

Today we have an old colleague and friend of mine.

Tony Lavoie actually got his engineering degree in the Aero-Astro Department here at MIT, so this is coming back home for him.

And he actually grew up in Massachusetts.

But, for the last 23 years, he has been in Huntsville, Alabama at the Marshall Space Flight Center and actually has been an engineer on quite a few of the projects that I have flown with, including the Astro Observatory and the Tethered Satellite.

And Tony has continued to rise up in the ranks of the NASA engineering community.

He is now, I guess, the Project Manager for the Robotic Lunar Exploration Program.

If you have been following the space news, Marshall Space Flight Center is going to be developing the first robotic lunar lander that we have ever sent to the Moon since Surveyor.

It will be like 40 years, I guess, we haven't done that.

And we are going to have to figure out how to do it.

Tony was also the Chief Engineer on the Chandra X-Ray Observatory which was launched on the Shuttle.

And so, I thought it would be interesting to hear, since we've been looking at many other aspects of Space Shuttle operations, to hear something about what it was like getting a payload ready to fly on the Shuttle.

But also we were discussing, before class began, that he may have some comments on the Robotic Lunar Exploration Program, which is now in the pre-phase A studies.

And it will sort of be interesting.

We talked a lot about what it was like in the pre-phase A and the early days of coming up with the requirements for the Shuttle.

And clearly, in the systems engineering, getting the requirements right is one of the key goals for success.

Tony may have some comments on that.

I have talked enough.

Tony, you've got it.

Hi there.

We are going to start talking about Chandra.

And, as Jeff pointed out, I am also going to talk about my new challenge, my new assignment as Project Manager and in pre-phase A what that means.

And I think it is a lot different, being a student, where you have a single closed form solution and you go and you work a problem and it is done.

In the real world you seldom, if ever, work a problem one time and have it finish.

And we will get into that as we go along.

I am going to talk about Chandra X-Ray Observatory.

And these are just a few images that Chandra has produced.

It is an outstanding x-ray imager, and it has been flying since July of '99.

And you are going to get all the gory details and all the lessons learned that we had in working that project.

I will do an overview Chandra history and then the challenges we faced and lessons learned associated with that, and we will cover those topics.

In terms of how we operate, feel free to ask questions as we go along.

It is kind of lose and easy.

Chandra, or back then it was AXAF, Advanced X-Ray Astrophysics Facility.

NASA, by the way, loves acronyms.

I mean, if you are interested in working with NASA or associated with NASA, they are just acronym crazy.

AXAF was one of the four great observatories.

The four great observatories were meant to cover as close to the entire spectrum as NASA could.

And so, the visible great observatory is Hubble.

And Ultraviolet, that's true, little Ultraviolet.

The X-Ray Observatory is AXAF, which was later named Chandra.

The Gamma Ray Observatory is the Compton Gamma Ray Observatory.

And the last one was Infrared.

And that originally started off as SIRTf.

And now it is Spitzer.

And that was launched a few years ago.

In each of the four great observatories, they were higher class than most of the other science missions, even the science telescopes.

And NASA did invest quite a bit of money on each of those.

And so, you will see that the performance from the four great observatories are, in terms of quality, probably the best in the world and the best NASA has ever done.

For the x-ray region, Chandra or AXAF was built.

And its objective, you start with program objectives and you work down with science objectives.

The science objectives were really to understand the nature of the universe and determine the nature of celestial objects, in particular those that are hot.

And hot objects will produce x-rays.

And, of course, Jeff knows all about that.

If you have any questions related to that you can ask Jeff.

But, from an engineer perspective, what happens is you are given a set of objects and maybe a program objective, and now you have to craft the mission and craft the project from those.

And one of the problems that you face is that sometimes they are not crisp and sometimes they are flexible based on cost.

And also sometimes they are not very specific.

Like determine the nature of celestial objects from stars to quasars.

Now what?

But this is the kind of thing that typically you start with.

And you have to iterate with, in this case, the science community.

If you are building part of the Shuttle, you have to iterate with, for instance, the users, the astronauts, the payload developers that would use the Shuttle, the operators, et cetera.

But the process is the same.

You start with the top level objectives and you work your way down.

Now, again, because of the nature of what we are talking about, it is usually a very iterative process that can last for several months.

And even several years if there is a technical challenge that causes you to stretch out.

And I will mention when it started.

In terms of AXAF, kind of the key parameters are the percent in circled energy, i.e., how sharp is the image?

Registration meaning where is the target in the sky?

And you have got to remember that x-ray astronomy is a relatively new branch of astronomy.

Because of the atmosphere it absorbs x-rays.

So we, humans, only discovered x-ray astronomy in the last 40 years.

We need to go outside the atmosphere to get information on x-ray astronomy.

One of the things that we did know is if there are some x-ray sources there they may or may not be the same sources that also emit invisible light.

And, if they are not, then you have a question of you kind of know there is a source out there, but you don't know where in the sky it is.

And so, one of the challenges that AXAF had was even if I have no stars around this x-ray source, I have got to be

able to pinpoint where that x-ray source is.

That was a particular challenge that we had, trying to figure out how to do that.

And that is what registration means.

And then effective area is related to how many photons you can collect to make a good image.

And recognize that if you are looking at the sun, the sun is very close so you can get a lot of photons, but if you are looking at something that is ten billion light-years away you are not collecting very many photons so you have to wait a long time to get those photons, unless you have a big collecting area.

Those are kind of the key performance requirements that you have to address when you are building an x-ray telescope.

From those there are key derived requirements that have to be derived basically to be able to maximize the scientific performance requirements.

And so, from these pieces of information, you derive mirror size and design, you derive focal length.

Pointing and control requirements, how accurately do you need to point?

Thermal stability because that plays a role in orbit, a very big role.

Instrument sensitivity and the fiducial transfer system, which I will talk about later, which is the thing that allows you to do the registration, even the absence of local visible stars around it.

Once you have the bulk of the science performance requirements and the derived requirements from these performance requirements then you fill in kind of the rest of the pie, if you will.

And those can be safety requirements in the case of a manned program.

In fact, that is pretty significant.

And it is a pretty significant cost driver for things that fly in space.

Unfortunately, they also cost a lot to make sure they are not a hazard to the crew that flies.

There are also design and construction standards, which also turns out to be a big deal.

For instance, if you are building a satellite and you have circuit cards, you are going to have requirements on

soldering, you are going to have requirements on integrated circuits, you are going to have requirements on glazing, you are going to have requirements on bonding.

You are going to have requirements on how to do things across the whole gamut of the spacecraft. Not just the spacecraft but also the Shuttle and anything that flies up there.

For a given mission, there are probably on the order of 100, 150 separate documents that describe various things, design and construction standards that you have to go through for a particular piece of building the vehicle.

And so, those things are always a challenge for a project manager because the engineers want to get the latest standard that applies, and in excruciating detail as engineers love to do.

And yet, from the project manager's standpoint, we are trying to maintain cost.

There is always a tradeoff between how good is good enough?

From a program manager's perspective better is the enemy of good.

Yet, from an engineer, good is always better, or better is always better.

There is always a healthy tension between the project manager and the engineers.

And I would say that a successful project knows how to tradeoff.

The project manager and the engineers know how to balance that tension to be able to get the best product.

And we have done that before, and we have also imbalanced it and had some spectacular failures.

Once requirements are set then you can start kind of designing and starting to put the drawings on developing the system, if you will.

But that is kind of the general flow of information for a science payload or a science mission.

You start with the top level requirements for performance, the scientific performance that you are looking for.

From those you derive the direct performance requirements, like what you see here, and then everything else cascades below that.

And, once you have the requirements set, then you can start developing concepts.

Now, this was a preliminary design of AXAF.

I am not going to talk about it, but I am just going to show the next slide as to what it looks like now.

And so, you can see that there is a major difference between what it started out as and what it ended up as.

And this is typical.

It is not atypical.

And the reason is cost.

A lot of times when you are first starting out with requirements, you kind of don't know what the cost of the mission is going to be.

And so, you come up with an idea that has a lot of capability.

This was orbit serviceable, it was in low earth orbit, and it had four focal plane instruments that you could select.

And it would actually move, much like Hubble.

And it was a 15 year mission with reserivicing with the shuttle astronauts.

But, when you start putting that on paper and you start crafting the requirements, one of the things you do, in the early phase of a program, is you also cost, or you try and cost what those requirements result in, in terms of design.

Even though you don't do the final design or you don't start working on the design in detail until after you've got the requirements set, in practical terms you have to still put together a concept design so that you can cost it.

And that is exactly what happened on AXAF.

We put together a design in parallel with working the requirements and we costed the design.

And it turned out that that was more money than NASA could afford.

And so, as is typical, the iteration process begins with headquarters to change the parameters of the mission.

To, in a sense, compromise the scientific objectives a little bit in terms of implementing a requirement set.

And the result is finally you get something that you can pay for and that still meets the intent of the scientific objectives.

That is a key lesson to learn.

When you are first starting out, when you are first crafting what you want to do, expect it to change.

And the key there is it is usually cost driven as to what you can afford.

Not only that, NASA, just by the nature of what it is doing, since it has never built an AXAF before, it is really hard to get a good, accurate cost of what it costs.

So, a lot of times, the cost is derived from a parametric usually weight-based system such that, for instance, if the mirrors or if the solar rays, you want them this big to provide this much power.

That big means they have to weigh or they mass so many kilograms.

And so that is about X amount of dollars based on weight.

Now, there are a lot of additional factors that you put in like complexity, interfaces, et cetera, technology readiness, whether you are looking for state-of-the-art or something that has already flown.

But generally it is what is called a parametric cost model.

And it is not based on taking a look at all the design drawings and saying this cost this much from the vendor, this cost this much, and you put it all together.

That is called a bottoms-up assessment.

And it is more accurate, but it also relies on having an accurate design picture which obviously you don't have early on in a program.

I mean it is almost like a black art to try and take a requirement set and craft a design early on or at least bound the design and bound the cost.

Because, again, early on what you are trying to do is get into the cost envelope and still meet the scientific objectives or the objectives that you are trying to do.

Yes.

In this final design, how are its capabilities versus the original?

I mean because you didn't put as much money into it.

Correct.

To make it on-orbit replaceable means that all of the boxes that you see in the spacecraft bus, or a lot of them would have had to have been, on this picture, replaceable.

They were all replaceable.

You have to spend some money making sure that the drawers slide in, slide out, there are no sharp edges and it is accessible, et cetera.

So, packaging is important.

That is one area where you can reduce cost and not really have an effect on performance, per se.

However, one of the big areas that we did take a little bit of a hit on performances is on the mirrors.

Now, x-rays, you cannot just have a regular shaped lens.

What we use is grazing incidence mirrors.

And I will explain this later.

Because the x-rays are so highly energetic, you have to graze them gradually to a focus.

And the original concept had six mirrors nested.

They were basically hyperbola parabola shaped.

The final concept had four.

We had to remove, not the inside or the outside, but two of the middle mirrors.

So, the mirror set is four.

This orbit that this flies in is a highly elliptical orbit.

It is about 140,000 kilometers by 10,000 kilometers.

The original orbit was in low earth orbit which is about 350 kilometers.

And so that affects the viewing time.

But, in all of the other registration and encircled energy, I think that we still were able to meet the original

objectives of the scientific mission.

There are ways to compromise that don't really affect directly what the scientists want to do, but it is a tradeoff.

It is usually not their first desire but they signed up to it, they are comfortable with the final answer.

Explaining a little bit of how it is put together.

This is a rather simplified version.

Here are the solar rays.

That is where we get our power.

This is the spacecraft bus.

This is the HRMA which is the mirrors, the nested cylinders that I was talking about earlier.

This is a low energy grading and a high energy grading.

For spectroscopy, what you want to do is you want to bend all of the photons in a particular energy much like a prism does for visible light.

And the way you do that is with these facets.

There are probably a thousand facets on this.

And I have a picture of that later.

And you flip it into the beam when you want to get some spectroscopy data.

And you remove the grading when you want to get an x-ray image of what you are looking at.

And that is a function of what the scientists want to do at a particular target.

Some scientists want to get an image because that conveys more information for them than spectroscopy.

Others want to see a spectrogram of the image because that tells them what the photons are and what the energy of the photons are.

And that can tell you, for instance, what elements are present and what temperatures for a given gas cloud or neutron star or what have you.

And then we have an optical bench.

And all that means is it is a piece of structure that is very stiff.

Because obviously you can tell if you flex a lot that sure doesn't help your camera imaging the performance.

It has got to be very stiff, and that is why they call it an optical bench.

And the ISM is the Integrated Science Module.

That is where the instruments sit in Chandra.

And here is a picture, as I mentioned before, of how the x-rays graze off of the cylinders in the HRMA down to a focal point about ten meters away.

Jeff had mentioned that sometimes it takes quite a while for science to come to fruition in a mission.

And this is probably a good case.

It started in 1978 with some concepts.

We did get approval for new start in 1988, so you can tell there was ten years time where this was just an idea in a scientist's mind.

And typically the course of events is that the scientist lobbies and submits papers, submits proposals to NASA, lobbies Congress, lobbies the National Science Foundation to say this is a good idea, we should do this and here is the concept.

And, of course, that concept changes over time as you get new technology, et cetera.

And so, that is exactly what happened here.

We did get approval ten years from when the original concept started.

Authority to proceed for the prime contract was in January 1989.

What this means is that is when we hired a contractor to build the spacecraft.

And the start of that process is called ATP.

We had two separate ATPs.

Sometimes you can say I can let out one contract to a company and let that company buy the science instruments or I can compete the science instruments separately.

And there are various pros and cons, but in this case NASA decided to compete the instruments separately.

From an engineer's point of view, one of the things that should peak your interest or note that when you have separate contracts you always have a question of integration.

If you are going to compete separately, you are going to have two separate pieces.

Who is going to integrate them?

The integration job will be harder when you compete them separately.

Now, your performance is probably better when you have direct insight into the instruments and direct insight into the spacecraft contractor.

But the penalty is now, when you integrate the two, you have to make sure that works.

And so, as a project manager, as a systems engineer you have to make sure you apply enough resources to make sure they work together and continue to work together as they are defining the interfaces, for instance, between the two elements.

This was interesting in that in order to get it funded, Congress said we will let you build an AXAF but, along the way, you have to do some testing to show that you can meet the performance that you say you can meet.

And so, we were required by law to have a test of the mirrors in June of '91.

And we also had another test later on, the same thing, to verify that we could get the performance that we showed on paper with actual hardware before we committed to fly and before NASA committed to spend the rest of the money to fly.

And so, sometimes that happens, sometimes if you have a smaller program it doesn't go to Congress so you don't get into those things.

But, when you have a great observatory and NASA is spending a lot of money, a lot of times Congress will get in and say yes, you're constrained, but you have to show me along the way that you can do this.

I wouldn't be surprised if Congress does that for the CEV, the new vehicle coming up, or the CLV.

In fact, for the CLV, the launch vehicle, there is a push to have a demonstration test early just to verify that we can

do it.

Now, during this time, of course, we are still formulating requirements.

And we rack up the cost.

And, low and behold, NASA cannot afford it.

We went through a program restructure at that time where we dropped a couple of the mirrors, we changed its orbit, we removed crew servicing.

And you could tell that the whole shape of the spacecraft changed.

Again, that is unfortunately not atypical.

It happens quite often.

Yes.

What was the reason for the orbit change?

Good question.

In orbit there are Van Allen Radiation Belts.

And we really cannot observe too well within those belts.

Now, in the first mission, the low earth orbit mission we were below the belts so we could operate.

But the problem with low earth orbit is you have this big earth, you are close to earth, and so you need power from the sun.

And, guess what, the sun goes behind the earth.

And so, you are eclipsed for a good percentage of the orbit time.

When we had a 15 year mission in low earth orbit we said since we don't have servicing, the lifetime cuts down to five years.

Now, if I am five years in low earth orbit, I have just lost two-thirds of my mission time.

So let's see if we can crank up the orbit, get as far outside the radiation belts as we can.

And then we have basically 100% visibility, the whole orbit, if we can get completely outside.

The problem with that is that costs propellant because the radiation belts end, they breathe, but roughly around 60,000 kilometers.

We could raise apogee to 140,000, but we didn't have enough money to raise perigee.

It gets down to money again.

So, perigee stayed at about 10,000 kilometers.

The resulting orbit time, if you looked at the integral of time outside the radiation belts, was probably about 70%.

It was a compromise.

We did end up losing some overall time, but about 70% of the orbit is usable for viewing.

That is the strategy and that is why we went to a different orbit.

As you can tell, there are a lot of things that are traded off during that time to fit within the constraints that you have got.

Again, that is typical.

The three reviews that are kind of standard reviews for all programs and projects are you start with a requirements review.

And that is pretty basic.

It is held early on.

Requirements have started before that review.

For science missions, you start with science objectives.

And you are percolating science requirements.

And the system requirements review is really a review that baselines, if you will, and sets the requirements as firm as you can.

And it is usually after you have gone through the cost gyrations, but it is done prior to doing any of the design

reviews.

Once you have a requirements review done, it was in December of '92, almost two years later we had what is called a preliminary design review in November of '94.

The third review is called a critical design review in February of '96.

What makes a preliminary design review and critical design review has to do with the maturity level of the design.

For critical design review, supposedly you have on the order of 90% of your drawings complete and on the order of 10% of your hardware built.

And I think for PDR it is about 10% of your drawings built, the final drawings.

That kind of gages what you are talking about there.

And it is a typical milestone that NASA uses for all of its programs.

Usually these are set early on in the mission, in the program milestones.

And, for political and for programmatic reasons, you tend not to deviate.

You try and meet those milestones.

Even if the project is not as mature as you would like, usually that is one that you try and keep hold of.

Maybe you don't have 10%, maybe you have 5% or maybe you don't have 90% you have 70%.

Sometimes you do go and you hold the milestone, you hold the review, but when you are not mature enough sometimes you have to hold a delta review to catch up.

And that is, again, something that NASA does for various reasons.

If the situation is right, sometimes you can slip and allow a single review at the right time.

Most often it is driven by nontechnical things that require you, either by your customer headquarters, they don't want you to slip so you tend to hold it where he said he wanted it.

And then we had the mirror delivery to the calibration facility in November.

It shipped the whole thing to the Cape for launch in February and we launched in July.

You can see we spent about five or six months down at the Cape integrating into the Shuttle testing, and then we flew in July of '99.

Those are the orbit parameters.

We achieved the final orbit August 7th, so you can see about two weeks.

And the way we did that is the Shuttle only puts stuff to low earth orbit, so you need some additional propulsion.

We had an upper-stage that got us part of the way there, used most of the energy.

However, it wasn't at the final orbit.

So, integral to the spacecraft, we had a propulsion system that had to do five additional burns to get us to the final altitude and the final orbit parameters.

That is why it took us about two weeks.

And, again, that is not atypical.

When you are not operating in low earth orbit, it takes you while to get your final orbit.

And these are a couple of interesting parameters.

Safe mode, we will talk a little bit about that.

Spacecraft are designed such that if -- Yeah.

I am just wondering if the inclination affects the requirements at all or does it not really matter so it just stays [NOISE OBSCURES]?

The 28.5 is driven largely by the launch site at Kennedy Space Center.

Right.

I am just wondering if the science requirements could drive the inclination [OVERLAPPING VOICES].

For some missions, yes.

We were fairly incentive to that, as long as we took advantage of launching at KSC at the optimum inclination.

But there are missions like Space Station.

Because you have a launch from KSC and a launch from Russia, to optimize the performance of both launch centers then the orbit is 58.5.

Now, what that means to KSC launch is that you pay a significant penalty.

You pay about a 30% weight penalty for launch at KSC to the Space Station.

Sometimes the answer is yes, you can optimize inclination, and sometimes, when you have a large program and multiple launch sites it is a compromise.

So, it is not efficient for each site.

I was mentioning spacecraft typically have their avionics systems.

And typically they do not contact the ground.

They are not in communication with the ground all the time.

And so, for Chandra, that is the case.

We may have "contact" with Chandra probably about 5% to 10% of the time per day.

In the meantime, figure that 90% of the time it is out of communication range.

Now, this is something, by the way, that is somewhat different from a manned mission.

The manned missions tend to optimize and maximize communication coverage.

When Jeff is flying in the Shuttle we probably have 90%, 95% coverage over an orbit and over a day for communication.

But, in science missions, rarely is that the case.

And so what that means is the whole philosophy of operating a science mission is based on stored commands.

And so, we send up a stored command load.

And that load is good for a day.

And basically it steps through and automatically executes based on time.

What that also means is you have to design a spacecraft to recognize when it has a problem.

If a hardware box has a failure, it has got to be able to recognize that and go into a safe mode so that the ground can then recover and configure the systems to continue to operate.

That is one part of flying a spacecraft, is building a robust safe mode to be able to do that.

And our first problem occurred August 17, 1999.

Now, note it was a ground error.

As usually the case, spacecraft are pretty complicated.

And we probably had 400 or 500 separate procedures and commands on the ground for doing certain things.

And, even though we try and test them out individually, the combination of those procedures and in the execution, sometimes you put the spacecraft in the wrong configuration with one sequence, and then you follow with another sequence.

Low and behold, the software that you designed for safe mode says I am in the wrong configuration.

Sorry, I am going to safe mode.

And so, that happens.

We have had only probably three safe mode events for six years, so that is not too bad at all.

Hubble, for instance, has had probably an order of magnitude more than that.

I think we have gotten smarter based on our design of Hubble to design this.

But, again, typically spacecraft do have safe mode designs in them.

Given our orbit, we still have a small eclipse season.

Twice a year in the fall and in the spring we have eclipses.

And so our first eclipse season was then.

This is the translation mechanics for the science instruments.

This is where the telescope is.

I won't go through that.

This is kind of where we purchases or we contracted for the various instruments.

This is a CCD imaging spectrometer.

In fact, it was both Penn State and MIT that produced this one.

And that is one of the focal plate instruments and the other one was a high resolution camera using micro-channel plate technology instead of a CCD.

And that one was from SAO just down the street in Cambridge.

The high energy transmission grating, as I had mentioned, was also from MIT.

And the low energy grating was from the Netherlands.

And this is a good point to make also, is that a lot of times for science instruments or for science missions, NASA tries to lower the cost of the mission.

And the way you can do that is you can reduce requirements or you can get a partner and have the partner pay for an instrument that you want to fly.

And so, in this case, as is typically the case for science missions, NASA tries to go and get international partners that will bring to the table an element or a piece of hardware or some software that reduces the overall cost of the mission but, at the same time, maintains its capability.

That is what we did with the low energy transmission grating.

We got the Netherlands to basically design and build that and fly that in exchange for that kind of arrangement.

They get a percentage of viewing time that they can say for 5% of the viewing time we get to point the telescope to wherever we want.

And so, typically that is kind of the handoff between our international partners and NASA in trying to reduce the overall cost to NASA for a given mission.

And these are the gratings, as I had mentioned.

And there are a bunch of facets along these.

Now, notice these are the four mirror shells where they co-align with the four mirror shells, the inner most and the

outer most.

And this is kind of how it works.

The gratings get flipped in.

And this is the array of CCDs.

Each one of these is a CCD.

And the pattern of light or the pattern of x-rays follows along these lines when the x-rays actually come in.

That is what actually you see.

Now, Chandra does have a ground system architecture.

We communicate with the deep space network.

That is how we communicate because it is not in low earth orbit so we cannot use the satellites that are around the earth.

We get that signal from that satellite and it is beamed to the control center, here in Cambridge, in fact.

And we do various things on the data.

We process it out and send it to the general observers.

Now, the way this operates is that every year this control center or the science folks operating the control center solicit objectives.

OK, science community come and request time and give me a proposal for viewing time for what you want to do, what target you want to look at, what is the configuration of the instrument, which instrument, et cetera.

And so, the Chandra folks go through, and they probably collect about 800 proposals a year from all around the world.

They will select roughly about 200.

And it is not based on numbers, it is based on time.

You have a certain amount of seconds in a year that is available.

And so, some of these proposals have large times, some have small times, but roughly it is about a four to one.

About 200 are selected, 800 are proposed, so that's good, you're oversubscribed, you have a lot of interest.

NASA does fund those that are selected.

So if you are a PI and you say I want to go look at Crab Nebula or I want to go look at a Quasar somewhere, and here are the reasons, here is what that will tell me.

And it gets selected.

NASA will fund you some money to be able to do that observation and process the data and publish.

Now, NASA won't fund any that were selected from foreign countries.

So the foreign countries, they have to get their funding, but we will allow them viewing time.

And that is also typically the case.

And, by the way, all of these scientists across the world never have to come to Cambridge.

It is all done electronically.

They don't get their data real-time.

They get it probably a few weeks after the observation.

And the reason for that is there is a lot of processing.

You don't actually send, or you can, but usually you don't send the raw data out to the user.

It is heavily processed, and the products of that processing are what is actually sent out.

Now, there are tools that the Control Center also provides, but generally it is the process data that is sent.

Yes?

What sort of data is this?

Images or numerical data?

It is data that can be used for images.

It can be used for registering each photon, what energy it is.

There are a lot of different choices of data.

And some of the data is overlapping, but it does give you information on energy, location, number count, where in the field.

It can be used to then process an image or process a spectrogram.

And you do that by using the tools that are also provided.

It is like a viewer or it is like a piece of software that allows you to view what is in that data.

But the data is just the raw data.

It is just a series of tables.

And, as we said, it began in 1978.

I picked the prime and was given a new start.

I am not going to go through these.

It was restructured.

We went to four mirror pairs.

We dropped two focal plane instruments, dropped the servicing requirements.

I mentioned that.

Interestingly enough, one of the instruments we wanted to create a separate spacecraft just for that instrument.

And, believe it or not, the addition of this spacecraft and this one for the spectroscopy mission was going to be cheaper than the original concept.

But again, due to cost, we continued this spectroscopy spacecraft for about a year to two years.

And then that was cut due to cost.

So we ended up flying just this one spacecraft.

And the result is the Chandra Observatory that we have now.

Note that sometimes it is not within your control as to whether your program is cancelled or not.

We were doing quite well.

It was within our cost envelope that they had given us.

But, at the time, there were other priorities at NASA and there were other overruns at NASA.

And they have to weigh.

Sometimes they say, well, you can have this much money for these many years.

And every year they re-evaluate.

And that is because you don't know how much things are going to cost because NASA typically builds one of a kind things, and you don't know really how much it costs until it is already done.

So that re-evaluation process continues now.

And it will probably continue for as long as NASA flies.

So imaging and spectroscopy is what we do.

Why is imaging so important?

Well, clearly this image was a Rosat image with the highest technology at the time.

And here is the Chandra image.

And you can see the neutron star right there.

And there is no way you can see a neutron star in there.

And they can learn a lot more from this image than they can from this, so that is why scientists push to get good imaging resolution as well as good spectroscopy.

The biggest challenge for us on Chandra was the mirrors.

Now, recognize that due to the analysis that we performed early on, yes, this was all theoretically possible, but we had to be able to polish the surfaces of those mirrors, those grazing incidence mirrors to an accuracy of on the order of angstroms.

And that is pretty small.

Can we measure that?

The answer is no.

At the time, we couldn't even measure how accurately we needed to polish the mirrors.

So that tells you that, boy, we've got our work cut out.

We had to work with the National Institute for Standards to figure out how to first measure how accurately we can polish it and then ended up polishing the mirrors.

And polishing, as a process, getting down to that smoothness clearly is a huge challenge and takes a long time.

And so, that was probably our biggest challenge, is to polish the mirrors.

Metrology is the science of working mirror technology and getting the surface figure of the mirrors correct.

So, really, the key to success is developing three different types of metrology measurements that are independent, that allow you to cross-check.

Because you are talking about things you have never built before and you are pushing the state of the art and you don't know how to measure it.

The key there for an engineer should say, OK, I need not just one way of checking my work, but I need at least two ways, and preferably three independent ways of making sure that I am doing the right thing.

So when you're pushing the state of the art, a good lesson is try and get into a position where you have got more than one independent check of your analysis to make sure it is correct.

So that was a very big challenge for us.

And, again, that is kind of how it goes down.

This is the general shape of the mirrors.

And here are the mirror pairs being assembled.

And, of course, one thing about the mirrors is when you are talking about that surface finish, you cannot just assemble it in your garage with dust around.

You have to be extremely clean.

And, believe it or not, one of the dirtiest things in a laboratory is a human.

And so, a human has to be almost completely covered to prevent any kind of accumulation on the mirrors themselves.

In fact, we had to limit the exposure time of humans to the unpolished and polished mirrors during this timeframe because contamination is an incredible problem when you are talking about the atomic scales of polishing the mirrors.

Again, I mention that the important thing is to make sure that you are doing cross-checks in the metrology as you are going along.

Because, again, you don't know exactly what you are doing.

I mean on paper, yes, you've got an analytical solution that tells you the right answer, but you have never done it before.

You have never built it before.

You don't know how to measure it.

So you have got to work on building confidence that your analysis is correct.

How do I know that analysis is correct?

And so, you have got to think about a sanity check, if you will, to make sure that those are correct.

And, of course, the best way for proving that it is correct is test them.

And that is what we did.

Mandated by Congress, we tested the mirrors.

Now, the thing about that is so I test the mirrors.

I am testing the mirrors in a 1G environment so, even after the test, how do I know that those results are correct?

Because the 1G deflection, even though you think it is small, it does affect the mirror performance on the ground versus in orbit.

And another thing is the finite source distance.

You cannot get a source that is infinitely far away on the ground just because of the nature of putting a source out there.

Our source was about, I think, a quarter mile away.

And we had to make analytical corrections to compensate for how we expected the spot size to change because of that finite source distance.

We had to analytically correct the image on the ground to compensate for the 1G affects of the mirrors.

And so, even a test you would think, OK, go test it.

Well, it is not straightforward because, even in that test, that is not a true representation of how it would work on orbit.

So it is a lot more complicated than one would think.

What do you use your x-rays for?

Various things.

We had some monochromatic x-rays at various wavelengths.

We had registration x-rays.

Iron.

A certain wavelength of iron is a big one to use.

You don't have a picture of the big tunnel, do you?

I think I do in here.

I have got a picture of it and I can talk about.

But that is another thing, is when you are testing, how well do I know my source?

And that is all factored into verifying that you have got the performance correct, because if your source is all over the place and you haven't characterized your source well enough then how in the world can you say that the

mirrors are that good?

So we did end-to-end testing.

And, again, you don't just look at the raw data and say, yeah, I am there.

Even in the ground test, a lot of times you have to massage the data and compensate analytically for factors that you cannot control.

And the key to success there is good systems engineering.

And a lot of that is because of multiple separate parties, academia, nonprofits and foreign groups all involved in the workings of AXAF.

You have got to make sure that you have got a team that works together.

Lesson learned, as I had mentioned, perform multiple cross-checks either via test or analysis with a different tool as possible.

Another lesson learned is let more than one group perform the review.

A lot of times, for Chandra, we had the Marshall engineers, but we also contracted SAO down the street to kind of independently assess, using their own software tools, where we were.

And that was a good idea because sometimes we had good ideas that we could incorporate and sometimes they had good ideas, but it gave us a sanity check when we both matches and we both felt good.

There is also no substitute for direct test or measurement.

Analytically, we have wonderful tools nowadays, but they will never take the place of testing.

Another thing that was very important for us because of the iteration cycle with headquarters on funding is make sure that you keep the science or the scientists that are part of the mission informed and part of the decision-making process.

One would think, and one would like to think that this is a technical project, all things are technical.

No, there are personalities involved and there are human responses involved.

And so, to keep the team together you need make sure that the science is informed and supporting your decision-making process.

A lot of times that is overlooked, but it is, nonetheless, very important as you are building a project.

And the final lesson learned is, as we are going through and making those mirrors, you construct what is called an error budget.

You assume perfect performance and then you kind of step back and say I am going to allow imperfect performance in this region.

And I am going to allocate that imperfect performance to thermal design.

I am going to allocate imperfect performance to the inaccurate knowledge of knowing the source positions in the sky.

I am going to allocate some error budget to flexing of the optical bench.

I am going to allocate some error sources to the difference between how long the focal length is and how long I think it is.

And so, you take each of the error terms.

And for Chandra we probably had a hundred or more.

And you go through and you verify that your assumption in that error budget term was correct.

And usually, for the critical ones, you did it in more than one way, more than one method so that you were sure that you had those error terms defined.

And the end result was we had mirrors that performed better than expected.

And, as a result, the payoff is just wonderful.

This is a time-lapse series of images.

And, supposedly, this wave front was moving out from the source pulsar at significant fractions of the speed of light on the order of 20% to 30% of the speed of light.

That was pretty cool.

And if you need more information just talk to Jeff.

But that was a pretty big payoff.

The next challenge that we had was the programmatic challenge.

Again, as I had mentioned, the normal process of going through a program in a project is you are first trying to size the mission.

And, really, you are sizing it for cost.

And that is very hard to do when you are building something new.

And that was a very big challenge.

And so, we changed a lot of things.

But, eventually, the key or the critical thing that you are trying to do is I compromise on this but I want to still try and maintain performance.

We maintained our imaging resolution performance and our registration performance.

And we compromised on a few other things but we maintained that performance.

Finish this slide and then we will take a quick break.

Note down here that we also needed to reduce the weight by a factor of two from our original designs.

Our original designs were pretty simple aluminum structures.

Our final designs were the best composites that we knew how to make at lowest weight.

And so, that optical bench, that tube was all composite.

And so, that is the kind of weight reduction scheme we had to go through.

Now, also all of the fittings, that means all of the brackets that we use were also composite.

And another challenge that we had is analyzing the stresses in a complicated composite fitting.

Very hard to do.

Didn't have techniques to do it when we started.

As part of the process of building this telescope, we had to go through and learn that and develop those techniques.

Before the break, I just want to point out I hope you all are catching the incredible parallels between this project, a complex technological process and the sorts of things that we've heard about the systems engineering of the Shuttle.

Right from the weight problems and composites, new materials, new ways of analysis right into the intervention of Congress into the engineering process, I mean it is all there.

And I think whenever you get involved in a big space project you can expect that those things are going to happen, and they are probably going to happen in the Exploration Program as well.

Two minute break and then we will start up again.

Start up again.

Let's see.

We were talking about for us the next biggest challenge was programmatic in that it you have to work that early on.

And typically, especially the more expensive missions do go through this as a typical part of the lifecycle of a project.

I have a question on that last slide.

[NOISE OBSCURES] The budget for systems engineering was cut which translates directly to the number of folks at the contractor site.

Working systems engineering was cut so we were challenged.

And what we did at Marshall was we offset with NASA's civil servants doing functions and tasks that the contractor would normally do to compensate for that.

So that is what that meant.

Typically, systems engineering, a very important thing if you talk to folks in the trade, almost everyone has a different idea of what all is encompassed in systems engineering.

But the point is there probably are a number of things that almost everybody would agree are part of systems engineering.

And they are very key to executing the program and the project.

They are really the glue that holds all the subsystems together.

And I will probably talk more about it, but for good systems engineers you have to be able to communicate well with all the other subsystems and all the other stakeholders, including the science when you have a science mission, including the users when you have a manned mission, including safety and everybody to make sure they are all connected.

And be smart enough to be able to challenge the subsystems when they really have their own ideas on where this needs to go and it is not really in concert with everybody else.

Weight reduction did have some impacts, as I mentioned, using light-weight composites.

And that was a challenge in and of itself.

We also dropped two of the SIs.

Another challenge was the science instrument module which was all composite.

And it had the capability of adjusting focus and translating, so it had two degrees of freedom.

But it needed to be reproducible in terms of motion on the micron scale over the whole orbit environment.

That was quite a challenge in and of itself to be able to do that.

I mean just kind of a flavor for the things that you have to deal with is not everybody had the same analysis tools so universities like MIT and like SAO had different structural analysis tools for their hardware than the contractors did.

Now you do a couple loads analysis, which is you pull together all of your stress, all of your structure analysis models together in one place and you go through a loads assessment all through the whole element.

Well, to do that I need a model from you, I need a model from you, I need a model from

you. But you are in California, you are a university, you are another company and you have to make sure they all play together.

And a lot of times there is parochialness with company A, I am not going to use anything but my model.

And the university says, well, I cannot do anything but what we have been doing.

And university C says, well, this is the only thing that will work and you guys are crazy.

It is not that bad but that is the kind of thing that you have to deal with as a systems engineer.

You have to make sure everybody plays together.

And sometimes you have to go through and work the best solution that may not be the optimum solution for each piece.

But it is the optimum solution for the overall answer.

Again, getting the programmatic challenge behind us, a key thing is also setting allocations.

I talk about the error budgets.

Another good thing that the systems engineer does is allocates various important resources.

And one that every spacecraft does is weight.

Weight, power and data are usually three of the important resources that are allocated.

For scientific telescope missions another one is error budgets for scientific performance, but the key to that is good weight allocation and good iteration, which with each of the elements that you allocated the weight to, to keep up with them to make sure they can meet that or it looks like they are going to exceed.

So I have got to go borrow some of that allocation from somebody else.

And I have got to take away from Peter, if you will, to pay Paul.

So systems engineering and going through that making the initial allocations and keeping up with that is a pretty critical factor to making sure you get to the final answer.

And, if you do that well, then usually you have less of a problem getting to the final solution.

To compensate for the loss of systems engineering, we established technical oversight panels which were NASA employees at our center.

We controlled, at a project level, the internal ICDs.

Now, what is an ICD?

It is an interface control drawing.

Because when you have company A developing an element and university A, where they come together is an interface.

And you have to document that interface, not just mechanically but thermally.

What is the heat flow across the interface, the data across the interface, the power, the signal characteristics across the interface?

The data flow.

What data is going across the interface?

And, as you can realize, as both of those things change, you have to change the interface.

So you have to keep up with that.

And a lot of times good systems engineers will recognize the further away in location it is the harder it is to interface with the two parties.

When you are in California it is one thing when they are in the same town, it is another thing when they are in the same country and it is another thing where they are across the water when you are integrating an international partner.

So, as systems engineers, the factor is the closer you are the better it is because I can just go in a car and go across town to get out the drawings and talk about the drawings versus in California that is a plane ride.

But, in Europe, the time difference allows you to only be able to talk to them for a few hours a day because that is when your work shifts overlap.

So it is a lot more complicated interacting with folks in Europe, interacting with folks in Japan, in India because of the time difference and because it is very far away when you have to sit down.

Those are kind of the challenges as systems engineers that you should think about.

Lessons learned.

Set your resource allocations early and continue to monitor them and work with each of the owners of the elements of those allocations to make sure that they are meeting their allocation.

As you recognize, it is a zero sum gain.

So when somebody goes over their allocation somebody has got to be reduced.

And it is very difficult sometimes to negotiate that reduction because engineers typically want to hold their margin, they don't want to release it.

They want it because they are not yet confident in their final solution.

But to make everything play sometimes you have to take away from them and reduce their margin to be able to let somebody else survive if you will.

It is not an easy thing, but attention to that detail is important.

Maintaining strong systems engineering group is important for that reason.

If I have a lot of pretty dominant subsystem managers that say I am building a power system and I need this much weight and I need this much thermal capability that's good, and they may be very good, but what they are doing is if you have a lot of dominant subsystems and no central systems engineering, as traffic cop there is no check and balance.

And so that dominant thing may work very well at the expense of the integrated performance of everybody else.

That is why systems engineering is so important, that dominance and to balance weakness where you have weakness.

And for Chandra hard work pays off when you get it right.

Technology.

A lot of times NASA will build a science spacecraft.

And the spacecraft itself usually isn't pushing a state of the art.

For Chandra it was a little different.

We pushed the mirrors, and there are a few other things we pushed.

But typically it is the case where NASA pushes the state of the art on the instruments.

Better sensors, better CCDs, better processing electronics.

And Chandra was no different.

Usually it is an embellishment from something that has flown before on a different spacecraft, only now it is better.

We are getting it bigger.

We are getting more resolution.

It takes less power.

There is usually a steppingstone.

And for ASIS, the one that was built at MIT, it charged a couple of devices in an array size that they have never been built before.

And so that was a challenge.

Also a challenge was in order for the CCDs to operate they needed to be cooled to minus 100 degrees Celsius, 120.

And, in addition, they developed a new way of clocking out the data from the rest of the CCDs that have been built.

That had better spectroscopy performance and a very low noise signal chain.

There are always going and they are always looking at, well, on spacecraft X this CCD was flown.

It may have been a 512 x 512.

Well, you know, it is pretty easy.

I can stretch it a little bit and work the substrate in 24 x 24.

Well, usually it is the case where it is not as simple as you think.

And there are always problems.

In particular, technology should be a warning flag to systems engineers to say typically it is not as simple as you think so make sure you are paying attention to these new developments or these stretching of new applications of existing technology.

Because, for instance, one of the things that you get bit by, and we got bit by is radiation susceptibility.

Thermal extremes for multilayer circuit cards that have to operate at minus 120 C on one end and room temperature on another end.

Low yields for those new types of CCDs.

I can just stretch it and make it 1024 x 1024, but where I was using 90% of the batch, now I am only using 10% just because I cannot get one that big to work out and to be smooth and homogenous.

ESD sensitivity.

Because they are getting bigger and bigger and thinner and thinner it is much more susceptible to electrostatic discharge.

And simple things like my yield in the wintertime, because the humidity in the air is lower, I have less of a yield in winter than I do in summer.

And so that is what we found out.

And there are also mitigating effects that you can come and put humidifiers in and compensate for that.

But those are the types of things that you have to worry about that can come and bite you with new technology because you are doing something new.

Because of going down to 120 to optimize the performance and to get the very low noise, we said we need a radiator or a sunshade.

And, low and behold, flying on the Shuttle in certain attitudes that you have to verify, because in some off nominal events you can, instead of being deployed on day one of a mission you get deployed on day three, but you have to go across the sun terminator frequently and you get a little glint of direct sunlight on these sunshades.

And, therefore, you have to analyze what the temperature extreme and what the effect is going to be.

And, low and behold, we see that we have delamination on that case.

But here is a case where it is an off nominal event in the Shuttle that we have to design for and we have to compensate for, even though chances are we will never see that design case.

But a lot of times the programmer or project manager has to say I am going to suck it up and I am going to make the change for that small eventuality.

Or I am going to risk it based on cost and we are going to try and preclude that eventuality from happening.

A lot of times it is not cut and dry.

It is a guessing game.

It is a balance that is based on judgment, past performance and a lot of other factors to say, boy, that is so low probability of occurrence that I am not going to worry about it.

And that threshold is never defined.

You have to define it as a project manager.

Just kind of a picture of what the instruments look like.

There are the CCDs.

This is the long stretched away.

And, for the spectroscopy measurements, that is where the lower energies -- Actually, the higher energies in the middle and the lower energies go along this wing and along this wing because you get the energy spread.

And there are the imaging devices that are used when you just want a good image.

On HRC, there is the detector right there in this diamond-shaped thing.

And it is operating off a completely new and older technology than the CCDs.

The instrument of choice has been this for probably 70%, 75% of the observations just because of the capability that this allows you not just an image, but the CCDs also give you some rudimentary information on energy.

And so, that is why this is preferable over this one.

On HRC, we talked about ASIS.

We had, again, a low noise signal chain.

On HRC, we never had three micro channel plates tied together and linked together.

And we had that challenge.

We also had very accurate event timing.

So, when we have a pulsar, I know exactly in registered time when those pulses are occurring.

We had spacecraft charging protection for new technology problems and spurious noise susceptibility.

Spacecraft charging.

To keep the mirrors nice and warm at plus or minus half a degree C across the entire mirror surface all around it required thermal blankets on the outside.

Well, the thermal blankets were susceptible to charging up and discharge with popping.

And so, the question is, is that going to affect my instrument performance by incorrectly registering a photon event every time these blankets discharged?

So we had to go through a test campaign to verify that, no, this discharging of those blanket electrostatically would not affect the imaging performance at the detector level.

That is something we didn't even think about before.

But, low and behold, we came up with it.

Yes, these things are going to discharge, they are going to pop so is it going to affect our measurements.

Those are the kinds of things that you will run into.

Again, good systems engineering and sound communications are the key to getting through those.

Insuring participating.

Keeping parties informed.

Encouraging teamwork to be part of the team and to help contribute to its success.

Lesson learned is establish standing interface working groups with mandatory participation from the SIs.

The SIs kind of have a different paradigm.

Science instruments are usually built by universities, and the universities are typically staffed by grad students and some full time engineers.

But usually the paradigm and the thinking process is a university atmosphere.

When you have companies it is a completely different interface, very regimented, very by the book, we have done it before and this is how you do it.

And so different cultures, you have to make sure they blend together to produce the final product.

And a lot of times that is not technical.

That is human interaction and making sure they all work together.

When problems occur recognize that you need to go outside of the project office to get help.

And this is typically done at NASA.

If Marshall is managing a project and we have got a problem with CCDs, the key to success is recognize, well, Ames Research Center in California has an expert in CCDs.

Call him up on the phone or bring him down and let's talk to him to get his expertise in on this project.

NASA does good on that.

It does go to where the expertise is.

Not just at another NASA center, but it may call experts across the industry or across academia to come help with a problem.

The foam problem on the Shuttle is a pretty big problem, and we have called experts from around the country and around the world to try and help understand what was going on with the foam.

That is something that NASA does well, is it brings in experts and it doesn't think that it knows everything locally.

And that is something to keep in mind.

If you don't know an answer, don't just try and think of it yourself.

Go get the answers from somebody who might have experienced that before.

And encourage teamwork.

A lot of times programs suffer from the culture clash between different paradigms and different cultures. And that is especially the case between academia and industry.

It is also evident between international partners and industry.

And, of course, the payoff is this is a spectrogram, the counts per bin, number of photons per energy level.

And, as you can tell, very sharp and very well-defined peaks for where those lines are.

Next challenge was integrated testing.

Integration is always a challenge.

How in the world are we going to integrate this whole thing?

The logistics of who does what when, we must choreograph it.

Here we have instrument A coming in.

We have instrument B.

We have the gratings coming in.

We have the ISIM built at Ball Brothers.

We have California building the optical bench.

Where are we all going to put this all together and test it out?

That is always a challenge because, again, you have different players involved in different pieces.

And you have got to make sure it all plays together.

And the key there is to create a working group to make sure you address every aspect of integration.

You plan very much ahead of time.

You choreograph everybody coming in to make sure you have got everything covered.

And in integration always, always it takes longer than you think.

You always plan assuming nothing is going to go wrong.

This test takes this long and this test takes this long and this test takes this long.

When, in fact, as is always the case, you always have problems in testing.

Any time you are matching two boxes together, especially electronically, a lot of times you have problems.

So testing is always something that you should make sure you have enough time to do troubleshooting.

And plan for failures and not plan a success-oriented schedule.

Even the best laid plans, TRW planned the integrated testing in a vacuum chamber.

And during observatory integration TRW couldn't hold the schedule because, guess what, the first thing that goes when you are in a money crunch is testing at the end.

And you say why do that?

Because they say we had six months of margin in the schedule and you have eroded that margin.

But that margin you haven't really defined anything to go into that margin.

So, hypothetically speaking, if everything goes OK, I don't need that.

Well, so you take it off the table and it gets eaten away.

And that is what happened.

It got eaten away and couldn't hold schedule caused, in this case, by antiquated EGSE and software.

But it is typical of a paradigm that says I create schedule, you have got to try and preserve this schedule and fight tooth and nail to keep that reserve in every step of the way because in integration test is where usually you get behind.

When we got behind the recovery is daily focus on schedule.

Every day we were on the phone going over what was to occur that day.

24 hour operations one year before launch continuing all the way around the clock.

We changed the I&T integration and test lead out at TRW.

We brought Marshall NASA engineers into help TRW do it.

And we had a technical presence help out at the test site.

We have a contractor, we hire a contractor and the contractor is supposed to do the job.

And, low and behold, the contractor is not doing the job.

So we could take a step back and say whack them on the head, you are not doing the job.

Or, you can go in and say we have got to get this job done.

We are going to go and we are going to send folks over there that are technical folks that will help your guys work these problems and work the issues and work the integration and get this thing done.

And so, that is what we did.

That is kind of a typical schedule.

We have a Comprehensive Acceptance Test.

That is just a functional test going through all of the paces.

EMI is your kind of electrical susceptibility to noise.

When you have a palm pilot or a cell phone and it is near a computer or near some device you hear this chattering.

That is what that test is all about to look at interference patterns and look at whether you are susceptible to noise.

We have an end-to-end test.

That is where we get a compatibility van.

And we make sure that we can communicate with the deep space network dishes in Goldstone and Madrid and Canberra.

That is our pointing at control systems polarity test.

Did we mount the gyro correctly or is it upside down?

Did we mount the reaction wheel correctly?

Do I have the spin vector pointed this way or pointed that way?

You think that is pretty simple but, in fact, on the box it doesn't say arrow pointing up.

You have to actually test it out.

And NASA has screwed it up before.

We do an acoustic test and then a pyroshock test.

Acoustic is you blast it with acoustic noise and the noise couples and creates vibrations, especially you are worried about thin films and large surface areas that are exposed to the noise pattern, the wave front.

So we painted the whole spacecraft with a huge woofer.

I mean it was enormous.

Just to see the vibration and what would happen to make sure that it was OK.

And we have the solar ray full motion test and the low gain antenna release test.

We have mechanisms which are typically things that fail a lot or don't operate the way you thought.

We want to test those hopefully in the environment that they would see in orbit.

That means if it is cold, if it is exposed to a plasma we want to try and duplicate that on the ground.

Typically, that is what we do with mechanisms.

Yeah.

I was wondering if you had like a separate satellite that was forecasting and they kind of knew that [NOISE OBSCURES].

Good question.

Typically, in the old days, when we had a lot of money, we would build what is called a qualification article which is basically a test article for the whole thing.

And the qualification article would match one for one the flight unit.

And we would test the heck out of that qualification article.

Nowadays, with money problems, typically we don't do that.

For selected components where we are worried about the design, we will develop a qualification article at the element level.

But typically we don't do it with whole spacecraft, not anymore.

Although, the one exception to that is the structural test article.

Usually spacecraft do build a structural test article that does look like the final flight design from the overall structure.

And you get your mode shapes and your frequencies when you ring it and see how it performs.

And you fold that back into your analysis.

But that is probably the only case where you build a full test article for structural purposes.

It would be nice if we could do that.

The DOD does that when they have a production of a thousand units, but NASA nowadays cannot afford to build a whole complex ground test article.

It would be nice but sometimes we just cannot do it.

We will do it on a selected element level.

Hardly ever testing to failure.

We talked about that with the Shuttle because JR Thompson was pointing out, like with the main engines, testing to failure was really critical for success.

But, remember, that is a very different operating situation than a satellite where normally if the satellite can survive launch and the vibration, and the vibration is tested, the structural, you know, if that survives then the operating environment is a lot more benign.

Now, what do, though, is we compensate for not testing to failure by testing with two our expected flight environment plus margin.

Analytically, or based on previous data, we determine what the environment is going to be.

And we will add margin to that and we will test to those margins.

Most times we do pretty well.

Sometimes we don't, but that is probably few and far between.

And we have learned, after 40 years, of how to apply margin to our best analysis or our best tools that define what the environment is going to be.

So that kind of shows you starting in October '87, ending April '98, that was the plan.

We actually took about nine months more than what you see there because of all the problems that we had.

A key point is when you are pressed for time do not cut corners.

We did review all the testing and we did not delete any testing that we considered was mandatory.

And, in fact, we added systems testing to verify that we had total system performance that we knew what this thing was going to do.

We also added some end-to-end testing in thermal vac and added much more end-to-end testing with the control center that will eventually operate.

A lot of times you have, in the integration and test time, a different set of engineers testing this than will actually operate it in flight.

Probably not a good idea.

What you really want is to get your operators in flight operating and testing early enough that they can go through all of the gyrations of the testing environment and all of the idiosyncrasies of the hardware before flight so that they are trained on the hardware and they know how the hardware performs.

And so, if you can, you want to try and get the operators to do the testing prior to launch.

And a lot of times that isn't the case, unfortunately.

Lesson learned.

Try to keep one integrated database along that theme of the Control Center for testing and operations because what you don't want to do is to have a database for testing that then you discard and you create a separate database for operations.

It is much better to have one database so that you know what the parameters are.

And, once you have verified them in test, they are the same parameters that you are using in flight.

Also, to define an explicit test and integration lead, review the approach, encourage end-to-end testing

participations of the operations group early and often and, if you can, get the operations group at the Control Center to run the tests.

And give adequate time for box level testing and data system integrated testing.

A lot of times what will happen is you're running late on schedule and at the box level I have a multiplexer.

And I had scheduled, at the box level, probably 300 hours worth of testing.

Well, you know what, I am running late.

They are wanting me at the spacecraft level.

Well, let me jip it up now.

It has only had 50 hours but now I am going to the spacecraft.

Low and behold, now what you have just done is you have moved the problem further on downstream when it is much more expensive and it costs you a lot more money.

When I have a problem in integrated testing, guess what?

The whole spacecraft stops until I fix this box.

If you can, try and do it so that you have got enough margin that you can test at the small box level in parallel with all of the other boxes before you get to integrated testing.

Management.

There were plenty of challenges in the management.

And a term used is unknown unknowns.

We didn't know at the time that money was going to be a problem with program restructuring.

You don't know that you are going to have an integration delay.

There are a lot of things that you really don't know.

But, nonetheless, you have to expect for and you have to plan for.

That is why, at the beginning of a project, you create what is called reserve resources.

In terms of cost, depending on the nature of the complexity, the technology push, you may have 20% to 30% of cost held in reserve.

Not applied anywhere because you don't know where the problems are going to occur.

And you try, as a project manager, to hold onto that by tooth and nail and not release it if you can at all and only grudgingly allow reserves to be used.

Approach to success in that kind of environment is experience, getting top management that has background and experience, working closely with the science community.

Again, focusing on teamwork.

You will see that as a recurring theme.

Hold enough reserves so that when you are scoping out and you've got a lot of unknowns that you have enough money and time to do that.

Set it as high priority for the company or the organization that you are working for and set the schedule early and balance the need for maintaining schedule with the need to slow down and understand what your problems are rather than racing to the finish line.

Again, these are some of the good lessons learned.

In fact, I just covered that one.

This is what it ended up being, the AXAF Schedule starting around calendar year '92 with our SRR.

There is our PDR and there is our CDR.

And the launch is even off the chart.

But that is typically the kind of things that we go through, including in parallel your Control Center development back down here.

There is a mirror.

All of this will be posted.

I mean these are really valuable lessons.

In fact, any of you who eventually go into systems engineering and project management take some of these lessons learned with you.

Print them out and review them periodically.

And remember them.

They are lessons of hard knocks.

What I am doing is now is I am now the project manager for a project or a mission called RLEP 2, which is Robotic Lunar Exploration Program.

And that mission will be the first mission that NASA has, in a long time, to land on the Moon.

We are going to land the Lander.

And we are going to have some type of mobility device that will either drive into a crater or hop into a crater or crawl into a crater or walk into a crater and look for lunar ice that we think might be there at the South Pole.

Now, this is a great example of how things are done.

We are given a couple of requirements to say we want to verify our precision landing capability in an unmanned sense, and we are going to use that technique and those algorithms for the future human lander.

And we are also going to need you to go look at the ice.

Characterize it and see whether or not ice is there.

Now, why is ice so important?

Ice is important because of the oxygen that you can get from the ice.

And the oxygen is not so much to breathe but it is for fuel or for oxidizer for propulsion systems.

Instead of getting it at the earth's gravity, well, you get it at the gravity well of the Moon.

That is our challenge.

Now, when they asked that, NASA didn't exactly know what it would cost or what the cost profile would be.

They originally designed it for a \$400 million mission.

They said, well, why don't you look at a \$750 million solution?

They don't know how many objectives we can meet.

This RLEP is supposed to be robotic lunar exploration.

It is really a precursor to the human missions, so they are not even sure how to use this program to help bring humans there.

And so, this is typical.

NASA doesn't exactly know what it wants to do, but it kind of puts out feelers to say try this.

It says go start with this, go touch the lunar surface, characterize precision landing and go see if there is ice there.

Well, go see if there is ice there.

I can do that with one hop into the crater, one sample and hopefully, if the ice is there I can find it.

Or, on the other extreme, I can go with a Rover, spend a year in the crater, take a thousand measurements and characterize the ice.

Obviously, one cost a lot more than the other.

So how in the world do you answer that question?

Which one should you do and whether it should be a mix?

Those are the very things facing the project that I am working on now.

Not only that, one of the things is we are trying to get prepared for a human mission and we want to reduce the overall cost and reduce the risk.

We would like to bring risk forward from the manned lunar mission into the RLEP program and retire it in the robotics program.

What that means is if I have a technology challenge for a cryogenic engine, that means a liquid oxygen, liquid hydrogen engine for the lunar descent, I want to use that same engine, if I can, on this mission to work out the kinks.

But how much is that worth to NASA?

If that cost an extra \$100 million for RLEP is it worth it?

There is no set answer.

And so, what we are doing right now is what is called a pre-phase A.

Much like I showed you on Chandra where you had that initial design and a final design, we propose to headquarters with one design.

It happened to have a Rover.

It was nuclear powered.

We went into the crater.

And we stayed there for quite a while.

And we definitely characterized what was there.

But it was expensive.

Is that the right answer?

We don't know.

We are right now going through a pre-phase A that we are looking at other options.

We are looking at kind of three classes of sizes.

Different solutions to get into the crater.

Different solutions on the size of the lander.

Its extensibility to the human mission.

Like, for instance, right now we use Russian progress vehicles to resupply Space Station.

We hope to design a robotic lander that would do the same thing during human missions.

It is not man rated, it won't fly men but it can fly supplies to the Moon with one design.

But that cost money because now I am not optimizing design for a point solution mission.

So during this timeframe there are a lot of unknown questions.

And the answers are not straightforward.

So what we are doing is we are giving a series of options and rationale and capabilities for those options to headquarters.

And we will make a recommendation as to what we think is the best answer for NASA to go forward with this design and here are the reasons.

But it could very well be that they choose something different because there is no definitive solution to this problem or this challenge, if you will.

So it is very challenging.

It is going to look probably a lot different now than it will when we actually launch it in 2010 or 2011, but it is something that we face now.

And it is more defined by how much money you can use versus what the requirements are.

Sometimes the requirements do set pretty much the basic tenets of the mission, but other times it is driven by cost.

And so, right now we are in the middle of working that interesting and challenge problem for headquarters to develop what we think that mission should look like.

And, as you can envision, you don't really have requirements.

You have a few requirements.

You have a desirability to reduce the overall cost and reduce the overall risk of a human mission.

But that desirability comes with a cost.

The more extensible you get to the human mission, yes, you can make a big lander that gives you 2,000 kilograms of payload, but it is going to cost you more.

I can make 1,000 kilograms of payload or 500 kilograms of payload and get the definition of what is in the ice.

But then later on I will have to design a new lander for getting other things accomplished that I want to get to on the Moon.

And is that the right answer?

You can tell that there are a lot of challenges with that?

Any questions?

We talked about the cost, performance and schedule triangle.

Do you have a sense, inside NASA, of how much they are pushing schedule at this point/ Where is your flexibility?

In their solicitation to the centers, they allowed schedule flexibility.

They allowed cost flexibility.

Other programs I will say are generally more bounded than this one.

This one is a matter of NASA not knowing exactly what the reason and what the purpose of the program is for.

You've got this general idea of going back to the Moon, using robots first to prepare for humans, but exactly what does that mean and exactly how do you prepare for humans and how much is it going to cost and when is it going to launch and how many missions are you going to have?

They are all undefined.

Now, what we would do is we have already gone to the program that will fund the human missions to say what do you guys plan to do and what do you want to do when you get there?

So we can take their mission design and say how can we help?

We can help here.

We can help here.

We can pull this technology forward and fly it.

The problem there is that they haven't thought that far ahead.

They don't know exactly what they are going to do.

And a lot of that is, is there ice there or not?

If there is ice there, we will want to process it.

We will want to extract oxygen from it.

As part of our solution, we do know that we have to answer that question.

We also do know that we do have to demonstrate precision landing.

But, above that, they have given us a target date of 2010 to 2011.

But they said schedule is not fixed.

They have given us a cost generally between \$400 and \$750 million.

And, by the way, if you can leverage areas of the budget to help your mission, do that, in particular technology.

There is a separate technology budget.

If I can get technology to pay for a new technology development at no cost to the program then I will do that.

And, again, if we come in and say we have \$1 billion but it buys you \$3 billion when the crew is ready to come, we think that is worth the investment.

And then headquarters has to scratch their head and figure out whether it is really worth minimizing the long-term cost versus the yearly constraint that they are given by Congress and the pressures of Shuttle and the pressures of Stations on the overall cost of the program to NASA.

It is real challenging.

It is a lot of fun.

We are going to have a blast.

The next one is going back to the Moon, and you will probably hear more about it in the next couple of years.

OK.

Thank you.

[APPLAUSE] This was really the last of the lectures on systems engineering per se.

Thursday is the last external lecture.

Gordon Fullerton will be talking about test flying the Shuttle.

And hopefully by then I will have the schedule for the presentations next Tuesday.