Intelligent Design and Intelligent Failure

Introduction
Good Evening, my name is Greg Jerman and for nearly a quarter century I have been performing failure analysis on NASA's aerospace hardware. During that time I had the distinct privilege of keeping the Space Shuttle flying for two thirds of its history. I have analyzed a wide variety of failed hardware from simple electrical cables to cryogenic fuel tanks to high temperature turbine blades. During this time I have found that for all the time we spend intelligently designing things, we need to be equally intelligent about understanding why things fail. The NASA Flight Director for Apollo 13, Gene Kranz, is best known for the expression "Failure is not an option." However, NASA history is filled with failures both large and small, so it might be more accurate to say failure is inevitable. It is how we react and learn from our failures that makes the difference.

Engineering Evolution
Before I go forward into space, I would like to set the way back machine to ancient Egypt: 4600 years ago. Egyptian king Sneferu was known for building two pyramids. The first was known as the Bent Pyramid. As construction proceeded, his architects realized the walls were too steep and unstable. So they reduced the angle to prevent collapse. Lessons learned from the Bent Pyramid were employed in his next creation, the Red Pyramid. This was a true smooth sided pyramid. The failures and successes of Sneferu informed his son Khufu who built one of the 10 wonders of the ancient world: the Great Pyramid at Giza. The success of new engineering endeavors are always linked to lessons learned from previous engineering failures. Failure analysis is engineering's evolutionary mechanism.

Engineering Reliability
Now lets fast forward to the modern era to look at the something everyone can relate to: a car. A typical car contains a few thousand moving parts that are located in the engine, transmission, electric windows, and air conditioning. It is annoying if your window won't roll down. It is a bad day if a piston rod fractures and your engine quits.

The Space Shuttle had about 2.5 million moving parts. Having a component failure on a liquid rocket engine could lead to an explosion, loss of vehicle, and loss of crew. A really bad day.

Reliability of an engineering system is based on the reliability of its constituent parts. The more parts, the greater likelihood of failure. In the aerospace industry, we rely on the use of extremely high reliability parts since there are so many of them and failure of one could be catastrophic.

Perception and Reality

I am now going to introduce you to a simple yet critical part of the Space Shuttle. It is a cable that fired explosive bolts that initiated the separation of the Solid Rocket Boosters from the Shuttle. These cables were reused for many years, until one failed to fire its respective bolt. Fortunately, most critical systems have a backup, so in this case, the backup fired and safely separated the booster. As you can see in the x-ray picture, the cable broke in the conductor itself. How many of you pull a plug out of the wall by pulling on the wire rather than grasping the plug? This is what happens when you pull a plug out by the wire too many times. How could this happen to critical flight hardware? Flight hardware is extensively tested and handled with kid gloves, but a subtle change in perception led to unintended consequences. After an SRB was recovered, all the cables were removed, but the technicians considered everything "flown" hardware so they were not very careful about removal. They thought the rigorous pre-flight testing, before reuse, would weed out any problem cables. They didn't count on the elastic nature of the rubber coating pulling the broken wires together. This gave a positive electrical continuity check, and a failure at 2Gs of acceleration on launch. A simple shift in perception of flight versus flown hardware could have been disastrous.

An Aging Space Shuttle

As the Space Shuttle began its second decade of service, a number of issues arose related to its age. Each Shuttle had been designed to fly 100 times, but some components were failing early. One such component was used in the liquid oxygen feed lines to allow for thermal expansion and contraction. In the spring of 2002, a routine inspection of the liquid oxygen feed lines on the Space Shuttle Discovery identified a cracked ball in a Ball Strut Tie Rod Assembly (BSTRA) joint. This ball was about two inches in diameter and was made from an alloy known as Stoody #2: a cobalt matrix hardened by chromium and tungsten carbides. It is a very high temperature alloy used at very low temperatures. Why? It had been used in similar joints on the Saturn V rocket, so it was accepted as a heritage material for the new Space Shuttle. The cracked BSTRA ball from Discovery had actually been sand cast in 1978, so at the time of failure, it was 24 years old. The ball failed because in had sand inclusions that created a small crack that opened with successive thermal shocks on cool down in liquid oxygen. During the investigation we found three different casting methods were used and there were three different ball sizes. So to determine the safety of all BSTRA balls used in the Space Shuttle, a series of thermal shock tests were conducted to see if we could get any other balls to crack. We also conducted stress tests to see if we could get a ball to completely break apart. We worked through Christmas into the New Year to prove the safety of the BSTRA balls. With any flight that occurs after a significant failure analysis activity, we always have heightened anxiety. On January 16, 2003 we were relieved when the Shuttle Columbia was successfully launched. The feed lines and engines worked flawlessly.

Columbia Disaster

Sixteen days later, Columbia crashed after reentering Earth's atmosphere. Everything we had done to safely launch Columbia had become irrelevant. It was a Saturday, and I didn't know what to do. I wanted to immediately drive to Texas to help in the debris recovery, but that wasn't my job. My job was to wait for enough debris to be collected in order to start some meaningful analysis. One post launch concern had been a large impact from insulating foam off the External Tank. It hit orbiter's wing leading edge, but foam had been hitting the orbiter since STS-1. It was not a significant concern. During reentry, initial alarms indicated a problem with the landing gear before loss of communication. So our first focus was on the landing gear. It wasn't until the discovery of Columbia's flight data recorder that extra sensor data became available. Columbia was the first Orbiter to fly in space, so it had been wired with extra sensors connected to a data recorder. This computer was not hardened like a passenger airliner's black box, so it was a miracle any data survived. The first thermal sensor to show a temperature anomaly was in the left wing's leading edge. This was where the loose insulating foam hit during launch. So how do you prove foam caused damage when it would have all burned away? It took three months to collect enough wing leading edge debris to start analysis. We sectioned molten deposits at different locations, and compared their chemical analysis to the alloys used in the wing leading edge. The results were clear. Molten metal and insulation went down in layers. We found the molten remnants of attachment hardware for the wing leading edge panels at only one location on the left wing. This pinpointed the location where hot gas first crept into the left wing. Modeling of the foam impact trajectory coincided with this location and subsequent foam impact testing showed high velocity foam of a certain size could damage the brittle high temperature wing leading edge material. Unknown amounts of risk had been regularly accepted as foam continued to hit the Orbiter on each and every flight. The loss of Columbia taught us that past success did not guarantee future success.

Intertank Stringer Failure

After Columbia, NASA decided to set a retirement date for the Shuttle program. We would finish building the International Space Station and retire the fleet by 2010. We came close to meeting that deadline until a unique failure halted the launch schedule in November 2010. Stringers used on the intertank structure of the Space Shuttle External Tank had cracked. The damage was only found after a launch scrub, and inspection of the External Tank revealed protruding insulation foam. The construction of the intertank structure used standard skin-stringer construction methods employed in aircraft. Metallurgical and process evaluations found the material used in the construction was within specification. Since the material and construction process were good, why were some stringers cracking? The answer lay in defining a box around the material properties. The specification called out a minimum strength value, but processing variability led to some very strong stringers. As materials get stronger, they generally become more brittle. In the case of the intertank stringers, some were too strong which reduced their fracture toughness. Thermal contraction during the tank filling operation was bending the ends of the stringer feet resulting in cracking of a few stringers. The fix was very simple, although too late to affect the Space Shuttle program. Future intertank stringers would have a maximum as well as a minimum strength requirement.

Final Thoughts
So if I were to summarize my career in NASA failure analysis, I would say there are three critical concepts that are common in the resolution of a failure.

The first is that a person's perception is their reality. Each person has their own view on how they perceive their job. We all view the world through the lens of our experience. The complex aerospace systems we build are touched by the input of tens of thousands of folks with hardware manufactured by hundreds of companies. Reconciling all these differences is a challenge. Perception is neither right nor wrong. It just exists. The trick is melding perceptions so they are complementary rather than in opposition.

The second is challenging assumptions. In an increasingly technological world, we make and live by many assumptions that allow us to simplify our existence. In a perfect world all our assumptions would be valid and nothing would fail. So when hardware fails, obviously one or more of these assumptions are no longer valid. Challenging assumptions can be difficult because, by their nature, they are unseen and ignored by the very people who hold them. However they are critical to understanding the framework we construct around our materials and processes. When you can positively identify invalid assumptions, you are on your way to actually understanding why things fail.

The last concept is defining and living inside a box. How often are you told to think "outside the box?" To an aerospace engineer, living outside the box generally means failure. Our rockets are only as good as the materials and processes we use to construct them. We are always pushing the performance envelope with higher pressure, higher temperature, and higher stress. To successfully fly we must be very specific about the limits we impose as we are always flying up to the edge of the precipice. It can be thrilling to stand on the edge of a cliff and look out into the abyss. It is a lot safer if there is a railing to lean on so we don't lean too far.

Early Example of Failure Analysis

Bent Pyramid

Red Pyramid

Failure of a Booster Separation Ordnance Cable

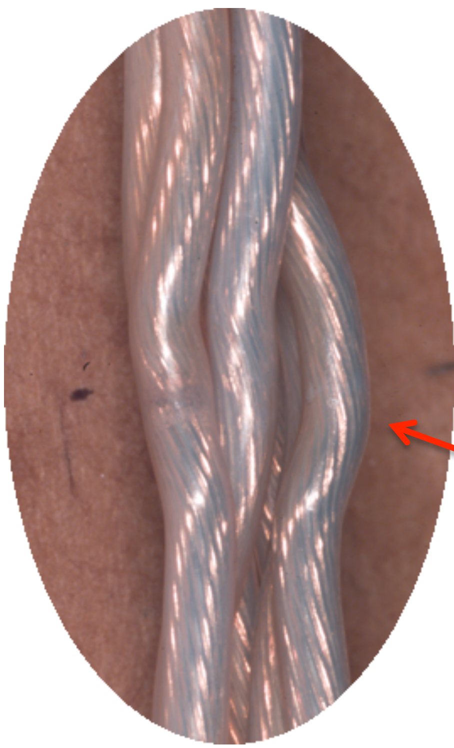

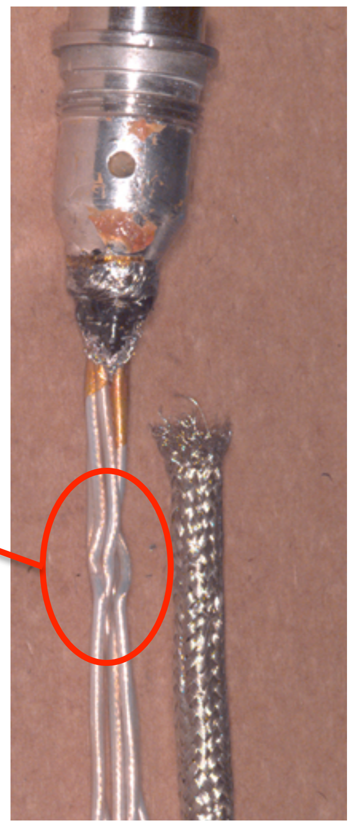

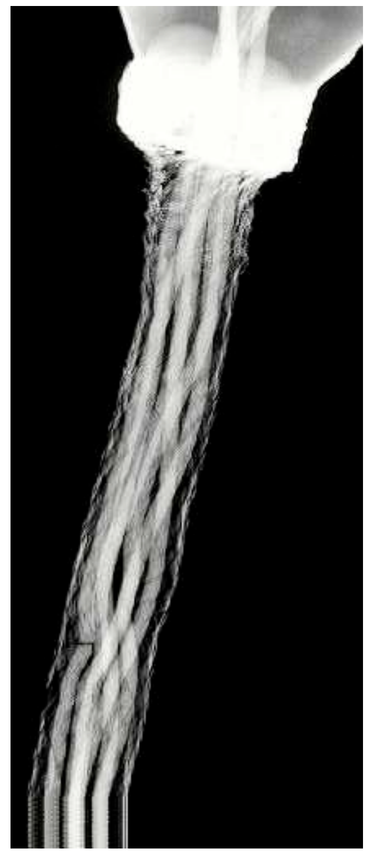

Cracked Ball from Orbiter Ball Strut Tie Rod Assembly

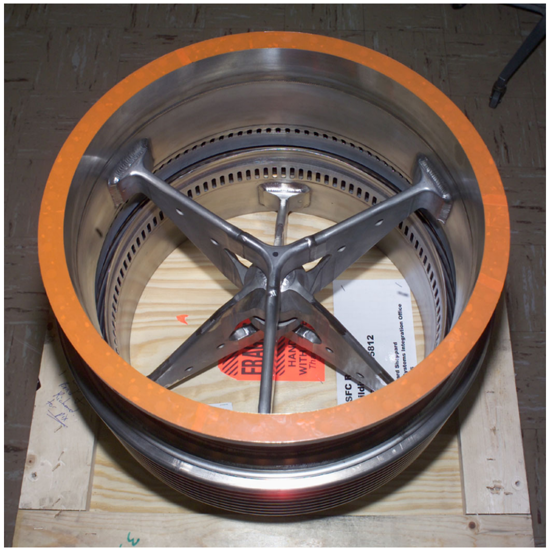

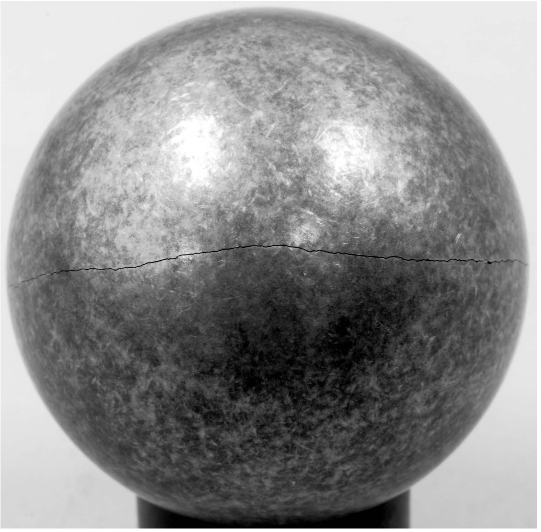

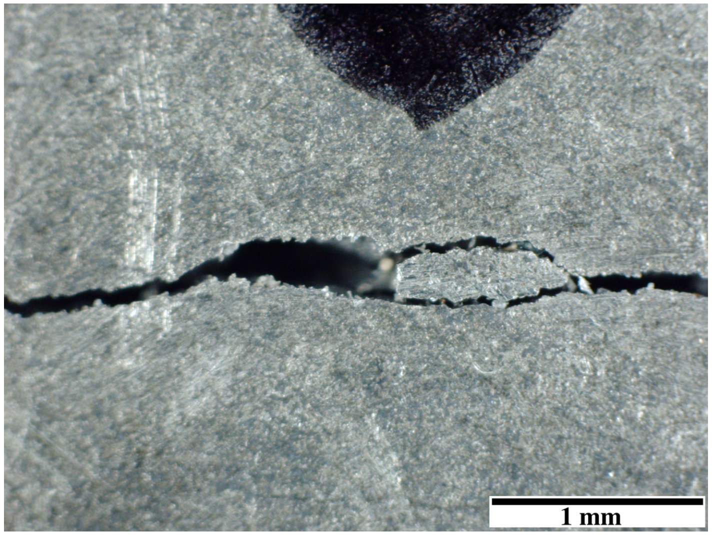

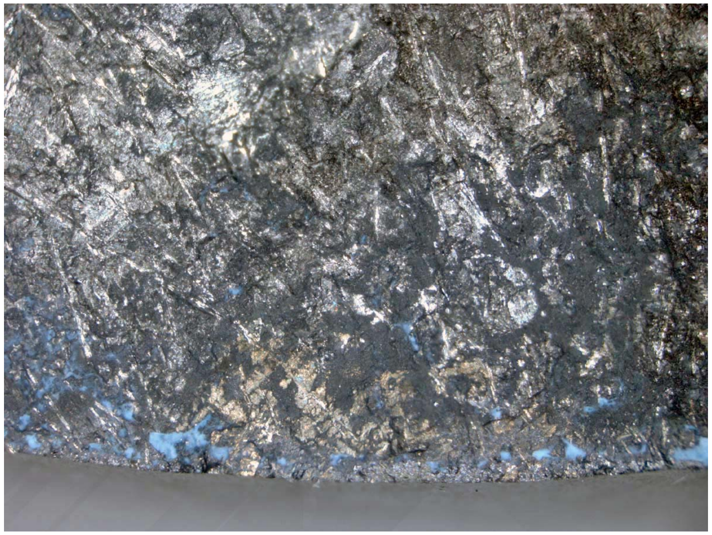

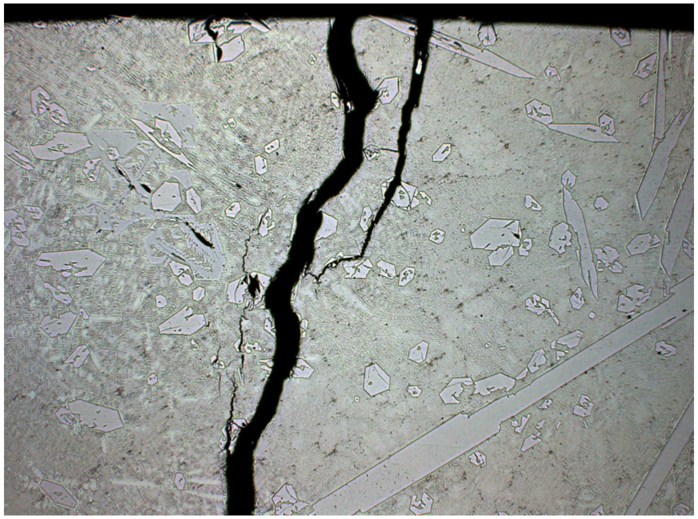

Columbia Investigation

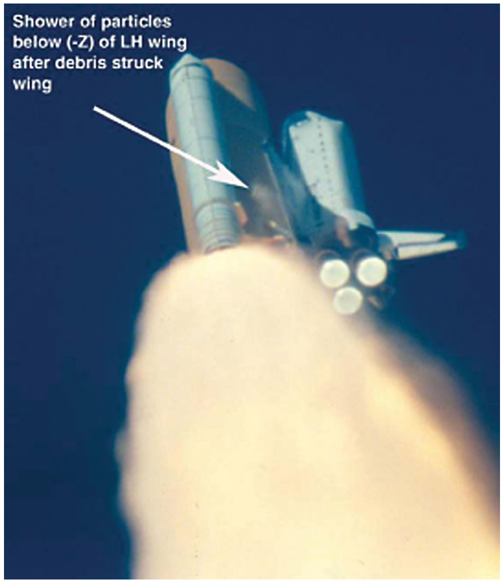

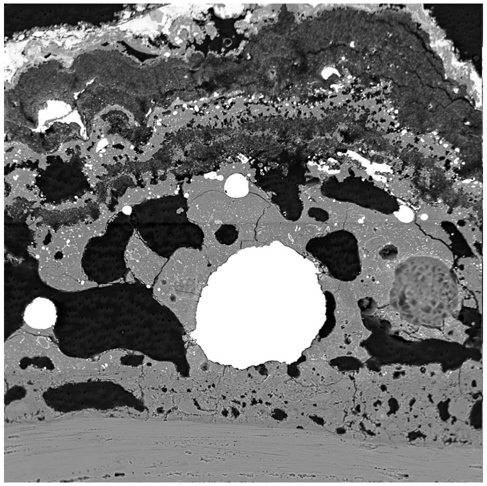

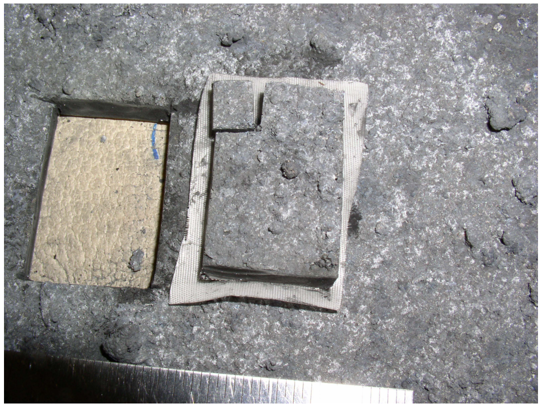

External Tank Stringer Cracking

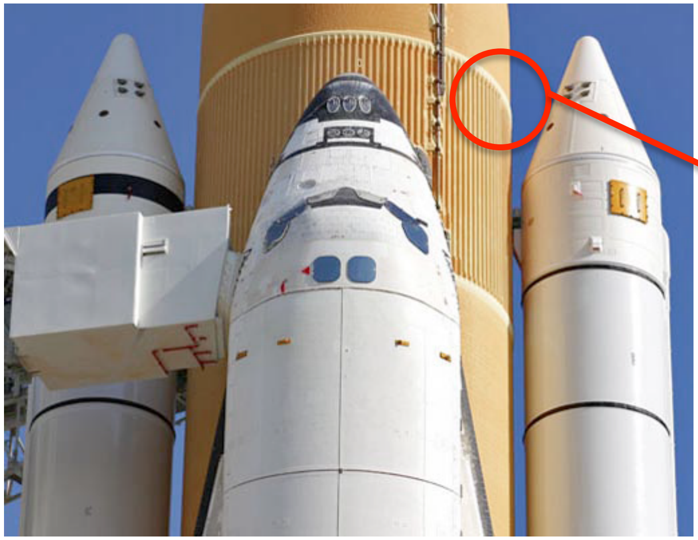

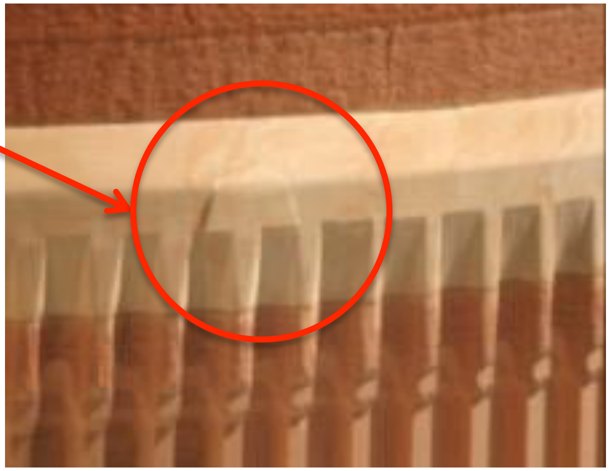

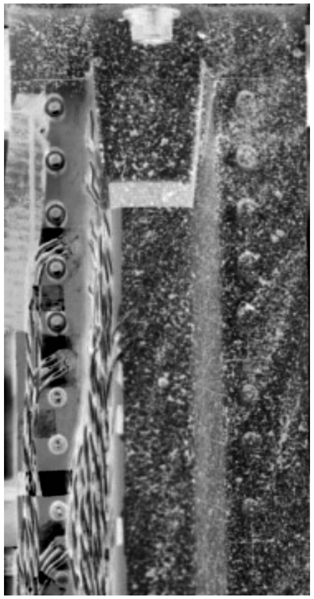

Fractured
Stringer Feet