Six Minutes of Terror

Wayne Lee

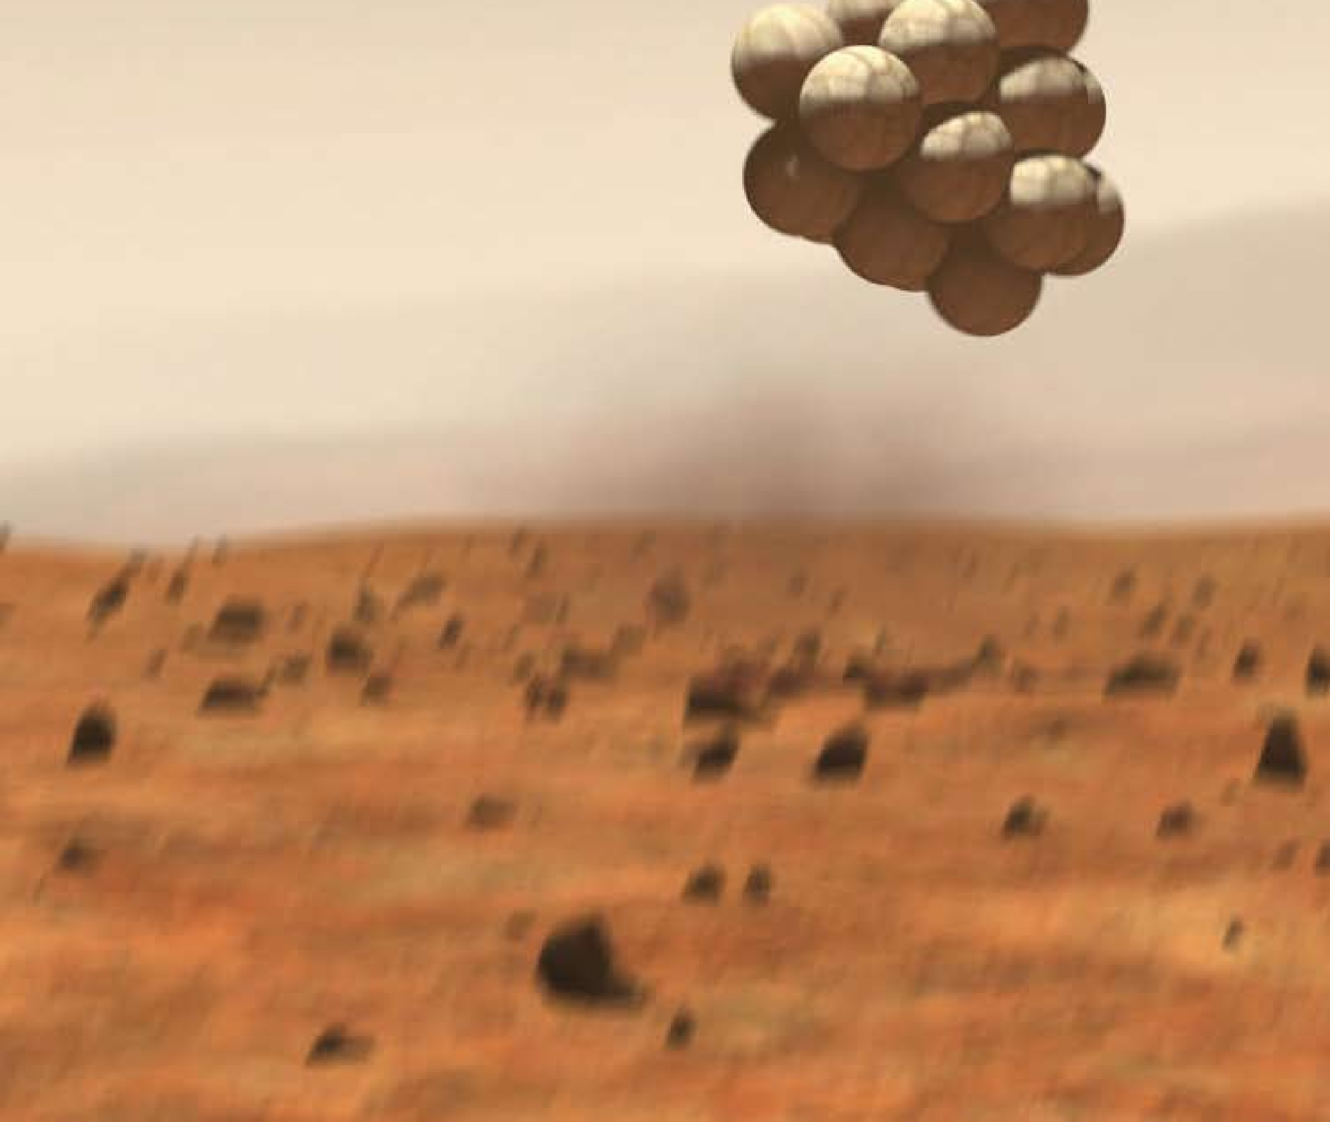

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Until that night, I had thought that moments frozen in time where your life flashes in front of you only occurred in the movies. I glanced down at my watch. In California, it was 8:40 p.m. on January 3, 2004, and it seemed as if we had hit the ground over an hour ago. In reality, our rover Spirit had landed on Mars only five minutes ago, but its radio beacon was nowhere to be heard. We had no idea whether Spirit was safely on the ground, but unable to communicate for reasons unknown, or whether the rover's tiny electronic heartbeat had been terminated by an untimely demise.

I looked down at my display console in the mission control center at the NASA Jet Propulsion Laboratory. My boss, Rob Manning, was sitting next to me. He called up a window with a plot of Spirit's signal strength as a function of time. I was frantically pointing to the flat line indicating a zero signal at the current time. Manning, an eternal optimist, and a virtual legend within the NASA robotic space flight community by virtue of having pulled off the only previous Mars landing in recent history, calmly pointed to a blip on the graph representing a point in time somewhere in the past.

Unfortunately, we had no idea whether the blip indicating a positive radio signal was before or after the expected landing time. If it were after, then we would have had positive evidence that Sprit had at least survived the initial impact. The problem was that in our haste to get the mission control display software ready for landing day, we had forgotten to program the computer to stamp the tick marks on the graph with the time of signal receipt. There was no way to tell for certain.

I looked around the mission control room and saw seemingly optimistic faces, but they were somehow unable to mask uncomfortable body language present only when one has a knot in the pit of their stomach. The supremely arrogant side of me wanted to tell everybody to have a little faith that Sprit was still alive. The other side of me was experiencing the so-called "life flashback" phenomena. I was mentally reviewing every decision we had made over the previous 40 months of designing, building, and testing the landing system for the rover and beginning to second-guess a fair number of them.

One of my first thoughts was to question whether we should have even attempted something this ambitious on such a short time scale. Unfortunately, we did not have much of a choice. The world of Mars exploration 40 months prior was clouded with uncertainty about the future. Spectacular success from the Independence Day landing of Mars Pathfinder in 1997 was followed by the embarrassing crash of the Mars Polar Lander shortly before its scheduled touchdown in December 1999. Subsequently, NASA management in Washington put all future Mars shots on hold due to a loss of confidence in the once-proud program.


NASA/nasaimages.org

In early 2000, a small group of respected engineers, including Manning and future Spirit mission operations manager Mark Adler, concocted a seemingly innocuous proposal to achieve redemption. Their theory was seemingly simple and foolproof. Why not re-fly the same landing system that led to the wildly successful Mars Pathfinder landing back in 1997? Unlike the ill-fated Polar Lander, Pathfinder was a proven landing system. "We even have spare parts left over," they argued.

The catch was that the payload Pathfinder put on the surface, a small pyramidal-shaped base station barely knee-height in size, and a six-wheeled rover the size of a small laser printer, was scientifically uninspiring with respect to a second flight. Back in 1997, the goal was simply to demonstrate that NASA was still capable of landing something on Mars despite having not done so since the Viking missions back in 1976. However, the unwritten laws governing efficient use of exploration funds mandated that the next mission following Pathfinder not only land safely, but deliver more advanced science.

"No problem," countered Manning and Adler. What if the small base station and tiny rover was replaced with a single, larger rover capable of roaming a kilometer from the landing site over a period of 90 days? For a science payload, Manning and Adler proposed to tap into the ingenuity of professor Steve Squyres from Cornell University. Squyres was a well-respected geologist who was in the midst of developing a sophisticated set of tools for a future mission to return Martian rocks to Earth using robotic vehicles. These instruments would allow the rover to both remotely sense the chemical composition of rocks and drill into their interiors for a microscopic examination. Overall, the concept sounded great, but I have to admit that I believed there was no chance that the mission would be approved.

I was quite surprised in August 2000 when senior NASA management not only announced approval for the mission, but also decided to fund two rovers with each one flying on a separate launch. The reasoning was simple. In theory, the odds for achieving at least one successful landing would dramatically increase if we sent two rovers. In practice, this strategy resulted in a tremendous amount of pressure on the flight team to go "two for two" because the natural tendency of the media would focus attention on the failed mission rather than the successful one. As an added measure of pressure, NASA management asked for a launch in June 2003. Spacecraft normally take 48 months or more to design and build. We were given only 33. The chance for redemption was not going to come easily.

The mission was officially dubbed the "Mars Exploration Rovers," or MER for short. Just prior to launch, NASA would christen the two as Spirit and Opportunity. However, for the first three years of the effort, we knew them as an impersonal MER-A and MER-B, respectively. While most of the project's engineers went off to figure out how to design a sophisticated six-wheeled rover the size of a small ride-on lawn mower, I was given the assignment of leading what was initially a small group of engineers challenged with figuring out how to land the vehicles in one piece.

Strangely enough, a good way to visualize the enormous challenge we faced in landing the rovers on Mars is to think about a rocket launch. These events are quite spectacular to watch due to the sheer amount of energy released. In fact, in order to send Spirit and Opportunity to Mars, enough energy was released by the two Delta 2 launch vehicles to propel the vehicle to a speed over 30 times faster than a speeding bullet at the time of rocket burnout. At that speed, one could fly from Sydney to Los Angeles in about 15 minutes. By the time the two rovers reached Mars about seven months later, they were still moving at a respectable speed of nearly 22,000 kilometers per hour. Put simply, that is a speed in excess of 25 times the speed of sound on Mars, or what aerodynamicists refer to as Mach 25. The challenge was in removing all of that remaining energy from the system in under six minutes.

Each Delta 2 rocket utilized a stack of fuel 35 meters high in order to propel Spirit and Opportunity to Mars. It would have been practically impossible for our tiny rovers to carry an equivalent amount of rocket fuel for

NASA/nasaimages.org

deceleration at Mars. Instead, we relied on atmospheric drag to slow the vehicles. However, this seemingly clever solution was not without its drawbacks. When an object moves through an atmosphere at hypersonic speeds, the collision with the air molecules slows the vehicle, but at the expense of generating an enormous amount of heat. Our computer simulations indicated that the rovers would be subject to heating of about 60 watts per square centimeter. Although this amount may not sound impressive, imagine holding onto 25 incandescent light bulbs in the palm of your hand. If the rovers had been directly subjected to the heat from atmospheric entry, they would have been incinerated.

The key to survival was to encapsulate the vehicles in a protective shell. Each rover was designed so that its six wheels and wing-like solar panels could fold up to allow the entire assembly to fit in a pyramidal shaped volume. Once folded, the vehicle was placed on the bottom face of a metal tetrahedron split open with the other three edges folded down. The edges were then folded back up to encapsulate the rover with the tetrahedron. Next, the encapsulated assembly was placed inside a composite backshell structure shaped like a blunt cone with an open bottom. A saucer-shaped heatshield was then attached to bottom, open face of the backshell to seal the tetrahedron inside. This entire system was designed to fly into the Martian atmosphere seemingly backward with the blunt end facing forward.

Starting with Mercury, followed by Gemini, and ending with Apollo, three generations of astronaut ferrying spacecraft utilized this blunt-body capsule shape to successfully reenter the Earth's atmosphere. We put the same idea to use at Mars. Harvey Allen, who was one of the foremost American aerodynamic geniuses of the 1950s, pioneered the theory that this seemingly unintuitive shape from an aerodynamic perspective would both provide a tremendous amount of drag, while limiting the amount of heat generated by the deceleration process. Nevertheless, at the point of maximum deceleration about two minutes after entry into the Martian atmosphere, our analysis indicated that the outer skin of the heatshield

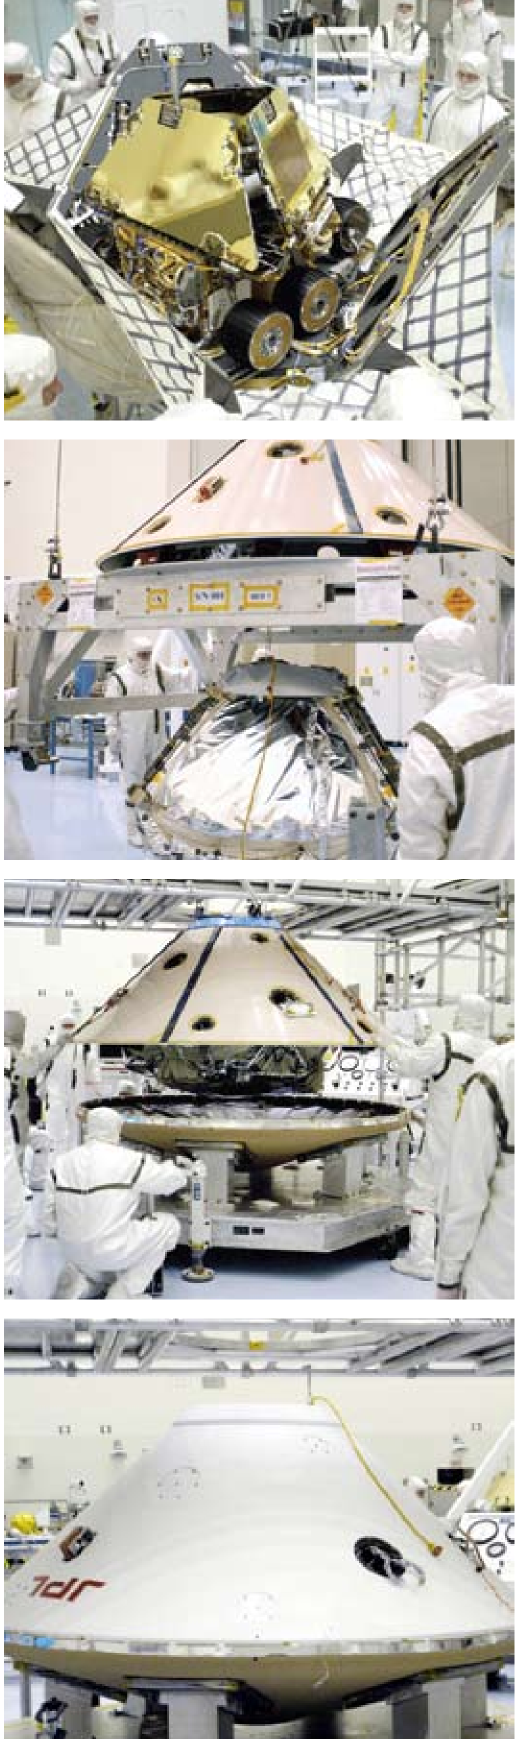

NASA/nasaimages.org

would be exposed to a temperature of nearly 1500 degrees C.

As implied by the name, the job of the saucer-shaped heatshield was to take the brunt of the heating from flying through the atmosphere at hypersonic speeds. The skeleton was constructed from carbon composite, but coated with a special cork-like material dubbed SLA-561V. Within this code-name, the "A" stands for "ablator" which is a type of material that chars at low to moderate heat rates, and then dissociates and flakes off at higher levels of heating. This process protects an entry capsule in two ways. First, the heating energy goes into charring the heatshield rather than the tetrahedron or rover within. Additional heat is subsequently carried away from the capsule within the dissociated material.

One of the initial problems we faced was in determining an adequate amount of SLA-561V with which to coat the heatshield. If we constructed the ablator layer too thin, we would risk the danger of burn through. Unfortunately, we faced a lot of pressure from management to reduce the mass of the heatshield because the estimates for rover mass were coming in heaver than expected. We knew that the computer programs that sized the requisite thickness for the SLA-561V always erred on a thicker answer due to the uncertainty in estimating both the external heating environment and the response of the ablator material to heat. The question was, "by how much?"

After months of debate, we convinced our heatshield experts to agree to reduce the thickness from the originally recommended 1.9 centimeters down to 1.4 in an effort to shave about 20 kilograms from the system. The next step was to prove that such a thickness was adequate to prevent the capsule from overheating. Our solution involved using a remarkable facility at the NASA Ames Research Center in California called the arc jet. This machine is about 20 meters long and with a 60 MegaWatts rating, consumes enough electricity to light a small city. Essentially, the arc jet shoots a super-hot stream of gas into a small test chamber at the end of the machine. We verified our design assumptions by placing small samples of SLA-561V,

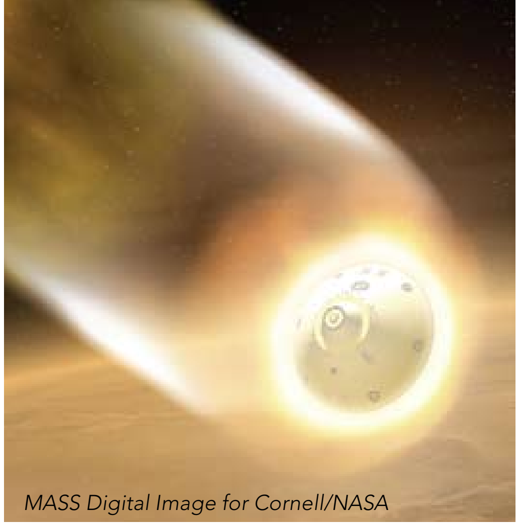

MASS Digital Image for Cornell/NASA

cut to the desired thickness, in the chamber and then measuring its performance.

Besides keeping the capsule cool, we faced another key challenge in ensuring survival during the fiery deceleration through the Martian atmosphere. Somehow, we needed to ensure that the capsule maintained enough aerodynamic stability to keep the blunt heatshield end facing forward at all times. Our aerodynamics team at the NASA Langley Research Center in Virginia initially utilized a technique called computational fluid dynamics, or CFD for short, to make the initial predictions. CFD is a technique where the area around the capsule is divided into tiny squares that form a grid. Then, a powerful supercomputer computes the flow of air around the capsule by numerically solving the equations of fluid motion within each part of the grid.

Numerical techniques such as CFD are extremely powerful tools, but limitations exist. The Langley engineers warned us that the computer solutions provided somewhat reasonable estimates for the capsule's aerodynamics at speeds greater than Mach 10, but less so at Mach 5 and below. In order to verify the accuracy of the CFD, they proposed that we conduct a test at a ballistic range facility. This sort of test would involve the use of a naval-warship-like cannon to shoot a small, palm-sized, tungsten model of the capsule down a 200-meter interior corridor. Laser activated cameras spaced at even intervals down the range would be used to photograph the model during flight. By looking at the orientation of the capsule within the photos, the Langley engineers would then be able to infer its aerodynamic properties and match them to the CFD results.

We selected the ballistic range at the Eglin Air Force Base in Florida as the location to conduct the test. The Air Force was happy to oblige us, and they invited us to come down in the Fall of 2001. In an unfortunate twist of fate, we selected September 11th as our test date. I will always remember the moment that morning when Prasun Desai, our lead flight dynamics engineer from Langley, came running into the hotel lobby to tell me that a “small plane” had just crashed into the World Trade Center. We had no idea of the magnitude of the tragedy that would unfold that day, and we had a deadline to meet, so we decided to continue onto Eglin to conduct the test.

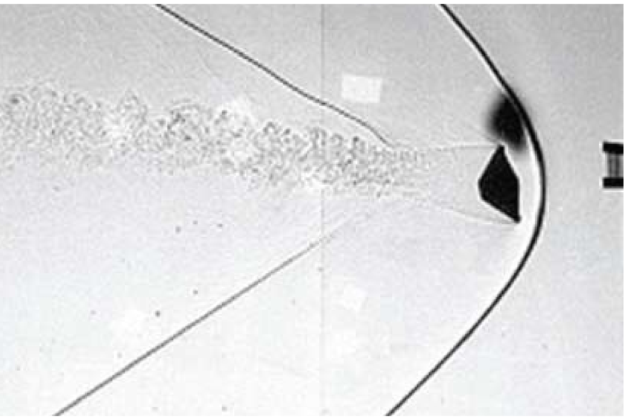

NASA/nasaimages.org

In retrospect, running the test that morning was a poor decision. We did not realize the magnitude of the large boom set off when the cannon shot the capsule down the range at Mach 5. The explosion shook the building, which was a normal effect, but also managed to rattle the nerves of base personnel who subsequently called the commander’s office to ask if the facility was under attack. Not surprisingly, the Air Force shut us down for the day. Over the next few months, we learned the bitter reality of working for the space program. We had a launch date to meet in less than two years, and a schedule slip was practically impossible because launches are possible only every 26 months when Earth and Mars are perfectly aligned. It was extremely difficult to continue to work and not be emotionally distracted by the events of the world. But, there was no choice.

Fortunately, the Air Force invited us back to Eglin a month later to complete the ballistic range testing. By then, we had begun to turn our attention to the next issue in our long queue of problems to address. Specifically, the Martian atmosphere is so thin that insufficient air exists to fully decelerate the capsule. In fact, even if the rovers survived the fiery deceleration from Mach 25, they would still impact the ground with a velocity greater than the speed of sound without further intervention. The solution involved deploying a large parachute close to the ground to further increase the drag on the capsule.

Designing a parachute suitable for use on a Mars landing presented somewhat different challenges than crafting one for skydiving on Earth. One primary issue was ensuring a sufficiently strong chute. No matter how low to the ground we chose to open the chute in order to minimize deployment speed and therefore force seen by the fabric, there was no way to avoid the violence of a supersonic deployment. Our initial calculations indicated that we required a chute with a diameter of about 9 meters and capable of withstanding the possibility of over 11 metric tons of force at a predicted inflation speed up to 1,600 kilometers per hour. Due to space limitations, we were forced to limit the mass of the chute to about 20 kilograms, and fit the entire fabric assembly in a volume barely bigger than a small, bathroom-sized trash can. In order to put a large chute in such a small space, our engineers were forced to utilize a super-thin blend of nylon and polyester for its construction.

By late spring of 2002, the team was ready to test the chute to determine whether the fragile design would withstand the forces of deployment. Unfortunately, we did not have the means to conduct a realistic flight test at supersonic speeds. Instead, our plan was to take a flight prototype chute to the Orchard Proving Grounds in Idaho, anchor it with a

2,700-kilogram weight, and then drop it out of a helicopter. The large weight was designed to subject the chute to the requisite force at the time of inflation. When the appointed test time arrived, we held our collective breaths, and then watched in shock as the chute fabric ripped to shreds. The entire test article, anchored by the huge weight, hit the ground at high speeds, and the team spent the remainder of the morning digging the contraption out of the ground. To make matters worse, post-test analysis indicated that the anchor weight was too light to generate the force expected during parachute deployment in the Martian atmosphere!

We were now faced with two problems. First, the team needed to find the weakness in the design of the fabric and the stitching. Then, we also needed to determine a method to adequately test the chute. Our solution involved going back to the Ames Research Center to utilize the world's largest wind tunnel. NASA engineers refer to this place as the "80 x 120" in reference to the cross-sectional dimensions of its test chamber in feet. I remember walking into the tunnel for the first time and feeling my jaw drop when I realized that the interior was almost as large as the Staples Center Arena that is the home to the Los Angeles Lakers, and that the six fans that pumped air through the tunnel compared in height to a two-story house.

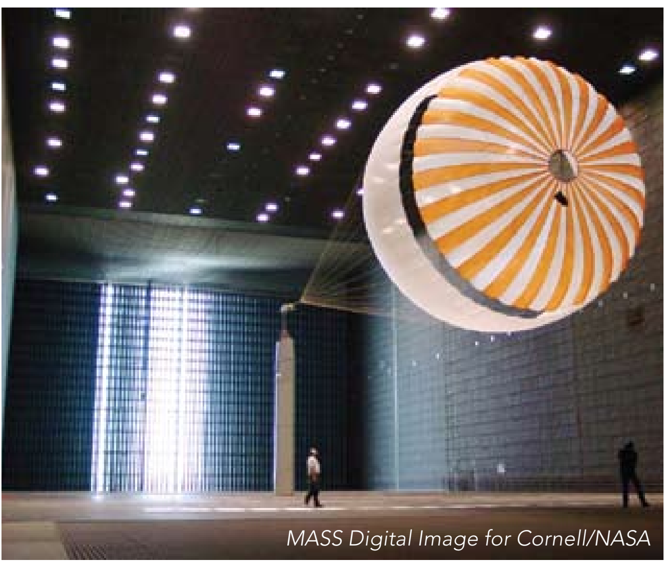

MASS Digital Image for Cornell/NASA

In theory, our two-step wind tunnel test strategy seemed foolproof. During the first step, we would inflate a prototype test chute using a low-speed wind flow, and then walk under the inflated canopy to look for suspected stress points. After applying extra stitching to reinforce those weak points, we would then put the chute back into the tunnel and turn the wind flow up to hurricane-like speeds of 80 knots to generate the same magnitude of forces expected in the thin Martian atmosphere at supersonic speeds. Unfortunately, our streak of bad luck continued as we encountered yet a third problem. During the first test in the tunnel, the chute canopy did not inflate. It simply opened and collapsed in a motion similar to a jellyfish propelling itself through water. We now had to prove that this effect, nicknamed "squidding," would not occur on Mars.

Nobody would have guessed at the time that we would spend the next eight months in a mad dash of trial and error to determine how to debug the inflation problem, strengthen the weak points in the canopy, retest the prototypes to prove to everybody's satisfaction that all was well, and then finally manufacture the actual chutes that would find their way onto the two capsules bound for Mars. Normally, manufactured components would have been delivered to the rover's assembly room at the Jet Propulsion Laboratory in California. The tightness of the schedule forced us to deliver the flight chutes for integration almost directly to the launch pad at the Kennedy Space Center in Florida. As an added measure of franticness, we executed the final verification test of the updated design prototype in June 2003, the same month as the launch.

Despite the long hours put into the design of this enormous chute, our calculations showed that the size was still too small to fully decelerate the capsule to a safe landing. In fact, we determined that after the flight-computer-commanded chute deployment at an altitude of about 8 kilometers and speed of 1,600 kilometers per hour, the capsule would decelerate to a terminal velocity of about 270 kilometers per hour in less than a minute. The laws of physics, conspiring with the thinness of the Martian atmosphere, guaranteed that the capsule would decelerate no further than terminal velocity even if we waited the entirety of the remaining 90 seconds prior to ground impact. In stark contrast, a person falling out of an airplane on Earth without a parachute reaches a terminal velocity of only 195 kilometers per hour.

So, after four minutes of fiery flight to decelerate from 22,000 kilometer per hour, and another two minutes on the parachute to slow down from 1,600 kph, we would still need at least one other way to remove the final 270 kph from the system to achieve a safe landing. In reality, our design employed two more deceleration devices. The first of the two was a relatively simple set of three, downward-pointing rocket motors mounted to the inside wall of the backshell. These "retros," as we called them, contained just enough propellant to theoretically slow the capsule to zero velocity. However, getting the vehicle in a configuration to fire the rockets was a complicated matter.

Since the rockets were mounted inside the capsule for protection from the heat of entry into the Mars atmosphere, we first needed to jettison the heatshield to expose the exhaust nozzles to the exterior environment. Then, the tetrahedron containing the rover would be mechanically lowered on a tether-like bridle into a position suspended 20 meters below the backshell. This complicated sequence was designed to complete in 40 seconds and ensured that the rover stayed well out of the way of the hot exhaust gasses expelled from the retros during firing. We were reminded of this important fact during rocket testing when a technician left a small metallic object near the test stand. The exhaust cut through the metal easier than a hot knife through butter.

Unfortunately, thousands of computer simulations told us that enough uncertain variables existed that zero velocity after retro firing would never be achieved. For example, the landing radar measuring altitude and velocity to allow the flight computer to ignite the rockets at the precise time had a 1% error potential, the amount of thrust imparted by the retros was temperature dependent and could vary by 2%, and an onboard camera system used to detect and correct wind-induced lateral motion by firing tiny horizontal-pointed rockets had an error potential of 25 kilometer per hour. With all of these sources of performance uncertainty, we calculated that the tetrahedron-encapsulated rover could still hit the ground at a speed of up to 90 kilometers per hour on a bad day.

As crazy as it sounds, we elected to allow the tetrahedron to slam into the ground and cushion the impact, as opposed to actively attempting to reduce the residual post-retro-firing velocity to zero. The second of the two final deceleration devices, a set of protective airbags encapsulating the tetrahedron, served as the key to this extreme strategy. Our space-fairing airbags operated on a conceptual principle loosely similar to automobile airbags, and were demonstrated as feasible for use at Mars during the Pathfinder landing in 1997. However, a big difference between Mars and car airbags is that ours were constructed from the bullet-proof vest material Kevlar in order to resist rock strikes on the Martian surface.

When we started working on the landing system design back in the fall of 2000, our initial plan naively involved using the Pathfinder airbag blueprints to fabricate new bags with the old design. One of our first priorities involved determining whether the old design was capable of cushioning the impact of a heavier vehicle. During Pathfinder, the total mass of equipment that hit the ground, including

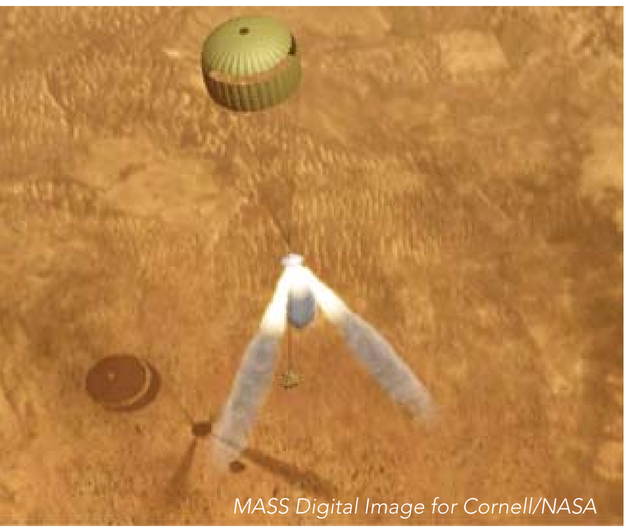
MASS Digital Image for Cornell/NASA

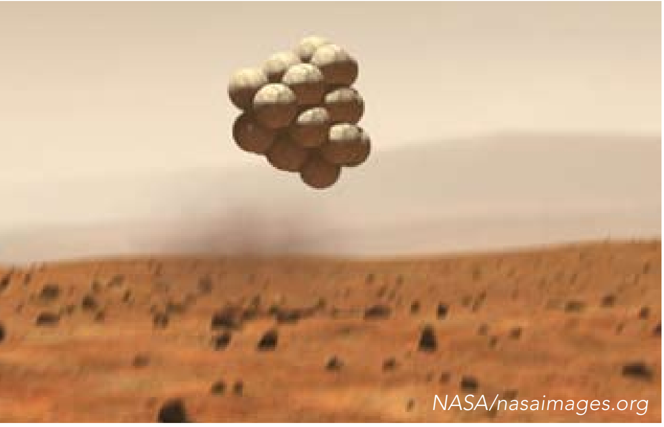

NASA/nasaimages.org

spacecraft, tetrahedron, and airbags, weighed in at less than 400 kilograms. For Spirit and Opportunity, this figure was expected to increase by an additional 150 kilograms. In order to verify performance with this extra mass, we raided the National Air and Space Museum in Washington to retrieve the Pathfinder flight spare airbags out of a display, and then pressed them back into service as a test article.

We attempted to conduct our first airbag test the week before Christmas of 2000 in the large vacuum chamber at the NASA Plumbrook Station in Sandusky, Ohio. This chamber towers nearly 100 meters high and looks almost like the containment dome of a nuclear reactor when viewed from the outside. Our test strategy involved the use of bungee cords to propel the airbags from the top of the chamber onto jagged Hawaiian volcanic rocks at the bottom. The rocks were an integral part of the test because most of our potential Martian landing sites were littered with small rocks. There was no question as to whether the airbags would encounter rocks at the time of landing. The question was, "how many rocks would the vehicle strike?"

During the first test, it was extremely sobering to see the airbags hit the rocks at the bottom of the chamber, pop, and then instantly deflate. In hindsight, this failure was the first warning that the airbag development process would be extremely challenging. At the time, we somehow refused to acknowledge reality and attributed the failure to the fact that the airbag material was weakened by prolonged exposure to ultraviolet light emitted from the powerful spotlights in the Pathfinder spacecraft assembly room. That foolish theory was disproved when the test team showed up in Ohio early the next year with a set of freshly manufactured airbags. New bags yielded the same results as they subsequently ripped upon striking the rocks in the chamber.

Now, we were in a real quandary. Airbags, by their very nature of being constructed out of fabric rather than metal, are extremely difficult to analyze using computer simulations, especially when it comes to proving that they will be resilient to tearing when striking rocks at high speeds. We quickly realized that the only way to gauge whether a design concept would work was to take a prototype into the test chamber. Unfortunately, testing was a long, tedious affair. Each test cycle consumed nearly two days and required hoisting the bags to the top of the chamber, waiting almost eight hours for the vacuum pumps to reduce the air pressure to Mars levels, retrieving the bags from the chamber after the drop, and then sewing up the ripped fabric in preparation for the next test.

After each test, the airbag design team would enter the chamber to inspect the damage induced by the rocks and ascertain the robustness of the current design concept. Each rock was of a unique size and shape, and was covered with a unique color of chalk dust. By looking at the color residue on rips in the airbag, we were able to determine the culprit rocks causing the damage. Our nemesis was a small, 30-centimeter tall rock powered with black chalk. Although seeming innocuous in size, this rock contained a sharp, tooth-like projection at the top that ripped through many design concepts. I often fell asleep at night worrying about the "black rock."

Each test failure compounded the time pressure on the team to arrive at a working solution prior to launch. A test airbag could only be patched up for eight drops before the toll of abuse rendered it useless. And, if a design concept failed, the process of redesign and manufacturing a new test prototype consumed nearly four months. In total we executed over

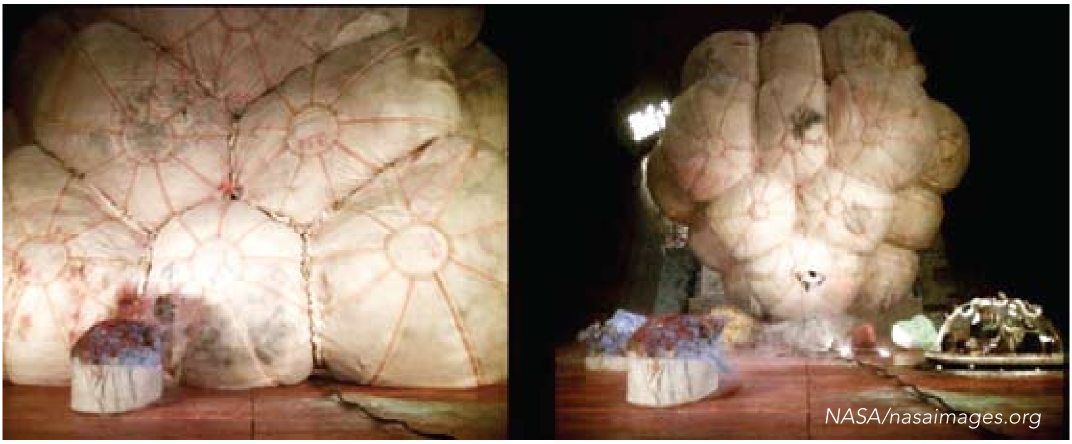

50 test drops between 2001 and 2003 to arrive at a viable concept. That final design consisted of eight layers of Kevlar in vulnerable areas to keep rocks from penetrating the inner bladder of inflation gas. And, in a desperate schedule situation similar to the parachute, the actual flight airbags bound for Mars were delivered directly to the Kennedy Space Center rather than to the spacecraft assembly facility in California.

With the myriad of technical challenges to overcome on both the landing system and rover side of the design, Spirit and Opportunity barely made it to the pad in time for their lift-offs in June and July 2003, respectively. After our hardware left the Earth, I had mistakenly thought that we were looking at seven easy and quiet months in transit to Mars. In retrospect, I should have realized that our experiences over the past 33 months were an indicator that nothing came easily on this mission. In fact, the time between launch and landing amounted to some of the busiest moments during the mission.

One of the first things we discovered after launch was that our predictions of the vehicle dynamics during retrorocket firing failed to account for all the force disturbances in the system. So, we went out into the California desert over the summer of 2003 to perform full-scale test firings in order to gather data that allowed us to reprogram the flight computer to compensate. Then, a few days prior to Christmas, a huge dust storm developed on Mars. This storm effectively thinned the Mars atmosphere and put the vehicle at risk of deploying the parachute too close to the surface. After a lot of debate with uncertain facts, we made the decision to reprogram the flight computer again. This time, we asked it to deploy the chute earlier, and while at a faster velocity, to compensate for the thin atmosphere. The down side was an increase in risk to ripping the chute due to excessive forces during inflation.

Yet another serious problem surfaced with one week to go prior to landing. Jason Willis, one of our lead avionics engineers, discovered a serious flaw in the electronics responsible for triggering pyrotechnic initiated events such as parachute deployment, heatshield jettison, airbag inflation, and retrorocket ignition. Normally, the flight computer arms these pyrotechnics for firing only seconds prior to use for safety reasons. Test results from our high fidelity electronics testbed indicated a subtle timing bug in the circuit that caused the arming command to be ineffective. The only viable solution was to order the flight computer to remove the safety inhibits and enter the atmosphere with all the pyros dangerously armed.

And, just when I thought we were finally ready despite the risky solutions we were forced to implement, I received a phone call from lead flight dynamics engineer Prasun Desai the night before Spirit's landing. Right when I was sitting down to watch a football game to unwind, Desai informed me that the team had just found a programming error in the sophisticated simulations we had been using to prove to ourselves that the onboard software would be able to fly the vehicle through the atmosphere. Although this revelation was not exactly the same as saying that the software was not going to work, it was nevertheless not a reassuring phone call. He promised me that they would work through the night to fix the simulation.

Less than 24 hours later, I discovered first hand why our Mars Exploration Program manager had dubbed landing the "six minutes of terror." At precisely 8:29 in the evening on January 3rd, 2004, the capsule containing the Spirit rover plunged into the top of the Martian atmosphere moving at 25 times the speed of sound. During that first minute of atmospheric flight, very little deceleration occurred due to the extreme thinness of the upper Martin atmosphere. I remember looking at the altimeter on my display console in mission control and watching the altitude tick off alarmingly quick at a rate of one mile every second – 70 miles, 69, 68, 67, 66.

"We're dropping like a falling rock," I muttered to myself. In reality, I probably used another word other than "falling" that began with the same letter. Fortunately, my headset microphone was off, and the comment did not get broadcast over the loop and onto national television. For the previous three years, we had studied the simulations results, and we knew just how fast the vehicle would fall on Mars. However, studying graphs and numbers falls woefully short in terms of preparing for the shocking reality of watching it happen in real time. One way or another, Spirit would reach the ground in six minutes, and the outcome of exhausting and stressful work depended on the autopilot, 40 months of sound engineering judgment, and the hands of fate. There was nothing we could do other than watch and pray.

During most of the fiery plunge through the Martian atmosphere, a tiny radio beacon from Spirit chirped out simple electronic beeps to let us know that she was still alive. When the signal terminated near the time of expected ground impact, we were left wondering whether a distant cousin of our "black rock" had sliced open the airbags, or whether the winds were too high, or whether the radar had provided a bad retrorocket firing solution, or heaven forbid, we had simply overlooked a careless mistake somewhere in the system. While Manning and I stared at my computer console looking for an answer we knew we would not find, our communications engineer, Polly Estabrook, was on the phone talking nonstop with the ground crews of our radio tracking stations around the globe.

After 15 minutes of awkward silence in the control center, Estabook startled me with an excited, "they see it, they see it!" That proclamation was followed by instant mad celebration inside mission control worthy of a winning goal scored in the final minutes of a World Cup final. I have to admit that I was probably the only one in the room who missed the celebration, at least initially. My responsibility was to call out "safe touchdown" once we established proof, and the skeptic in me wanted more evidence from Estabrook other than an excited proclamation. I never gave the call, but it hardly mattered. A little faith had already delivered the answer the others were awaiting.

Just prior to the successful landing of the Opportunity rover three weeks later, Desai assured me that the simulation was working this time around and asked if I had any plans for the evening before landing. "Don't know," I replied, "but I'm getting too old for this stuff, so I'm turning off my cell phone in case anybody calls."