too difficult: The computations required were too great for a computer to manage for Resource Prospector's proposed two-week mission.

For VIPER's 100-day mission, that computation problem was compounded. As they continued to develop SHERPA for VIPER, they decided to break the AI model's planning into smaller pieces. Instead of calculating every possible decision branch ahead of time, they tried having SHERPA calculate only the future branches that would extend from VIPER's current state.

That was also too complex for a computer to calculate, so they tried breaking up SHERPA's planning into individual lunar days, generating mission plans for the two-week periods when VIPER would be active between hibernations, rather than for the entire 100 days all at once.

But even focusing on just the two-week periods was too large of a computational problem — it took hours to run through just one proposed route tree. So they tried a different track. "My team is probably sick of me saying this by now," Balaban says, "but I always tell them that when we are faced with a difficulty, let's try to think of how a smart human would" solve it.

Instead of planning out precisely how much science VIPER would accomplish during a two-week period, they tried to capture how much potential there was to do good science: to reach the best drilling locations,



▲ This composite of a plateau above the Nobile Crater, VIPER's planned landing site, was assembled from 147 images taken by NASA's Lunar Reconnaissance Orbiter.

NASA

and the most locations, during the 100-day mission without putting the rover at risk of missing a safe haven for hibernation if something goes wrong.

To capture the idea of science potential, they came up with a model based on approximations with "just enough" information on the risks and benefits of possible routes for the rover. This avoided the complex computations required of their previous approaches. Then they tested the outline with simulations of random problems that could occur on the lunar surface and filled in the plan with more details.

Breaking down the problem this way also allowed the scientists to weigh in with their preferences, as potential routes with the biggest science potential were identified first, without an exact route and timetable established at first, Balaban says.

"When a planetary scientist looks at a map and he or she sees a little crater here and from their experience, they can tell, 'Well, by the shape of this crater, I can tell it's a relatively fresh crater and since we're nearby, it would be kind of nice to drive by it and look at the ejecta of that crater.'"

Once the humans are happy with the template of the basic route and science targets, SHERPA will calculate — before VIPER's journey to the moon — the precise route and how long to pause at each science station.

There's also the issue of unforeseen circumstances. Once VIPER is on the lunar surface, conditions could be different than what was expected. Operators expect the rover will drive within a range of speeds, but the motor on one of VIPER's wheels could malfunction, for instance, slowing its pace below expectations. Or initial drilling could reveal that the best



places for finding ice are different than what the scientists predicted.

“We may need to replan quickly, based on both good events and bad events,” Balaban says.

Before VIPER hits the lunar ground, the plan will be constructed based on probability distributions for unknown quantities, such as how long the science activities will take, the speed of the rover and its power consumption inside the permanently shadowed regions. As the mission unfolds, the model will be updated.

With this outline-first, details-later construction of the route planning, SHERPA can compute the best plans much more quickly and react to unplanned circumstances, with some contingency branches planned ahead of time. For instance, if time is running out to reach a preplanned safe haven before one of the two-week communications blackouts, contingency plans could have an alternative safe haven site lined up, along with the subsequent plans for starting a new route from that point after the

“All these permutations that I was thinking about for a mission that’s supposed to be fast-paced and very productive — it was hard to think through it just with a pen and paper. It just blows up exponentially.”

— Edward Balaban

◀ To practice maneuvering VIPER, test versions of the rover were driven around a soil bin at NASA’s Glenn Research Center in Ohio, representing lunar terrain. Here, an engineer wearing a suit for protection from dust places rocks in the soil to simulate the consistency that VIPER could encounter as it traverses near the lunar south pole. NASA

blackout concludes. SHERPA will also identify weak spots and bottlenecks in the original plan and automatically construct contingency branches that could be chosen to avoid those issues, if needed. The route plan models will be updated during the two-week hibernations, based on new information gathered during the actual operations on the moon.

After VIPER, Balaban sees potential applications for SHERPA with a variety of other NASA missions. Because SHERPA was designed to be modular and reusable, “we think of it as a Lego set for mission planning and execution,” he says.

One possibility is applying SHERPA to unusually complex missions. For example, SHERPA could plan routes and tasks for multiple rovers or human-robotic teams as they work in coordination at multiple sites, perhaps for an astrobiology mission or the construction of a lunar base.

SHERPA could also generate an offline AI for a future mission like Europa Lander, a proposed spacecraft that would look for signs of life on the icy moon of Jupiter, where extreme radiation would be expected to fry even a heavily shielded onboard computer within 20 days.

They’re also developing SHERPA for future missions to become more conversational and interactive, Balaban says, so it’s easier for humans to offer their input and for the AI to automatically explain its decisions and the pros and cons of alternative decisions.

“If SHERPA can help us sort through some of these options and make it easier for human operators to see what are some of the better or worse options, then humans still have the final say,” he says, “but they are hopefully operating from a much more informed position.” ★