

Some Funny Things Happened on the Way to the Moon

Richard H. Battin

Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

In the Beginning

WITH THE TYPESETTING program $\text{T}_{\text{E}}\text{X}$ installed in my MAC II, I begin the von Kármán Lecture for the AIAA. The date is December 21, 1988, exactly 20 years to the day from the launch of Apollo 8. Eleven years earlier, the world had been enthralled by the Russian Sputnik and the Space Age had begun.

I have at my fingertips several orders of magnitude more computing power than the Apollo Guidance Computer which was carried onboard the Apollo spacecraft. And this marvel of modern technology sits on my desk at home!

This is the year for celebrating Apollo 11—the 20th anniversary of the first lunar landing. But the earlier flight of Apollo 8 was also a dramatic milestone. It was the first manned spaceflight beyond an earth orbit. The astronauts, Frank Borman, Jim Lovell, and Bill Anders, were the first human beings ever to see the entire earth as a ball. Said Norman Cousins

“On the first flight to the moon we really discovered the earth.”

Indeed, who can ever forget that picture of the earth rising above the lunar landscape?

To many of us who were part of the Apollo program, it was the most thrilling flight of all. We demonstrated the feasibility of onboard, self-contained space navigation for the very first time.

It all began on October 4, 1957 when the Russians launched that first satellite. I had started life at MIT—first as a student, then working with Hal Laning at the MIT Instrumentation Lab. But on that fateful day I had been away from MIT for a year working in industry. When I learned that Hal had a simulation of the solar system running on the IBM 650 and was “flying” round trips to Mars, I could hardly wait to rejoin him.

A report by Hal Laning, Elmer Frey, and Milt Trageser on the feasibility of a photographic reconnaissance flight to Mars had just been published at the laboratory. They were dead serious that such a flight was possible within the next five to seven years.

In those days, the contract between MIT and the Ballistic Missile Division of the US Air Force contained a clause which allowed our laboratory, within certain limits of course, to work on whatever struck its fancy—a sort of government IR&D program.

The few of us lucky enough to be involved were very excited about the Mars probe. We studied the problem intensely for a year or so and produced a three volume report together with a full-scale wooden model of the spacecraft.

Today, that model is displayed in the lobby of The Charles Stark Draper Laboratory.*

Navigation data for the Mars probe was to be gathered by an onboard sextant and processed by a spacecraft digital computer. Observation data would be used to determine vehicle position and a correction to the onboard clock. Periodically, changes in velocity would be made by a small propulsion system as directed by the computer. Spacecraft attitude was to be maintained by momentum wheels also under computer control. Power would be obtained from the sun using solar panels which would unfold like venetian blinds after launch.

To ensure reasonable launch velocities, the round-trip flight time was to be three years. Computer activity would be minimal during the long coasting periods between velocity corrections. Most of the time the computer would be “asleep” to conserve power. Indeed, the principal requirement for the computer was that it have a long shelf life.

Only one picture of the Martian surface was to be taken to eliminate the mechanical pitfalls of a film transport system. The spacecraft would make one pass by the planet a few thousand miles above the surface, orient itself for the photograph, open the shutter to expose the film plate, and coast back to earth.

The capsule housing the film had a shape similar to the Apollo Command Module and, in like manner, would dive through the atmosphere to splash down in the Gulf of Mexico. A radio transmitter, a yellow dye, and a flashing light would aid in its recovery. A repellent would hopefully discourage any shark who might think of it as dinner.



Richard H. Battin received an S.B. degree in electrical engineering in 1945 and a Ph.D. in applied mathematics in 1951—both from the Massachusetts Institute of Technology. He received an Honorary Doctor of Science Degree in 1999 from Texas A&M University. Currently, he is a Senior Lecturer in the MIT Aeronautics and Astronautics Department. He retired in 1987 from The Charles Stark Draper Laboratory, Inc. In 1972, he and David G. Hoag were presented by the AIAA with the Louis W. Hill Space Transportation Award (now called the Goddard Astronautics Award) “for leadership in the hardware and software design of the Apollo spacecraft primary control, guidance, and navigation system which first demonstrated the feasibility of onboard space navigation during the historic flight of Apollo 8.” He received the AIAA Mechanics and Control of Flight Award for 1978, the Institute of Navigation Superior Achievement Award for 1980, the AIAA Pendray Aerospace Literature Award for 1987, and the von Kármán Lectureship in Astronautics for 1989. He was presented by the American Astronautical Society with the 1996 Dirk Brouwer Award and the inaugural 2000 Tycho Brahe Award by the Institute of Navigation. For his latest book *An Introduction to the Mathematics and Methods of Astrodynamics, Revised Edition*, published in 1999 by the AIAA in their Education Series, he will receive the AIAA Summerfield Book Award for 2002. He is an Honorary Fellow of the AIAA and a Fellow of the American Astronautical Society. He is a member of the National Academy of Engineering and the International Academy of Astronautics. “In recognition of outstanding teaching” the students of the MIT Department of Aeronautics and Astronautics honored him in 1981 with their first Teaching Award.

Presented as the von Kármán Lecture at the AIAA 27th Aerospace Sciences Meeting, Reno, Nevada; Jan. 9–12, 1989, Paper 89-0861; received Jan. 9, 1989; revision received March 8, 1989. Copyright © 2001 by Richard H. Battin. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/02 \$10.00 in correspondence with the CCC.

*The Draper Lab and the Instrumentation Lab are the same. Only the name was changed in 1970 as a fitting tribute to its founder.

All in all, it was a rather sophisticated project considering the state-of-the-art. Milt Trageser was and is to be congratulated for his engineering prowess and ingenuity. The program was well conceived and carefully planned. But, as has happened to so many other “best laid plans,” it too went astray.

In the Doldrums

The Mars probe preliminary design was complete in the summer of 1959. The Air Force had been its sponsor, but a new government agency—the “National Aeronautics and Space Administration” would control its destiny.

With view graphs, reports, and the wooden spacecraft model, we arrived in Washington on the same day as Nikita Khrushchev. Our presentation was well received. But the high-level NASA audience we had anticipated, including Hugh Dryden, was busy entertaining the Russians.

NASA did not immediately write us a blank check for the Mars probe, but we were promised some future study money. Our small team survived but much of the original enthusiasm did not. Now we were simply doing “interplanetary navigation system studies.” We had no reason to anticipate what lay ahead.

To support the Mars project, we developed appropriate trajectories with flight times of roughly three years, and launch dates in 1962–1963 time frame. In this case, the spacecraft makes two orbits of the sun while the earth does three. Later we found round-trip missions to Venus having flight times of only a year and a quarter.

I also discovered on January 26, 1961 the first multiple flyby orbit—earth to Venus to Mars to earth—which is traversed, at least theoretically, without additional propulsion. This game of celestial billiards is played by proper control of the orientation of the orbital plane and altitude during the swingby of each planet.

Today, the Voyager spacecraft on its Grand Tour of the solar system is a spectacular demonstration of such missions. Soon, hopefully, the Galileo spacecraft will be making dozens of similar close encounter flybys of the Jovian moons.

The tools we needed for this work were not easily achieved. One did not study Celestial Mechanics unless he planned to be an astronomer and astronomers were not designing orbits for missions to Mars. On the contrary, in 1956 the British Astronomer Royal declared “Space travel is utter bilge!”

Some astronomers, whom we did consult, had reservations about the success of the project. “How do you expect to send a spacecraft to Mars when you don’t know exactly where it is?” (In those days the uncertainty was several thousand miles.) I suppose it was not easy to abandon their familiar earth-based reference coordinates.

Navigating the Mars probe consisted of measuring angles between planets and stars. We linearized those measurements about a reference point and used Gaussian weighted least squares to obtain the celestial fix.

The terms “estimator” and “state vector” were not in vogue so we couldn’t yet say that we had designed an “estimator” for a four-dimensional “state vector”—three for position and one for the onboard clock correction.

Later, for the Apollo system, the state space had nine-dimensions. In addition to position and velocity, we would also be estimating the rendezvous radar antenna angle biases on the Lunar Module and estimating lunar landmark locations as observed from the Command Module.

Gaussian least squares requires batch processing. All data is collected before the computation begins. In a small flight computer with data gathered over long periods of time, the method is cumbersome and impractical.

The situation was remedied by a recursive form of the estimator which allowed measurements to be incorporated as they are made.

It was not important for the Mars probe, but it was essential for Apollo navigation.

This recursive estimation process is now known as a Kalman filter. Today, it is widely used for all sorts of purposes. Every student of control systems studies the subject in school.

But it was used for the first time to navigate the Apollo 8 Command Module in cis-lunar space on its way to the moon!

With NASA funding, Hal Laning and Ramon Alonzo began, in earnest, the design of the Mars probe computer with its unique characteristics for space applications:

- variable speed to save power
- relatively few transistors
- parallel word transfer
- automatic counter incrementing
- automatic interrupt

The program and constants were wired in a so-called “core rope”—a memory with unusually high bit densities which could not be altered electronically.

This was the computer concept and architecture that would one day take man to the moon.

In the Race

Our NASA contract ended and there was a nine month hiatus before another six-month contract began in early 1961—this time for a preliminary design study of a guidance and navigation system sponsored by the NASA Space Task Group.

Later that year on May 25, 1961, President John F. Kennedy in his Special Message to Congress on Urgent National Needs said:

“I believe that this nation should commit itself to achieve the goal, before this decade is out, of landing a man on the Moon and returning him safely to Earth.”

Jim Webb, the NASA Administrator, knew Doc Draper and asked him to develop the Apollo guidance and navigation system. Of course Doc agreed.

“When will it be ready?” asked Webb.

“When you need it,” said Draper.

“How do I know it will work?” Webb persisted.

“I’ll go along and operate it for you.”

And he most certainly would have done so, had they only let him.

It all became official on August 10, 1961—exactly eleven weeks after Kennedy’s speech.

The first major Apollo contract awarded by the space agency was to the MIT Instrumentation Laboratory!

It sounds incredulous but that’s the way it really happened. Everyone who knew Doc Draper believes it and it is the story just the way Draper always told it. Doc is no longer with us but his spirit and inspiration live on. †

Today, many people wonder “Was there ever really a Space Race?”

It was certainly real in the beginning. There was a strong concern that the Russians would interfere with an Apollo flight by jamming the telemetry signals.

Our lab was noted for developing autonomous systems in missile guidance—the self-contained backup system for the Atlas intercontinental ballistic missile, the Thor IRBM, and the Polaris fleet ballistic missile guidance system. Of course, needless to say, the Mars probe would have been self contained.

The challenge of providing an autonomous guidance and navigation system for Apollo was right up our alley. So when Charlie Frick, our NASA boss from the Apollo Spacecraft Program Office, announced:

“There will be absolutely no ground communications with the Apollo spacecraft! Don’t even think about it!”

†On September 28, 1988 the Charles Stark Draper Prize was created by the National Academy of Engineering and funded by the Draper Lab Board of Directors. It is a major new international award to recognize achievement in engineering and technology.

The Prize is similar to the Nobel Prize. It will recognize extraordinary engineering accomplishment in the service of human welfare and freedom, and will emphasize those aspects of engineering that are essential to a better future.

Doc would certainly have been pleased.

it was difficult to suppress our smiles. We felt just like Br'er Rabbit:

“Please don’t throw us in that briar patch, Br’er Frick!”

Security was also very real. Almost everything was classified—even schedules. Fortunately, that particular phobia soon petered out.

At the other extreme, there were many who advocated cooperation with the Russians. During a panel discussion at an AIAA conference, Wernher von Braun addressed the question: “Why don’t we work together on the Apollo program?” His response went right to the heart of the matter:

“If there were cooperation with the Russians on space-flight, there wouldn’t be a program in either country.”

To start our part of the race we had first to assemble a team. But finding the right people was frustrating. We had the most challenging guidance system imaginable to develop. Recruits should have been pounding at the door. But that didn’t happen.

In the long run, though, the best people were already at the Lab—working in other divisions. Appropriate transfers were made. Nevertheless, for a very long time to come, the software task was not adequately staffed.

We were beginning to learn that “people problems” were often much more difficult to solve than the technical ones.

Although we couldn’t seem to get engineers, we did get advice—both technical and theological:

“Apollo should be launched on a clear night when the moon is full to provide the best possible target.”

and

“If the Lord had intended man to go to the moon, He would never have created Senator Proxmire.”

In the Software Jungle

The MIT Instrumentation Lab was designing software for the Apollo onboard guidance system even before the word “software” was invented.

I still remember the first time I told my wife that I was in charge of “Apollo Software.” She exhorted me: “Please don’t tell any of our friends!”

I suppose real men do “Hardware” just as real men don’t eat quiche.

It was an attitude that prevailed a long time in many organizations. Salaries for computer programmers did not keep up with the salaries of engineers. Engineers did engineering. The programming (or coding) was more menial work and should be left to others.

I wanted no such distinction. Our best engineers should design *and* program the software for the flight computers. The reasoning was simple. A good engineer can learn to write programs. But a computer software specialist would find it far more difficult to do the engineering without considerably more training.

The basic architecture of the Apollo onboard computers was the design of Hal Laning—certainly one of our most creative engineers, and the early application programs were written by our best system engineers. We tried hard to keep our standards high.

Even so, with the best talent, computer programs seldom perform as intended the first time. If they work at all, they may do unexpected or “funny” things. We called them “FLT’s”—short for *Funny Little Things*.

This was before “Murphy” discovered his famous law. He must have been secretly watching our software development efforts.

The Apollo Guidance Computer, or AGC as it was called, evolved from the design for the Mars probe. In 1961, when our Apollo contract was signed, it had 4096 words of fixed memory and 256 words of erasable.

A word was 16 bits (that’s *bits* not *bytes*) with one bit for sign and one for parity check. Hence, double precision was required for most calculations. The cycle time was modest, approximately 24 μ sec, which was more than sufficient for the Mars application.

The fixed memory was called a “core rope memory” since the early models clearly resembled lengths of rope. The high density of storage was achieved by “storing” a large number of bits in *each*

magnetic core. A stored bit is a “one” whenever a sense wire threads a core and a “zero” when it fails to thread a core.

In the operation of the rope memory, a core is switched which induces a voltage drop in every sense line which threads that core. Sixteen sense wires were connected to sense amplifiers to detect which had voltage drops. With the addition of an appropriate switching network, each core could then hold several words.

All this information was permanently wired at Raytheon by “LOL’s”—literally, “little old ladies” who slowly and painstakingly threaded the cores by hand. Later an ingenious adaptation of a textile loom was used for this purpose which was far faster and certainly more reliable. The loom was driven by a punched paper tape created by the same program that produced the mission software. Once a memory module was manufactured, not a single bit could be changed, either intentionally or unintentionally.

There was a dichotomy of opinion regarding the fixed memory concept. Once the memory had been programmed, it took about six weeks to manufacture. Add to that the time required for testing and you discover the rule:

No flight-computer memory changes within two months of a launch!!

It was far too late to change the design so the argument was purely philosophical: “Was it bad or was it good?”

Each side made its own case:

- There must be a way to make last minute changes.
- But, last minute changes may be ill-considered and/or hurriedly tested. They are dangerous.

In the long run, despite objections, it was good discipline and kept everyone honest. If you know you can’t make last minute changes, then, by golly, you’ll be that much more careful in your design and testing.

Inevitably, of course, someone would discover that erasable memory could hold more than just data. You could actually write and execute programs from the erasable memory.

By that time, though, strict NASA approval was required to do so.

The physical size of the computer was one cubic foot. That could be changed only at great cost since the spacecraft manufacturer, North American Aviation, had allowed that much room and no more in the Command Module. If we needed more capability, it would have to come from advances in computer technology.

Eldon Hall, who had designed the Polaris missile computer, was in charge of the Apollo Guidance Computer development. I can still remember the day he asked if I could use twice as much memory. (He had figured out how to stuff twice as many sense wires through a core.) I was ecstatic. Our prayers had been answered.

Charlie Frick, however, was not so pleased:

“You told us that 4000 words was enough. And now you want to double it!!”

He was, obviously, a man of experience in dealing with unscrupulous contractors. Doubling the size would surely mean doubling the cost.

The memory size was, indeed, doubled. And it doubled twice again after that before it went to the moon. The final count was 36,864 sixteen-bit words for the fixed memory and 2048 for the erasable. The cycle time had also been cut in half to 12 μ sec. But the physical size never did change.

We were able to cope with limited memory, limited instruction repertoire, and short word length for mission programs by using a powerful interpretive language.

Charlie Muntz created the “Interpreter” to manipulate the 28 bit data words. Memory was conserved (dramatically, to be sure) but now the time for numerical computations was measured in milliseconds:

- Double-precision add—0.66 millisecond
- Double-precision multiply—1.1 millisecond
- Double-precision square root—1.9 millisecond
- Double-precision sine—5.6 millisecond

Although the data words now had sufficient length, they were expressed with a fixed decimal point. We couldn't afford the luxury of a floating-point arithmetic. Fixed point is much faster and many of the mission programs had to function in real time.

The Master Control of this software maze was the Executive and Waitlist Program skillfully designed and written by Hal Laning. It was a sophisticated piece of software that:

- permitted time-sharing of erasable memory
- allowed orderly interruption of programs by those of higher priority
- accommodated as many as seven programs in suspended animation with sufficient information saved to enable each to resume at a later time as though nothing had happened.

It was a supreme triumph of Hal Laning's ingenuity and perspicacity. (Remember that this was 1961 and control computers were barely in their infancy.)

I can still remember when Hal first tried to explain to me just how all this was supposed to work. It was mind boggling. "Hal," I said, "This is much too complicated. There has to be a simpler way." He responded with profound authority:

"There is no other way to do this job!"

And, of course, he was right.

But the same mind which could conceive this logical masterpiece preferred not to cope with endless meetings, in large conference rooms, with too many people, and too much contention. This was to be our lot for years to come and Hal avoided it like the plague. Fortunately for us, he was always there for advice and counsel.

Mission programming began in earnest as soon as the software tools were in place. We had no firm requirements from NASA but we knew, generally, some of the things that had to be done.

We would have to navigate with onboard sensors; to make course changes outside the atmosphere; to reenter the atmosphere at the proper angle; and we would have to guide the Command Module safely to its splash-down site in the ocean.

These activities started before the decision favoring Lunar Orbit Rendezvous and the invention of the Lunar Module. In fact, the Grumman Aircraft Engineering Corporation wasn't even identified as the LM contractor until the end of 1962.

In June 1964, Robert C. (Cliff) Duncan, Chief of Guidance and Control in Houston, directed that spacecraft autopilot functions be performed digitally in our computer. It was rather late for such a big change. By that time, we had been on the job for three years and a lot of code had been written.

To accommodate these new and time-critical functions, he authorized a design change in the computer—we could double the speed and increase the repertoire of codes. (At that time the computer had but eight basic instructions and "divide" was not one of them.) But the software people were not to worry. Their carefully-crafted programs would still work on the new computer. "Upward Compatibility" was to be the name of the game.

But the new computer could be programmed much more efficiently with the new codes and our precious computer memory must never be squandered. Everything was redone for the Block II system.

Unfortunately, the Block I system didn't go away. It would still be used for the first orbital flights. Now we had two different software systems to maintain.

That wasn't the worst either. Soon we would have both the LM computer as well as the CM computer for which to provide software. The two were the same but they had radically different tasks to perform.

The Apollo software job was escalating rapidly. The question was: "Were we really up to it?"

In the Trenches

The first manned Mercury flight was May 5, 1961 and the last on May 15, 1963. The Gemini program, which was to proof-test the concept of orbital rendezvous, had its first flight on March 23, 1965—the last took place November 11, 1966.

During this period, NASA was totally focused on the success of those missions. The next flight always has everyone's undivided attention.

At one time during those Gemini flights, a software change was required in their reentry guidance program. IBM, the software contractor, told NASA the cost of the change would be one million dollars!

George Mueller was appalled. As the new Associate Administrator for Manned Space Flight he wondered "If it was that costly to change Gemini software, how about Apollo? Who was doing the Apollo software anyhow?"

It was our first indication that top NASA management had ever thought about the MIT software effort. Until then, we were virtually unimpeded by outside influences. It was pure heaven! So many things were decided by one or two engineers which, today, would take many trade-off studies and large committees.

When the spotlight finally fell on us, the conceptual part of the job and many of the mission related programs were essentially finished.

Later, I reminded George Mueller of the incident and asked "How much did it really cost to change the Gemini reentry program?" His answer: "One million dollars."

Soon we were made keenly aware of a basic fact of life: There is a vast difference between getting the Apollo software job and keeping it.

We frequently heard the opinion that MIT should not do the Apollo software. Production software is not the kind of job for a school. They don't have the resources, the right people, or the motivation. (I guess nonprofits can't be rewarded or punished—at least not in the traditional ways.)

The NASA creed was evident everywhere—on desktops, on posters, and bumper stickers:

Better is the Enemy of Good

But, "MIT doesn't know when to quit designing," they said. "They have a bunch of prima donnas who want to make everything perfect regardless of how long it takes."

We seemed to be in a kind of trench warfare—constantly defending our job against all assailants.

An amusing incident took place at a meeting between our MIT software people and an industry group who was, obviously, probing to find our weak spots. I presided at the conference table and was absentmindedly toying with a pair of scissors when an annoying fly flew past.

Suddenly, with one slash of the scissors, I cut that fly in half in midair. The visitors swallowed hard and the tone of the meeting immediately changed. Once again, the day was saved.

George Mueller took a personal interest in the Apollo Guidance Computer and it was no wonder that he did. He had recently experienced the dire consequences of a programming error.

A missing "hyphen" in a flight computer program caused the loss of the Mariner I at the start of its mission to Venus on July 21, 1962. Although I never saw it, I understand that he had the symbol framed



and prominently displayed on his office wall lest anyone forget the importance of flight software.

Early in his tenure at NASA, he formed the Apollo Software Task Force with representatives from many of the Apollo contractors. Fortunately, I was a member. When George saw computer memory disappearing at a rapid clip, he solicited informal proposals from Task Force members to program the Apollo computers using just half the memory.

Inevitably, the "half-memory computer" would be privately labeled the "half-assed computer." At least, we hadn't lost our senses of humor.

But now we were really nervous. The door was being opened to compete the Apollo software job and Laboratory policy did not permit us to compete with industry.

We were confident that the Apollo software requirements could not be met with only half the memory. But we were also certain that there were those who would say that it could be done. Bellcom and IBM were among those interested in bidding.

The time soon came for proposal presentations from the various bidders. When it was Bellcom's turn, Gordon Heffron, the Bellcom representative, said that MIT was doing an outstanding job and they would not submit a proposal.

It all ended right there and then. Our job was secure—or so it seemed.

The biggest complaint NASA had about our software effort was quite straightforward:

“MIT doesn't have enough people!”

Some tasks were obviously one man jobs. We used those examples in feeble attempts to fend off the criticisms. We loved the analogy:

“How many people do you have to squeeze into a booth to make a phone call?”

The Michelangelo analogy was another favorite:

“When he was painting the Sistine Chapel ceiling, would Michelangelo have accepted help from a gang of house painters just to finish on schedule?”

Our real problem was paradoxical:

1. Good people are hard to find.
2. If you find them, they must be trained.
3. The people who must train them are too busy.
4. Why are they so busy?
5. Because we don't have enough people.

For the first Command Module flight, Apollo 3, Alex Kosmala was the official “Rope Mother.” With a small dedicated group, he spent 15 months preparing the flight program, called CORONA.

The original estimate was 6 months. But until then, no one had any idea just how difficult and time-consuming the job would be.

Although we were 9 months late, NASA was even later in implementing the Real Time Control Center. CORONA was released in January of 1966 but didn't fly until August.

We were saved this time, but the specter loomed that one day the Saturn V would be perched on its launch pad carrying the Apollo spacecraft and waiting for those MIT guys to deliver the flight program.

We could well become the “long pole in the tent”—a colorful label no one wanted pinned to *his* back.

The people problem was solved when the System Development Corporation became our subcontractor. They supplied organized teams with competent leadership. We didn't have to worry about how to employ 30 or 40 individuals, but could just deal with a few team leaders. Other companies, Raytheon, AC Electronics, ARCON, and CDC also supplied talent but not in such large numbers.

Flight Operations in Houston under Bill Tindall took charge of the MIT software contract. His first action was to deal with computer memory problems. It was Friday the 13th of May, 1966—“Black Friday” for many of our favorite programs.

Since the Russians probably wouldn't try to mess us up, the flight would be controlled from the ground. We could just as well strip out all those lovely algorithms which made the spacecraft self-sufficient.

But wait! Most of those programs were needed for another reason. We could not, for instance, dispense with “Navigation” or “Return to Earth” or “Powered Flight Guidance.” After all, we still could lose contact with the ground for non-sinister reasons.

After that, when new and absolutely essential requirements surfaced, they could only be added when something else of lesser importance was removed.

Things had gotten that serious.

Then came that terrible day of the fire and the loss of the crew: Gus Grissom, Ed White, and Roger Chaffee. It was January 27, 1967 during ground testing of the Apollo spacecraft in Florida.

The launch schedule was now a shambles. The spacecraft was completely redesigned under the superb leadership of Aaron Cohen. Three unmanned guided flights, which had not been planned before the fire, were flown and supported by our software teams before the first manned flight, Apollo 7, on October 11, 1968. Then we really went into high gear!

I still can't believe how fast events happened after that 20 month hiatus:

- Apollo 8 on December 21, 1968—To the moon
- Apollo 9 on March 3, 1969—Flight-test the LM
- Apollo 10 on May 18, 1969—The dress rehearsal
- Apollo 11 on July 16, 1969—The lunar landing

NASA was shooting them off like Roman candles!

The rest is history. We never held up a flight; we never had a system failure; and I am proud to say:

“The MIT Instrumentation Laboratory did all of the on-board software from the first Apollo guided flight on August 25, 1966 through the three SkyLab missions in 1973 to the Apollo-Soyuz rendezvous mission with the Russians on July 15, 1975.”

In Route

Certainly the longest and most thrilling 5 minutes of my life was the 5 minute burn of the S-IVB engine to boost the speed of the Apollo 8 spacecraft to the 24,200 mph necessary to escape the earth.

“You are on your way,” said Chris Kraft, from the Mission Control room, “you are *really* on your way.”

A few weeks before the launch, the Navigator Command Module Pilot, Jim Lovell, spent a few hours practicing on the earth-horizon sextant simulator at MIT. He consistently identified the “horizon” about 20 miles above the real horizon.

Great! Jim Lovell could be calibrated and his bias number loaded in the flight computer.

He was recalibrated in real time on the way to the moon. His first eleven sextant angle measurements, made early in the flight, were compared with what they should have been according to the RTCC.

After the horizon calibration, came the first midcourse correction of almost 25 ft/sec. It was a fairly large one due to earlier maneuvers to get the spacecraft safely away from the third stage of the launch vehicle.

Following the course change, the onboard computer's version of the state vector was made to agree with the value obtained from ground tracking.

From that time on, Jim Lovell made dozens of star-elevation measurements using both the earth and lunar horizons. These were processed by the Apollo Guidance Computer using the recursive estimation algorithm. The final set of 15 sightings was made about 35,000 miles from the moon.

The onboard and ground tracking estimates were almost identical! Chris Kraft even suggested that we use the onboard state vector for lunar orbit insertion. But the flight plan dictated that a state vector update from the ground was required before any maneuver. And there was no overriding argument to deviate from the plan.

Nevertheless, the evidence was conclusive.

The astronauts could have done it on their own without any ground assistance whatsoever!

Early in the morning of December 24, Apollo 8 disappeared behind the Moon. For 34 minutes there was no way of knowing what had

happened. During that time a 247-second burn took place under the control of the MIT guidance system and Apollo 8 was in lunar orbit. The astronauts announced their orbital parameters provided by the AGC when voice contact was resumed. The Mission Control folks were obviously disconcerted:

“How do they know? We haven’t had time to track them yet.”

They were not yet experienced with self-contained inertial systems.

At 8:40 pm Christmas Eve, 1968, the Apollo 8 astronauts were on television broadcasting from lunar orbit. “For all the people on Earth,” said Bill Anders, “the crew of Apollo 8 has a message we would like to send you.” He paused a moment and then began reading:

*“In the beginning God created the Heaven and the Earth
... and God saw that it was good.”*

The commander Frank Borman added:

*“And from the crew of Apollo 8, we close with good night,
good luck, a Merry Christmas, and God bless all of you—
all of you on the good Earth.”*

Days later, in the Washington Post, there appeared an editorial:

*“At some point in the history of the world someone may
have read the first ten verses of the Book of Genesis
under conditions that gave them greater meaning than
they had on Christmas Eve. But it seems unlikely . . . This
Christmas will always be remembered as the lunar one.”*

In Retrospect

The euphoria following the success of Apollo 8 was quickly overshadowed by the departure of most of our top software talent. One of the many “spinoffs” of the Instrumentation Lab was happening again. Like so many others before them, they were starting their own company. It was a hard blow indeed.

But we did survive. By that time we had very strong team members and they were ready for the promotions to follow.

Apollo 11 was a magnificent triumph! But the actual landing was more exciting than we had counted on. Just before touch down the Apollo Guidance Computer almost caused a near panic by displaying alarms. The NASA flight controllers remained cool and did not order an abort.

During the short stay of Neil Armstrong and Buzz Aldrin on the lunar surface, everyone who could possibly contribute to an understanding of the problem was hard at work. There was concern about the upcoming liftoff and rendezvous in just a few hours.

The culprit was an erroneous mode setting—the rendezvous radar was transmitting pulses to the computer at maximum rate. During landing, the computer has plenty of work to do. The additional task of counting all those extraneous pulses was just too much.

There are still a lot of people, who should know better, who describe that event as a computer malfunction. Nothing could be farther from the truth. It was, in fact, a triumph of software design.

When the computer was operating near capacity and sensed an overload condition, it was programmed to stop everything, clear the table, and restart only the top priority jobs. The display of the alarm was to tell the astronauts and ground controllers just what was happening.

The software designers had the foresight to anticipate such a possibility and the skill to cope with it should it ever happen. No one ever thought it would happen—except Murphy, of course.

Neil Armstrong remembered the Apollo computer when he was later assigned to NASA headquarters to head the research and development programs. One such program involved an F8 airplane to be used for a fly-by-wire experiment. NASA needed a reliable computer to fly that airplane. A computer failure would mean the loss of the plane.

Neil suggested the Apollo Guidance Computer:

“It’s the most reliable computer I know. It got me safely to the moon and back.”

And that is exactly the computer that was used. Phil Felleman, an outstanding leader and systems engineer at the Draper Lab, was given the job to do. And he did it well.

The Apollo Guidance Computer never failed in any flight. There was concern when lightning struck Apollo 12 during launch. It wiped out the erasable memory load and initiated a computer restart. But the memory was reloaded in earth orbit and all went well.

The quality-control people couldn’t calculate the “mean-time-between-failures” since it would have required dividing by zero.

Generally, engineers stare in disbelief when told that the Apollo computer had absolutely no built-in redundancy. There were many possibilities of single-point failures which would have disabled the computer. But none of them happened.

Even the computer itself was not redundant. In the Space Shuttle today there are five guidance computers but there was only one each in the Command Module and the Lunar Module. (There was a back-up computer of limited capability in the LM for emergency use, designed and built by TRW, but it was never needed.)

It wasn’t a matter of luck either. Reliability considerations were uppermost in the design right from the start.

For example, only one type of integrated logic circuit, a three-input NOR gate, was used. This meant that Eldon Hall and Raytheon, the computer manufacturer, could concentrate on quality control of a single circuit rather than several. A wider variety of circuits would have reduced the number of components per computer but reliability would have been adversely affected.

In the wake of the Apollo 11 flight, Dave Hoag, the MIT Program Manager for Apollo, and I were invited to Russia as guests of the USSR Academy of Sciences. We took with us the NASA film of the lunar landing.

It was an immediate success. No matter where we went, everyone wanted to see the film. If there was no screen, they would hang up a bed sheet. When the lights went out, cleaning ladies, janitors, and other extraneous folk would squeeze into the room. It was shown so many times on so many strange projectors that the sprocket holes were damaged beyond repair.

During the question and answer period following the show, there was always the question:

“When will the Americans go to the moon again?”

Fortunately, we knew the planned date of Apollo 12 and told them. They were incredulous. No one ever announced launches ahead of time.

As it happened, we were dead right. Apollo 12 did fly on schedule November 14, 1969. That must have done a lot for our credibility back in the USSR.

The Soviet penchant for secrecy has always baffled me. Just before we flew to Russia, the Soyuz 6 with two cosmonauts aboard was launched on October 11, 1969. When we changed planes in London, we learned of the launch of Soyuz 7 on October 12. “Wonderful,” we thought. “We will be there in person and learn all about this first hand.”

As we drove toward Moscow from the airport with our host, I asked about the purpose of the double launch. The reply was an irrelevant comment about some housing project we had just passed.

The next day, October 13, Soyuz 8 lifted off to join the other two. Now there were three spacecraft and six cosmonauts in orbit—a real spectacular! We knew of this only from pictures in *Pravda*. None of us read Russian so we were at the mercy of our Russian host for details. But, to our dismay, we again learned nothing. Any question was countered by a quick change of the subject.

Each of the three missions lasted about five days. When all of the cosmonauts had safely returned, we speculated that there would be a big celebration and parade in Moscow. “Oh no,” was the reply. “Spaceflight is now so routine.”

At the time, we were in Tbilisi, Georgia where we stayed an extra day because of, alleged, bad weather—no planes to Moscow.

When we finally did arrive in Moscow, we were met by a new host. “You should have been here yesterday,” he said. “There was a big parade in Red Square and we had seats for you with Brezhnev in the reviewing stand.”

When we returned from the Soviet Union, Howard Johnson, then president of MIT, told us a marvelous story which always brings tears to my eyes.

It seems that a friend of Howard’s was between flights in London after the lunar landing. To pass the time, he visited Westminster Abbey and came upon the tomb of Sir Isaac Newton. (If you haven’t seen it, you should. It is an impressive memorial.) Someone had left a note on the tomb. It read:

Sir Isaac—the Eagle has landed!

AMERICA'S PROGRAM FOR ORBITAL AND LUNAR LANDING OPERATIONS

Apollo 11 Splashdown Party
MIT Faculty Club
July 24, 1969



STANDING FROM RIGHT TO LEFT: Fred Martin, Eldon Hall, Gerry Levine, Dick Battin, Tom Fitzgibbon, and George Schmidt.