National Aeronautics and Space Administration

SYSTEM FAILURE CASE STUDIES

MARCH 2010 *VOLUME 4 ISSUE 3*

Island Fever

On the morning of March 28, 1979, a combination of tightly coupled equipment malfunctions and signal misinterpretations led to an accident at the nuclear power plant known as Three Mile Island. The incident resulted in no fatalities and no proven damage on the outlying community of Dauphin County, PA. However, it is recognized as the worst civilian nuclear accident in the U.S., and caused drastic overhauls in government regulations for nuclear power plant operations.

BACKGROUND

The Nuclear Power Plant

The Three Mile Island Nuclear Generating Station, commonly referred as Three Mile Island (TMI), is a nuclear power plant located near Harrisburg, PA. It is comprised of two pressurized water reactors, referred to as TMI-1 and TMI-2. At peak operating capacity, each unit could produce about 900 million watts of electricity, enough to power almost 1 million homes.

Nuclear Reactor: How it Works

Figure 1 shows a brief overview of how nuclear reactors work. First, heat is generated as atoms split inside a nuclear **core**. The core in TMI-2 is 12 feet tall and weighs 100 tons. To control the rate at which atoms split and the resulting heat, **control rods** are raised and lowered into the core. More heat is generated when the control rods are raised, and less when they're lowered.

The **primary loop** pumps water through the core, absorbing heat generated by the nuclear fission process. In addition to being radioactive, water in the primary loop stays at very high temperatures, and must be pressurized to keep from turning into steam. The primary system has a pressurizer to control the pressures in the primary loop. The pressurizer has a pilot operated relief valve (**PORV**) that opens to release excess pressure from the primary loop. In the case of "loss of coolant" from the primary loop, an emergency core coolant system (**ECCS**) injects water into the system to cool down the primary loop. To contain the radioactive material, all of these components are surrounded by a **containment structure**, which is a four-foot-thick concrete encasement.

The **secondary loop** carries water that is heated by primary loop through a heat exchanger. Though the two loops are next to each other, the water in the two loops never comes in direct contact. This prevents water in the secondary loop from becoming radioactive. As water moves through the

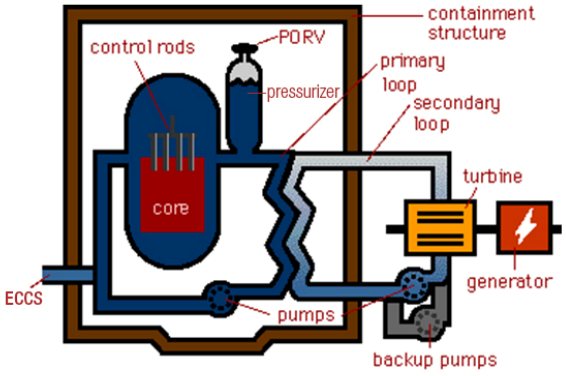

Figure 1: The nuclear reactor uses nuclear power to turn water into steam, generating electricity.

secondary loop, heat from the primary loop turns secondary loop water into steam. This steam goes into a **turbine** and expands, turning a **generator**. This generator then produces electricity. After moving through the turbine, the steam condenses back into water and continues its cycle through the secondary loop.

WHAT HAPPENED?

Mechanical Failure

At 4:00 a.m. EST on March 28, 1979, a mechanical failure occurred in the secondary loop, causing the pumps to stop running. This failure triggered two events: (1) the turbine shut down and (2) the control rods lowered into the core to shut down the nuclear reaction. Both procedures were in line with proper plant operations. When the main pumps shut down, the backup pumps activated automatically. However,

Mechanical failures spark nuclear meltdown in Three Mile Island power plant.

Proximate Causes:

- Mechanical failures caused malfunctions within the radioactive system
- Operators misinterpreted warning signs and left an important valve open
- Radioactive materials flowed out of the valve, causing a hazardous material containment breach

Underlying Issues:

- Mechanical Failures
- Errors in Interpreting Warning Signs
- Unclear Control Indicators

the valves in the backup pump had been closed during routine maintenance two days prior, and were not re-opened, as they should have been, violating Nuclear Regulatory Commission (NRC) rules. The closed valves prevented the backup system from pumping water. This stopped the secondary loop water flow completely, also stopping the cooling for the primary loop.

While the nuclear reaction was shut down, residual heat from the core built up inside the primary loop. Since the secondary loop was no longer working to remove this heat, the temperature and pressure rose inside the primary loop. In response, the PORV automatically opened to release pressure. After the pressure fell to normal levels, the PORV should have closed again automatically, but a second mechanical failure caused the valve to remain open.

Leak of Nuclear Materials

In the control room, operators did not realize that the valve was still open. As water from the primary loop flowed out the stuck-open PORV, pressure inside the loop continued to drop. When it became too low, high-temperature water boiled and turned into steam. This steam and boiling water caused the pressure to increase in the primary loop, even though the valve was still open.

On the control board, the pressure gauge showed high readings. Since operators assumed the PORV valve was closed, they concluded that the high pressure came from excess water in the primary loop. They believed the core was overflowing, and turned off the ECCS. In reality, the injected water was escaping from the PORV, steam was building up inside the core, and the core was experiencing what's known as a loss of coolant accident (LOCA). Turning off the ECCS only worsened the situation, and the core continued to overheat. The steam build-up in the core increased pressure and pushed more radioactive water through the PORV into a connecting tank. This tank overflowed and ruptured at 4:15 a.m. in the morning. Radioactive water began to leak into the general containment building. In the core, rising temperatures and excess steam caused the primary loop pumps to cavitate and vibrate excessively. To prevent damage, plant operators turned off the pumps. Water stopped circulating in the primary loop and, as a result, was converted to steam. The steam continued to increase pressure in the core, pushing more coolant out of the PORV.

Serious Damage Occurs

Around 6:00 a.m., the core became exposed to the intense heat and steam building up in the primary loop. The fuel rods reacted with the steam, melting exposed portions of the core (Figure 2). As the core melted, more radioactive material was released to the coolant, producing hydrogen gas bubbles in the core. This hydrogen bubble was discovered later, and became a serious concern because it prevented water from flowing through the core. At 6:00 a.m., there was also a shift change in the control room. The new shift noticed that the temperatures in the PORV pipe and the holding tanks were too high. They correctly realized the system was experiencing a loss of coolant, and shut the valve. However, by this point, 32,000 gallons of radioactive water had already spilled out of the primary loop.

Around 6:50 a.m., radiation alarms began to ring, signaling excess levels of radiation in the containment system. By this point, radiation levels had reached 300 times the expected values. Ten minutes later, a site emergency was declared. NRC regional and national offices were notified, and the DOE and EPA were also alerted. By 11:00 a.m., all non-essential personnel were evacuated from the premises, and Pennsylvania Governor Richard Thornburgh advised pregnant women and pre-school-aged children within a five-mile radius of the plant to leave the area.

After the emergency was declared, personnel immediately began working to cool the system. However, the pockets of gas that had built up in the loop prevented normal water flow. Operators also noticed the hydrogen bubble in the core, and had to take extra care to diffuse it safely throughout the week. Pockets of steam and hydrogen were slowly released, and after 16 hours of damage control, the primary loop pumps were turned on once again, and the core temperature began to fall. To shut down the plant completely, operators had to wait for the residual heat in the core to decrease to the point where the coolant water pumps could be turned off. This was finally accomplished a month after the crisis, on April 27. After a lengthy shut-down process, TMI-2 was permanently deactivated.

Proximate Cause

The accident started when pumps in the secondary loop of the nuclear reactor system stopped working. Water in the

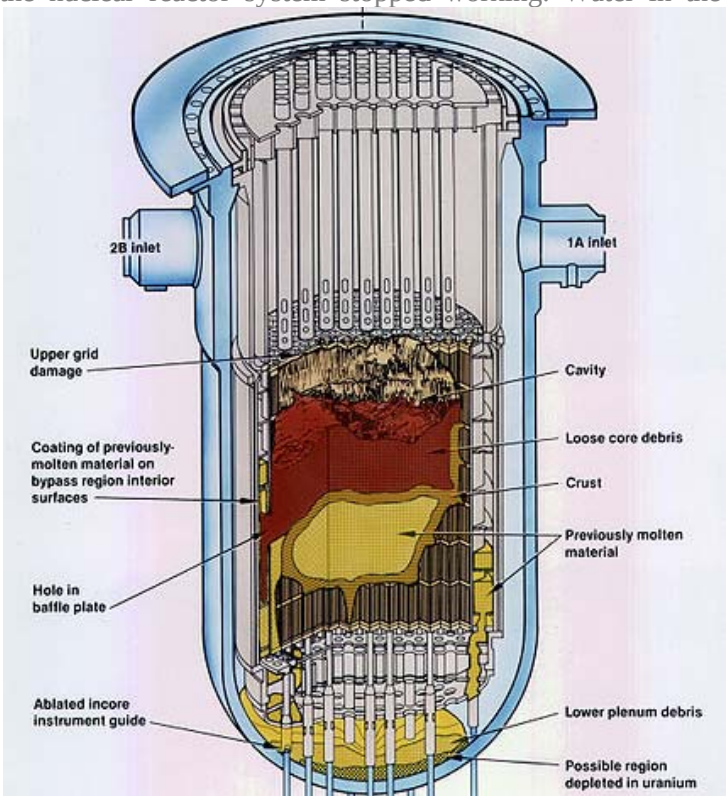

Figure 2: Illustration of the end-state of the core. Half of the core had been melted from the loss of coolant.

secondary loop could no longer remove heat from the radioactive primary loop. As heat and pressure built up in the primary loop, the PORV opened to relieve the pressure. A mechanical failure left the PORV open. Radioactive water flowed out of this open valve into a tank, which burst, leaking water into the containment unit.

Meanwhile, unpressurized water in the core turned to steam, severely reacting with the core and further raising primary loop radioactivity levels. Leakages into the containment building posed a contamination problem to the facility and the surrounding areas.

Underlying Issues

Mechanical Failures

Two failures occurred over the course of this accident. The first failure was the mechanical problem that caused the pumps in the secondary loop to stop working. This should have been easily fixed, since there was a system of backup pumps meant to take over in this situation. The second, and more significant failure, was the PORV failing to re-close after the pressure had decreased to acceptable levels. Within the countless parts of the power plant, these two failures caused a domino effect throughout the system.

Errors in Interpreting Warning Signs

The secondary loop had a set of backup pumps to take over should the main pumps fail. During a routine maintenance check two days before the accident, isolation valves were closed for testing. After the maintenance check was complete, the operator forgot to re-open the valves. This action violated NRC rules, since it prevented the backup pumps from operating properly.

As the accident progressed, plant operators in the control room relied on a PORV indicator light to tell them whether the valve was open or closed. Additionally, the operators were nearing the end of their all-night shift, which had an impact on accurately interpreting the many dials and figuring out the complex series of malfunctions that had taken place.

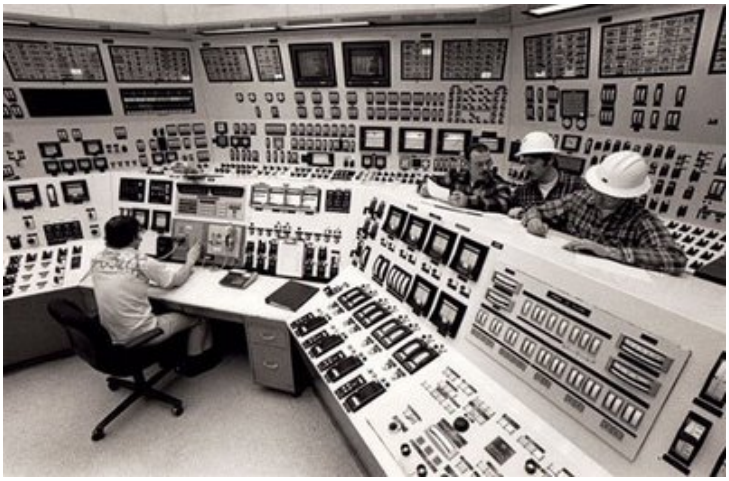

Figure 3: A typical nuclear power plant control room. Note the numerous arrays of dials, controls, and indicators.

The dials showed that pressure inside the reactor was falling, while pressure inside the primary loop was rising, two seemingly conflicting facts since the two are connected. Because the series of malfunctions that had occurred were very unlikely and were not part of normal plant operation, the operators didn't think to check the PORV or the valves, and natural assumed that one of the dials was just wrong.

Additionally, the operators had also been trained to never allow pressure in the primary loop to climb too high, to prevent excessive leaks and release radioactivity. The high pressure in the primary loop prompted the operators to turn off the ECCS to avoid "going solid," or flooding the unit with too much water and increasing pressure further. Though the operators believed they were taking safety precautions, they were unknowingly making the situation worse.

Unclear Control Indicators

The systems in a nuclear power plant operate with numerous backup systems to check and control performance. The mechanical failures at TMI were not complicated, but were difficult to solve because of the intricate ways in which the components were connected.

The key to this incident was the unclear nature of the PORV indicator. The indicator light was linked to the power line that powered the PORV, not the PORV itself. This light did not give direct indication about whether the PORV itself was open or closed; it only told whether the power source to the PORV was on or off. Even though the power source had shut off and the PORV should have closed, it actually remained stuck open, without indication.

The operators were not trained to look for other clues that could have indicated that the valve hadn't closed; they were used to the assumption that, if the light is off, the valve was closed. As Figure 3 shows, the control boards offered no easy answers to the problem, and the operators received seemingly conflicting signals from the gauges. It was not until a fresh shift came in at 6:00 a.m. that a new set of eyes put the pieces together to pinpoint the problem.

Another shortcoming in the plant instrumentation was the lack a direct indication of water level in the core. The operators had to rely on information from temperature and pressure gauges to estimate the water level.

In light of all these plant intricacies, the tightly coupled interactions made the problem difficult to solve. It was only a matter of time before these malfunctions would occur like a rolling snowball, compounding into a catastrophic event.

Aftermath

After clean-up, the facility was defueled and shut down, a process that cost nearly a billion dollars and lasted years. Though the general population was rightly concerned with the health effects, later studies showed that very little radioactive material leaked from TMI-2. The exposure to this material proved to be equal to 1/6 of the radiation exposure in a standard chest X-ray.

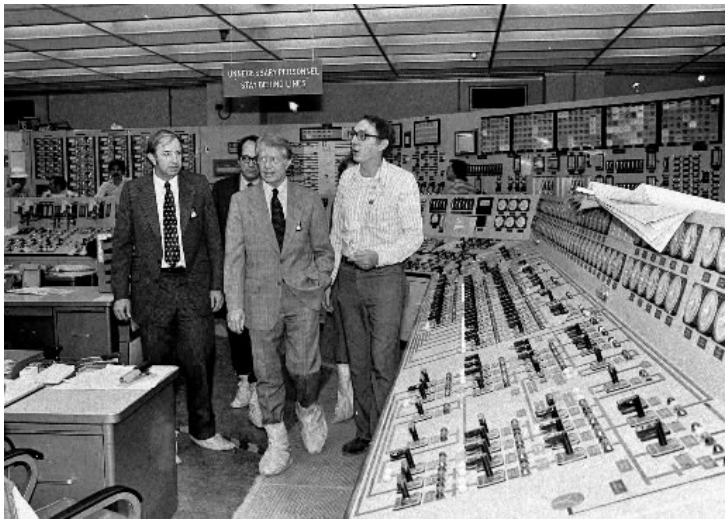

Figure 4: The accident sparked national concern and prompted a visit from President Jimmy Carter (center).

This incident caused many sweeping changes in the NRC's rules and regulations. The NRC made drastic changes to how it regulates its licenses. Plant operation and management also became more closely watched. A new emphasis was placed on staff training and performance, and both are now viewed as critical in analyzing plant safety. Control room layout, and instrumentation and displays were also upgraded.

Finally, subsequent to the TMI incident, all valves had to have physical limit switches that showed an actual valve status, as opposed to showing the "commanded" or "inferred" position. Similarly, thermocouples were added to piping downstream of the PORV exhaust, so that plant operators could determine, with a redundant measure, what was really happening within the system.

For Future NASA Missions

TMI-2 is a classic example of a complex operation in which a couple of small hiccups were lost in an intricate system, going unnoticed and resulting in near-disaster. NASA's operations are similarly complex. Much is on the line, and employees are extremely important in interpreting data and looking for potentially obscure signals that, if missed, can lead to serious consequences.

The experience at TMI-2 highlights the drastic impacts of minor mishaps, especially in an environment with numerous interactions and hazard potentials. Small accidents are always a possibility, but operator actions are crucial to preventing the escalation of these accidents. Machines in themselves are not perfect creations; often the message and data they generate needs to be closely analyzed.

NASA's machines and processes also interact with each other in complicated ways, and as with TMI-2, it can be nearly impossible to accurately point the one component that is broken, out of thousands. With repetition, the conduct of tasks also tends to rely more and more on assumptions.

Though assumptions can get the job done faster, sometimes they prove to be completely wrong, and blindly following them can have damaging results.

Questions for Discussion

- Are indicators, alarms, and warnings for critical system parameters clear? Are they as closely aligned to the actual measure as practical?
- What decision-making processes does your organization follow? How do you protect against misinterpretation?
- What potential failures could your project experience? How might you guard against them?
- Are there areas in which your backup procedures could be strengthened? Is there a proper level of redundancy?

In high-risk situations, nothing is minor, and attention to detail can mean the difference between a great success and an utter catastrophe. Though malfunctions happen, employees' eyes and ears are essential in taking precautions to prevent these malfunctions from happening within NASA.

References:

"Aging Reactors a Worry for Regulators," Ohio Citizen Org., 2004. Retrieved from www.ohiocitizen.org/campaigns/ electric/2004/nuc2004c.html.

Backgrounder on the Three Mile Island Accident, United States Nuclear Regulatory Commission. Retrieved from http://www.nrc.gov/reading-rm/doc-collections/fact-sheets/3mile-isle.html.

Kemeny, John G. (October 1979). Report of The President's Commission on the Accident at Three Mile Island: The Need for Change: The Legacy of TMI. Washington, D.C.: The Commission. ISBN 0935758003.

NRC Image of TMI-2 Core End Stable State. Retrieved from www.nrc.gov/images/reading-rm/photo-gallery/20071114-006.jpg.

Perrow, Charles. Normal Accidents. Princeton: Princeton University Press, 1999.

"Plant Operations," Three Mile Island Info, 2010. Retrieved from www.threemileislandinfo.com/reliability/plant-operations.aspx.

"Three Mile Island: How a Nuclear Reactor Works," American Experience, PBS, 2009. Retrieved from http://www.pbs.org/wgbh/amex/three/sfeature/tmihow.html.

SYSTEM FAILURE CASE STUDIES

NASA SAFETY CENTER

Executive Editor: Steve Lilley steve.k.lilley@nasa.gov
Developed by: ARES Corporation
Thanks to Mr. Peter Prassinos for his insightful peer review.