http://www.cs.wm.edu/~coppit/other-papers/mco-oberg.php

Why The Mars Probe Went Off Course

By James Oberg

SPECTRUM Magazine

December 1999

AIRCRAFT ACCIDENT INVESTIGATORS have a special term for a particularly insidious type of accident--CFIT, or controlled flight into terrain. It occurs when human error in the cockpit, in the traffic control tower, or in the flight planning process in effect flies a perfectly good airplane right into the ground.

In the past 40 years, space flight has encountered all sorts of failure modes. Propulsion systems have leaked and exploded. Power systems have short-circuited. Observation instruments have failed to work or have been pointed in wrong directions. But until this year no CFIT had occurred in outer space.

Then, on 23 September, through a series of still-baffling errors, flight controllers at the Jet Propulsion Laboratory, a California Institute of Technology facility under contract to NASA, sent erroneous steering commands to the Mars Climate Orbiter as it neared the target planet. Obeying blindly like all true robots, the probe, metaphorically speaking, marched off the cliff and was destroyed.

NASA assigned three separate teams to investigate the embarrassing, US $125 million debacle and determine its cause. Preliminary public statements faulted a slip-up between the probe's builders and its operators, a failure to convert the English units of measurement used in construction into the metric units used for operation.

After six weeks, on 10 November, NASA officials released their preliminary findings. However, an *IEEE Spectrum* investigation had been going on separately, using unofficial sources associated with the program and independent experts. *Spectrum* quickly learned that far more had gone wrong than just a units conversion error. A critical flaw was a program management grown too confident and too careless, even to the point of missing opportunities to avoid the disaster.

As reconstructed by *Spectrum*, ground controllers ignored a string of indications that something was seriously wrong with the craft's trajectory, over a period of weeks if not months. But managers demanded that worriers and doubters "prove something was wrong," even though classic and fundamental principles of mission safety should have demanded that they themselves, in the presence of significant doubts, properly "prove all is right" with the flight. As a result, the probe was about 100 kilometers off course at the end of its 500-million-kilometer voyage--more than enough to accidentally hit the planet's atmosphere and be destroyed.

Edward Stone, director of the Jet Propulsion Laboratory (JPL), in Pasadena, Calif., did not try to dodge responsibility for events. "Our inability to recognize and correct this simple error has had major implications," he stated in a 24 September press release from the lab. "We have under way a thorough investigation to understand this issue." As is normal in such cases, all data was impounded for use in the accident investigations, and all participants in the mission were ordered not to talk to the press.

Even at that time, NASA managers hinted that much more had gone wrong. "People sometimes make errors," said Edward Weiler, NASA associate administrator for space science. "The problem here was not the error; it was the failure of NASA's systems engineering, and the checks and balances in our processes, to detect the error. That's why we lost the spacecraft."

Carl Pilcher, science director for solar system exploration at NASA headquarters in Washington, D.C., agreed: "Human error occurs all the time. But even so we have a tremendous success rate because we have systems that detect and correct the errors. The problem here is that our system failed to do that."

And Thomas Gavin, deputy director for space and earth science at NASA's Jet Propulsion Laboratory, added: "A single error should not bring down a $125 million mission."

In an internal memo dated 18 October, not intended for non-JPL readers, a laboratory official summarized the way thinking was developing: "There might have been some overconfidence, inadequate robustness in our processes, designs, or operations, inadequate modeling and simulation of the operations, and failure to heed early warnings." No, it was not a simple mistake at all, as NASA finally explained in detail on 10 November.

One down, one to go

NASA had urgently needed an explanation for the failure since a sister probe was also nearing Mars. Due to arrive on 3 December, the Mars Polar Lander needed to approach with even greater precision, since it would directly enter the Red Planet's atmosphere for a landing near the southern pole. A trajectory error even a tenth the size of the one that had doomed the climate orbiter would also destroy the polar lander.

The Mars Climate Orbiter --the first weather satellite for another planet--had not been intended to enter Mars's atmosphere. With an orbital observation program several years long stretching ahead of it, the aim had been for the probe to swing around behind Mars just above its atmosphere. Then its directional rockets were to fire for 15 minutes, to slow it down so that it was captured by the planet's gravity.

Next, a series of skimming motions across the upper atmosphere were to lower the probe's orbit and adjust it till it became circular, the better to chart the planet's weather patterns and search out signs of past and present water. The climate probe was also to serve as a communications relay (in the UHF band) for other probes on the planet's surface.

The launch of the Mars Climate Orbiter had gone according to plan. A Delta-II booster lifted it from Cape Canaveral, in Florida, on 11 December 1998. At the beginning of its interplanetary cruise, its 629-kg total weight included 291 kg of rocket propellant, the amount needed to slow it down on arrival. It carried a cargo of two science instruments, namely, a color camera and an infrared radiometer (a copy of an instrument lost when the Mars Observer probe disappeared in 1993); a computer with a RAD6000 processor (a radiation-hardened version of the PowerPC chip used in some Macintosh computer models); and other standard spacecraft systems for thermal control, attitude control, and propulsion.

Two features dominated the probe's appearance: its high-gain antenna and its single solar-power array [Fig. 1]. The 1.3-meter dish used microwave X-band signals with a 15-W transmitter, giving a peak data rate of about 110 kb/s. Uplink was at 7172 MHz, and downlink was at 8427 MHz; the carrier signal was phase modulated and the subcarrier was phase-shift keyed. The solar panels, 5.5 meters from tip to tip, relied on gallium arsenide cells to provide as much as 1000 W of power, supplemented by a 12-cell 16-Ah nickel-hydrogen battery.

The route planned

On its arrival near Mars, the probe was aimed to pass above the planet's north pole, within 200 km of the surface but beyond the atmosphere. This last, though much thinner at the surface than Earth's atmosphere, still extended fairly far out because of Mars's lower gravity. It was not a particularly challenging route. Previous probes had been guided very accurately--for example , in 1997 the Mars Global Surveyor, still orbiting the planet today, missed its target altitude by a mere 4 km.

During the long cruise outward from home, flight controllers had navigated using the spacecraft's radio link to Earth for orbit determination. Motion along the line of sight was measured using the doppler shift in the radio link, and a series of range measurements over a period of about two weeks could be accurately converted into the probe's actual flight path, or ephemeris.

Using knowledge of the forces acting on the spacecraft, computer programs could then calculate the flight path forward in time to see how the craft moved relative to Mars. Controllers could also "target" a change in the probe's motion in order to shift its future position relative to Mars, and then command the probe's rocket engines to carry out the prescribed course change. These were called trajectory correction maneuvers, or TCM burns.

As mentioned, accurate orbit determination and targeting any course changes required a complete knowledge of all forces acting on the probe. With this, the computer programs generate both an ephemeris and an error estimate for the flight path based on how widely scattered the navigation marks are. It was at the first step--where the forces acting on the spacecraft were input in the wrong units, according to preliminary NASA explanations after the debacle--that the computers were led astray. But what doomed the spacecraft, experts believe, were errors in human judgment that led to a poor grasp of the navigation uncertainty, in general, and of the meaning of the great range that existed in the calculation of the altitude over Mars, in particular.

Momentum wheels and jets

Like many other spacecraft in terrestrial orbit and beyond, the Mars Climate Orbiter maintained control of its attitude, or orientation in space, through the use of momentum wheels. These metal disks, measuring about 10 cm in diameter, and resembling the wheel in a child's gyroscope, are spun up or down by electric motors. Top speed is 3000 revolutions per minute. With one momentum wheel for each axis, they serve to turn the spacecraft in different directions or to steady its orientation against disturbing torques, and do so gently and with fine control. Gas jet thrusters, which are also available to turn the spacecraft as needed, do so forcefully and with coarser accuracy.

If the forces behind the reorientation of the spacecraft and its resistance to torque were random, the devices could spin up and slow down within their range of operating speeds, and this process could go on forever. But the primary torque on the Mars spacecraft was from sunlight itself (from photon pressure), which was not random. This was particularly significant because the asymmetry of the spacecraft's solar array gave rise to a disturbing force in a single direction. To counter this, some of the momentum wheels had to spin faster and faster.

Periodically, one or more of the wheels would come close to spinning too fast for safety, and the momentum would have to be dumped--a process that could happen as often as once or twice a day. Dumping involves deliberately spinning the wheels down, and so also turning the spacecraft, while at the same time firing small gas jets to counteract this turning force. JPL calls this an angular momentum desaturation (AMD) maneuver.

Momentum wheels are used on unmanned and manned spacecraft. The Solar and Heliospheric Observatory (SOHO) has them, and in mid-1998 lost control during a momentum dump. The Skylab space station (1973-4) had a larger version, as does Russia's Mir space station, on which they are called gyrodynes.

For dumping some of the momentum stored in these wheels, spacecraft near the Earth can dispense with the jet firings, and the propellant needed to supply them. Instead, they can perform the operation fuel-cost free by utilizing other known torques. Sometimes they use the force due to the gravity gradient along the length of the spacecraft that results from being very close to the massive Earth. Sometimes they use magnetic torquers--a device that has a current running through a wire loop--to push against Earth's magnetic field.

But some momentum dumps also occur that use propulsive attitude control by firing jet thrusters. In deep space the method is the only practical one, despite its imperfections.

The use of jet thrusters for attitude control raises further operational issues. In a world of perfect symmetry and unlimited payload size and budget, a spacecraft could rotate cleanly about its center of mass (often carelessly called its center of gravity) if opposing ends were equipped with jets and if those jets pointed in opposite directions and were set at right angles to the axis to be turned.

In the real world, rotational jets may not be arranged in such a theoretically perfect alignment. Only one set, at one end of the spacecraft, might be installed; symmetrically opposed thrusting would then be unavailable for achieving pure rotation. And even then, the jets may not be able to point precisely at right angles to the spacecraft's axis because of mechanical constraints or concerns over where the jet plume may impinge on spacecraft appendages.

On the Mars Climate Orbiter, four separate clusters of jets were located around the vehicle's waist [Fig. 1, again].

However, because of the large solar array extending from one side, the craft's center of mass did not coincide with the center point of the waist. Thus there was a significant imbalance each time these small thrusters fired.

This arrangement results in what space engineers call a cross coupling of forces between rotational axes and pure translation (where translation means moving the spacecraft away from its original location or course) [Fig. 2]. For the climate probe, it did indeed turn upon firing a jet that was mounted at some distance from its center and pointed at right angles to the axis to be turned. But the fired jet fails to do exactly as desired, for it also slightly pushes the spacecraft--it cross-couples into translation--in the direction opposite of the gas jet's thrust. (Most of the thrust does go into turning the spacecraft about the center of mass.)

At the 10 November news conference, NASA revealed even more damning information about this issue. Fully aware of the issue of asymmetric torques, the original designers of the probe had planned to neutralize them by slowly spinning the spacecraft about its long axis, perpendicular to the sun (the so-called barbecue mode). Some time later, concern over a potential shortfall in the electric power budget caused the design team to change this balancing spin to a constant face-onto-the-sun orientation. There were no navigation experts on the team at this point (they were in fact not added until two months before launch, and had no significant knowledge of the spacecraft's peculiarities even then). So the change was made without opposition.

Mars's two measurement systems

Because it used momentum wheels for fine pointing control, the Mars Climate Orbiter also performed momentum dump operations periodically during its cruise out to its destination. The flight controllers at JPL observed the jet firings that occurred to control the probe's orientation during these maneuvers. And they then would have taken into account the minor--but critical--translational cross-coupling forces that the jet firings had induced.

According to early NASA statements about the failure, the trajectory problem began at this point. The spacecraft experts back at the factory had calculated how much translational force each rotation jet accidentally induced when it fired, and the amount was proportional to how long it went on firing, which JPL controllers easily measured. They then could multiply the known force by the observed duration of its application, and update the spacecraft's navigation computer with the calculated course change.

NASA did not originally specify the actual units used, but the unit that JPL would have wanted was the newton (the force that accelerates a 1 kilogram mass at a rate of 1 meter per second per second).

The corresponding British unit was the pound force. If the values provided by the spacecraft engineers at Lockheed Martin Astronautics Co., Denver, Colo., had been in pound force, they would have been too large by a factor of 4.45.

But in engineering terms, these two values are still of the same order of magnitude. In other words, there is no really gross mismatch in the scale of calculations made with the one or the other--the kind of mismatch that can provide an intuitive hint that something doesn't add up. These two units were close enough in magnitude that the unintentional substitution of one for the other apparently rang no warning bells.

According to a JPL spokesman, every maneuver intended to dump momentum added a velocity error of about 0.001 meter per second, on a probe that was traveling at a rate of tens of kilometers per second. These deflections themselves were not the problem, but their incorrect modeling was, when the computer was told the spacecraft had received a force of four or five times as great as it really had.

"Every momentum dump will introduce an additional bogus force into the navigation software," *Spectrum*'s interplanetary navigation consultant explained. "Ultimately the curve fit will be warped by this constant application of a non-existent force."

Navigators back on Earth used radio doppler tracking to estimate the actual course of the probe and compute the required trajectory correction maneuvers, the TCMs. TCM-1 on 21 December was fairly large; it had to make up for booster insertion errors, as well as to correct a trajectory deliberately aimed away from Mars, to prevent the last stage of the unsterilized booster from hitting the planet. But TCM-2 on 4 March was very gentle, about 0.86 m/s, in keeping with a well-navigated on-course spacecraft.

For TCM-3 on July 23, JPL implemented a new feature of its navigation function. The official press release referred to it as "tuning the spacecraft's autopilot."

Outside observers correctly interpreted this as the use of new subroutines-- a so-called small forces model--to get extra accuracy by estimating minor forces such as the cross-coupled force from momentum dumps. But if inaccurate forces had been introduced into the orbit determination software, then the position of the spacecraft at the moment of the burn would be in error, and so would the desired size of the burn itself.

At the 10 November press conference, chief investigator Arthur Stephenson explained why the initial trajectory had appeared healthy. There had been a software error in the implementation of the small forces model that was not fixed until April. This delay had forced JPL to make its own estimates of the course deviation from the AMD maneuvers, and they had done this accurately. Because of the rush to get the small forces model operational, the testing program had been abbreviated, Stephenson admitted. "Had we done end-to-end testing," he stated at the press conference, "we believe this error would have been caught." But the rushed and inadequate preparations left no time to do it right.

Feelings of unease

Spectrum's investigation uncovered a report that one navigator wrote a memo describing some vague uneasiness he had about the trajectory as planning began for the TCM-3 burn on 23 July. Possibly, the unease was due to an intuitively correct appreciation of higher-than-expected course-error estimates.

The size of the midcourse burns was large, unlike the experience with the Mars Global Surveyor two years earlier. In that case, the third midcourse correction was canceled entirely, and the fourth and final one was only 0.29 m/s.

But TCM-3 for the climate orbiter was 3.3 m/s, and TCM-4 was almost 2 m/s.

NASA confirmed this navigational "unease." According to Stephenson, within days of the activation of the small forces model, navigators began to suspect something was not quite right with the spacecraft's ephemeris. "We even had a big meeting in Denver about this," he told the press conference on 10 November. "But the issue was never resolved."

Although the navigators continued to express concern about the spacecraft trajectory, NASA's Stephenson explained why there had been no management response. "They did not use the existing formal process for such concerns," he stated. JPL has a special form to invoke a so-called incident surprise and analysis procedure, and the navigators did not follow the rules about filling out that form to document their concerns.

Stephenson did admit that inadequate navigation team staffing was a contributing factor to the accident because they were responsible for three separate missions at the same time. In addition, their training in team operations was inadequate, according to Stephenson. "This was a problem of transition from the era of very large teams to when teams are very small," he explained.

Richard Cook, a trajectory expert at JPL, discussed the nominal flight path with *Spectrum.* Prior to the last midcourse maneuver in September, he reported, the probe was headed for an impact with Mars. But the TCM-4 rocket burn moved its fly-by altitude about 800 km farther away from the center of the planet, more than enough to clear the atmosphere--or so the navigation computers claimed.

But even if the targeting software worked perfectly--and some experts believed it did--the maneuver doomed the probe because the spacecraft was not really where the navigators thought it was. The trajectory errors induced by the wrong units in the small forces model made Earth think the spacecraft was several hundred kilometers from where it actually was. This mistake would displace the point of closest approach to Mars by about the same amount. This was confirmed by NASA on 10 November.

After the fourth rocket burn, navigators began taking new marks in order to determine if a final adjustment was required. This maneuver would have been TCM-5, two days before encounter. It normally took many days to accumulate enough marks to generate an accurate orbit, and as the probe neared Mars, geometry conspired to reduce accuracy. When it had left Earth, its velocity was mainly aligned directly away from Earth, so doppler ranging measurements provided good data. But now as it neared Mars, its velocity was mostly perpendicular to the Earth-Mars line, so doppler data was far less precise [Fig . 3].

Reportedly, the first trajectory measurements after the fourth burn showed the probe to be right on course for a 193-km close point (they had aimed for 224 km, just to be on the safe side). Still, an error of 30 km when the expected error was less than ±10 km may have been slightly disturbing.

In the following days, as additional tracking marks were accumulated, the navigation programs began to show the estimated closest point sinking lower and lower. It looked as if the probe was drifting off course, an impossibility if all the forces on it were properly modeled. What was really happening was that as the probe fell towards Mars, its increasing speed helped refine the probe's true path. But instead of believing that the newer values were closer to reality, controllers apparently chose to trust the earlier navigation, and suspect instead that something was now going wrong with the navigation software. The spacecraft itself, they assumed, remained on a safe trajectory.

No final trajectory correction

The 4 October issue of *Aviation Week* magazine provided detailed navigation data from this period, based on its inside sources at JPL and confirmed by other experts who talked anonymously with *Spectrum*. According to the magazine, "Several days after the TCM-4, the navigation calculations had relatively poor convergence. The new numbers were trending to 150-180 km--but with uncertain confidence." The phrase "uncertain confidence" probably means that the calculated error bounds had gotten quite high. Yet even though some predictions were up to 70 km off the aim point, navigators still acted as if they believed the aiming accuracy was within 10 km.

According to *Aviation Week*, project management decided on 19 September to forgo TCM-5, because of a flyby range predicted at 150 to 180 km and a belief that the space craft would be safe to 85 km.

But the 85-km figure was not based on actual engineering analysis. According to a control center expert who has spoken with *Spectrum*, JPL operators put their faith in a 160-km fly-by altitude until a few hours before encounter. Then a recomputation showed that the altitude would be 110 km. When Lockheed Martin engineers were asked to examine the effects at this altitude, they were frightened. Atmospheric drag would probably be enough to tumble the spacecraft and overheat it. As the spacecraft passed behind Mars, navigators came up with a new estimate: 95 km. This would generate heating equivalent to a bank of propane torches.

A sense of gloom descended on the controllers as they began listening for the reacquisition of radio contact on the other side of Mars. By then, JPL navigators had an even newer estimate--57 km, where the heating would be 10 times as bad as at 95 km.

Spectrum has been told that this decision to forgo the TCM-5 correction was flawed: "Given expected errors in altitude targeting of about 10 km, a spread of values over a 100-km range [from 70 to 180 km] should have people screaming down the halls," one navigation expert told us. "This tells you that you have no idea where your spacecraft is, and therefore your trajectory has an unacceptable probability of intersecting the planet's atmosphere. To me this says 'aim high' and put another 200 km in there to be safe."

Reportedly there were more indications of trouble. After the last burn, when navigation was hinting at a 100-km range of uncertainty of true position, somebody at JPL ran the data through the 1998 Mars Pathfinder navigation code (different from the Mars probe code). It showed the spacecraft was off course by hundreds of kilometers, which turned out to be correct.

Asked about these indications by *Spectrum* at the 10 November press conference, JPL officials denied any knowledge. When *Spectrum* then asked if the main actors in the navigation misjudgments could now be allowed to speak freely with the press, JPL director Edward Stone stated coldly that "everyone's full attention must be on the December 3 [Mars Polar Lander] landing." Another journalist jumped on what he called this "artfully unresponsive" answer and got Stone to concede that perhaps some time next spring, after the final report is issued, the investigation gag order might be lifted.

Rumors even assert that the leader of the navigation team, Pat Esposito, had recommended making the TCM-5 burn to raise the fly-by range "just in case." He declined comment when phoned by Lee Dye, retired *Los Angeles Times* science editor and now a contributor to the abcnews.com science Web page. Dye then asked Thomas Gavin, deputy director of JPL's space and earth science directorate, about rumors that Esposito had personally called Gavin to urge making the burn. Gavin vehemently denied it. "I thought the telephone line was going to vaporize," Dye reported.

NASA presented a strikingly different version of the dispute. Stephenson admitted that in fact the navigation team had verbally requested the TCM-5 maneuver be performed. Permission was denied-- "properly," Stephenson asserted--because the team was not prepared to perform the burn; no procedures for the fast response needed had been put in place or practiced.

Besides, Stephenson continued, "the navigation team was not clear on what the problem was, but they did not see it as a case of possible loss of spacecraft." He concluded: "Unfortunately, maybe they didn't see it as a big issue, so they didn't make it into a big issue."

JPL director Stone endorsed this. "Even on the day of encounter it was not clear to anyone that we were on the wrong trajectory," he insisted. "No one."

This is the remaining inconsistency between NASA's official version of what happened and the one reconstructed by *Spectrum*. Our conclusion is that adequate doubts had been raised to require the TCM-5 burn, even in an emergency mode. Further, according to participants in this tragedy of errors, by the time the probe reached Mars, those most "in the know" were persuaded it was already doomed by its sick trajectory--but by then it was too late.

All this information has modified perceptions of the widely published images of glum space controllers waiting hopefully for the probe to regain contact after it passed behind Mars. They were not astonished, or caught by surprise by an anomaly whose cause they could only guess. They had seen the most recent tracking data, which spelled doom, and they were hoping not for the spacecraft to emerge but for a miracle, which never materialized.

As for the Mars Polar Lander, JPL controllers suspended all correction maneuvers early in October, pending assessment of the navigation error on the Mars Climate Orbiter. But by 11 October they were satisfied they could avoid making the same mistake again, and after a few final delays, performed the burn on 30 October.

A NASA press release explained this new confidence this way: "Extensive analysis of spacecraft data by the flight teams at NASA's Jet Propulsion Laboratory in Pasadena, Calif., and Lockheed Martin Astronautics in Denver, Colo., has confirmed that the lander does not have the same unit conversion error that contributed to the loss of the Mars Climate Orbiter last month."

Although NASA suggested that because the lander had no reaction wheels, it was immune to the subtle problems that seduced the orbiter's navigators, this was not true. In its long cruise from Earth to Mars, the lander still had to fire attitude thrusters to stay pointed in space. True, the lander's structure was much more symmetrical than the orbiter's, and this greatly reduced the amount of photon pressure torque that had to be counterbalanced. But these unbalanced firings still induced small translational forces. Its small forces model still had to work properly if those disturbances were to be properly modeled. However, with the lesson of the orbiter behind them, there's every reason to expect that JPL navigators have avoided making the same mistake twice.

Roots of the error

Even if what ruined the Mars Climate Orbiter mission can be overcome, it should not be forgotten. The analogies with the Challenger disaster are illuminating, as several direct participants in the flight have independently told *Spectrum.*

In that situation, managers chose to cling to assumptions of "goodness" even as engineers insisted the situation had strayed too far into untested conditions, too far "away from goodness." The engineers were challenged to "prove it ISN'T safe," when every dictum of sound flight safety teaches that safety is a quality that must be established--and reestablished under new conditions--by sound analysis of all hazards. "Take off your engineering hat and put on your management hat" was the advice given to one wavering worker, who eventually went along with the launch decision.

Similarly, various versions of the trajectory debate in the final days of the flight indicate that in the face of uncertainty, decision-makers clung to the assumption of goodness; assertions of trajectory trouble had to be proved rigorously. Just the opposite attitude should have ruled the debate.

Other complaints about JPL go more directly to its existing style. One of *Spectrum*'s chief sources for this story blamed that style on "JPL's process of 'cowboy' programming, and their insistence on using 30-year-old trajectory code that can neither be run, seen, or verified by anyone or anything external to JPL." He went on: "Sure, someone at Lockheed made a small error. If JPL did real software configuration and control, the error never would have gotten by the door." Other sources commented that this problem was particularly severe within the JPL navigation team, rather than being a JPL-wide complaint.

In the meantime, out in space, where did all this leave the Mars Climate Orbiter? Behind Mars and out of sight from Earth, halfway through its braking burn, it encountered a far thicker Martian atmosphere than had been expected. Something catastrophic must have occurred. At first, experts thought it might have burned up like a meteorite, scattering its fragments across the Martian north pole.

Later, JPL calculated that rising engine temperatures from atmospheric impact would have triggered an automatic shutdown, followed by structural failure from deceleration loads. Alternatively, the aerodynamic torque on the spacecraft's solar panel may have twisted the craft into a fatal tumble.

In an analysis done by the spacecraft's builders at Lockheed-Martin, the spacecraft was almost certainly destroyed when its hydrazine propellant tank heated to the point of self-ignition of the remaining fuel. "There was enough explosive force there to level a city block," one engineer told *Spectrum,* so the spacecraft probably was blown apart into shards of scrap metal that soon burned up in the Martian atmosphere.

Whether the debris burned up, fell to the surface, or grazed the atmosphere fast enough to have passed out into space on the other side of Mars is not known. For now, it is truly "lost in space."

But if space mission operators learn their harsh lessons better, the exploration of Mars will push ahead. Sometime in the next century, or later, voyagers with far greater range and far sharper sensors than our own may find the orbiter’s remains and come to know exactly how it died. It's enough, now, to know just why it died.

To probe further

For more on the NASA Mars exploration program, turn to the agency's World Wide Web site at http//www.mars.jpl.nasa.gov/. For NASA's initial views on reasons for the loss of the Mars Climate Orbiter, issued 30 September, see www.mars.jpl.nasa.gov/msp98/orbiter/. More about the Mars Polar Lander is at mars.jpl.nasa.gov/msp98/lander/.

The Jet Propulsion Laboratory, Pasadena, Calif., maintains a picture archive for the Mars orbiter at www.jpl.nasa.gov/pictures/solar/mcoartist.html.

Non-NASA space news sources include three top sites: Lou Dobbs' new (and increasingly impressive) Space.Com home page, www.space.com; Keith Cowing's NAS A Watch gadfly site: http://www.reston.com/nasa/watch.html; and *Florida Today*'s online Space Today site: www. flatoday.com/space/today/index.htm.

Spectrum editor: Alfred Rosenblatt

IEEE Spectrum December 1999 Volume 36 Number 12