

NORMAL ACCIDENTS

Living with High-Risk
Technologies

*With a New Afterword and a
Postscript on the Y2K Problem*

CHARLES PERROW

Princeton University Press
Princeton, New Jersey

Copyright © 1999 by Princeton University Press
Published by Princeton University Press, 41 William Street,
Princeton, New Jersey 08540
In the United Kingdom: Princeton University Press,
Chichester, West Sussex

All Rights Reserved

First published by Basic Books 1984

Library of Congress Cataloging-in-Publication Data

Perrow, Charles.

Normal accidents : living with high-risk technologies / Charles Perrow
p. cm.

Reprint. Originally published: New York : Basic Books, 1984.
includes bibliographical references.

ISBN 0-691-00412-9 (pbk. : alk. paper)

1. Industrial accidents. 2. Technology—Risk assessment.

3. Accidents. I. Title.

HD7262 .P55 1999

363.1—dc21 99-32990

The paper used in this publication meets the minimum requirements of
ANSI/NISO Z39.48-1992 (R1997) (*Permanence of Paper*)

<http://pup.princeton.edu>

Printed in the United States of America

11 13 15 17 19 20 18 16 14 12

ISBN-13: 978-0-691-00412-9 (pbk.)

ISBN-10: 0-691-00412-9 (pbk.)

Contents

Abnormal Blessings	vii
Introduction	3
1. Normal Accident at Three Mile Island	15
2. Nuclear Power as a High-Risk System: Why We Have Not Had More TMIs—But Will Soon	32
3. Complexity, Coupling, and Catastrophe	62
4. Petrochemical Plants	101
5. Aircraft and Airways	123
6. Marine Accidents	170
7. Earthbound Systems: Dams, Quakes, Mines, and Lakes	232
8. Exotics: Space, Weapons, and DNA	256
9. Living with High-Risk Systems	304
Afterword	353
Postscript: The Y2K Problem	388
List of Acronyms	413
Notes	415
Bibliography	426
Index	441

1715 PD 983352

INTRODUCTION

Welcome to the world of high-risk technologies. You may have noticed that they seem to be multiplying, and it is true. As our technology expands, as our wars multiply, and as we invade more and more of nature, we create systems—organizations, and the organization of organizations—that increase the risks for the operators, passengers, innocent bystanders, and for future generations. In this book we will review some of these systems—nuclear power plants, chemical plants, aircraft and air traffic control, ships, dams, nuclear weapons, space missions, and genetic engineering. Most of these risky enterprises have catastrophic potential, the ability to take the lives of hundreds of people in one blow, or to shorten or cripple the lives of thousands or millions more. Every year there are more such systems. That is the bad news.

The good news is that if we can understand the nature of risky enterprises better, we may be able to reduce or even remove these dangers. I have to present a lot of the bad news here in order to reach the good, but it is the possibility of managing high-risk technologies better than we are doing now that motivates this inquiry. There are many improvements we can make that I will not dwell on, because they are fairly obvious—such as better operator training, safer designs, more quality control, and more effective regulation. Experts are working on these solutions in both government and industry. I am not too sanguine about these efforts, since the risks seem to appear faster than the reduction of risks, but that is not the topic of this book.

Rather, I will dwell upon characteristics of high-risk technologies that suggest that no matter how effective conventional safety devices are, there is a form of accident that is inevitable. This is not good news for

NORMAL ACCIDENTS

systems that have high catastrophic potential, such as nuclear power plants, nuclear weapons systems, recombinant DNA production, or even ships carrying highly toxic or explosive cargoes. It suggests, for example, that the probability of a nuclear plant meltdown with dispersion of radioactive materials to the atmosphere is not one chance in a million a year, but more like one chance in the next decade.

Most high-risk systems have some special characteristics, beyond their toxic or explosive or genetic dangers, that make accidents in them inevitable, even "normal." This has to do with the way failures can interact and the way the system is tied together. It is possible to analyze these special characteristics and in doing so gain a much better understanding of why accidents occur in these systems, and why they always will. If we know that, then we are in a better position to argue that certain technologies should be abandoned, and others, which we cannot abandon because we have built much of our society around them, should be modified. Risk will never be eliminated from high-risk systems, and we will never eliminate more than a few systems at best. At the very least, however, we might stop blaming the wrong people and the wrong factors, and stop trying to fix the systems in ways that only make them riskier.

The argument is basically very simple. We start with a plant, airplane, ship, biology laboratory, or other setting with a lot of components (parts, procedures, operators). Then we need two or more failures among components that interact in some unexpected way. No one dreamed that when X failed, Y would also be out of order and the two failures would interact so as to both start a fire and silence the fire alarm. Furthermore, no one can figure out the interaction at the time and thus know what to do. The problem is just something that never occurred to the designers. Next time they will put in an extra alarm system and a fire suppressor, but who knows, that might just allow three more unexpected interactions among inevitable failures. This interacting tendency is a characteristic of a system, not of a part or an operator; we will call it the "interactive complexity" of the system.

For some systems that have this kind of complexity, such as universities or research and development labs, the accident will not spread and be serious because there is a lot of slack available, and time to spare, and other ways to get things done. But suppose the system is also "tightly coupled," that is, processes happen very fast and can't be turned off, the failed parts cannot be isolated from other parts, or there is no other way to keep the production going safely. Then recovery from the initial disturbance is not possible; it will spread quickly and irretrievably for at

Introduction

least some time. Indeed, operator action or the safety systems may make it worse, since for a time it is not known what the problem really is.

Probably many production processes started out this way—complexly interactive and tightly coupled. But with experience, better designs, equipment, and procedures appeared, and the unsuspected interactions were avoided and the tight coupling reduced. This appears to have happened in the case of air traffic control, where interactive complexity and tight coupling have been reduced by better organization and "technological fixes." We will also see how the interconnection between dams and earthquakes is beginning to be understood. We now know that it involves a larger system than we originally thought when we just closed off a canyon and let it fill with water. But for most of the systems we shall consider in this book, neither better organization nor technological innovations appear to make them any less prone to system accidents. In fact, these systems require organizational structures that have large internal contradictions, and technological fixes that only increase interactive complexity and tighten the coupling; they become still more prone to certain kinds of accidents.

If interactive complexity and tight coupling—system characteristics—inevitably will produce an accident, I believe we are justified in calling it a *normal accident*, or a *system accident*. The odd term *normal accident* is meant to signal that, given the system characteristics, multiple and unexpected interactions of failures are inevitable. This is an expression of an integral characteristic of the system, not a statement of frequency. It is normal for us to die, but we only do it once. System accidents are uncommon, even rare; yet this is not all that reassuring, if they can produce catastrophes.

The best way to introduce the idea of a normal accident or a system accident is to give a hypothetical example from a homey, everyday experience. It should be familiar to all of us; it is one of those days when everything seems to go wrong.

A Day in the Life

You stay home from work or school because you have an important job interview downtown this morning that you have finally negotiated. Your friend or spouse has already left when you make breakfast, but unfortu-

NORMAL ACCIDENTS

nately he or she has left the glass coffeepot on the stove with the light on. The coffee has boiled dry and the glass pot has cracked. Coffee is an addiction for you, so you rummage about in the closet until you find an old drip coffeemaker. Then you wait for the water to boil, watching the clock, and after a quick cup dash out the door. When you get to your car you find that in your haste you have left your car keys (and the apartment keys) in the apartment. That's okay, because there is a spare apartment key hidden in the hallway for just such emergencies. (This is a safety device, a *redundancy*, incidentally.) But then you remember that you gave a friend the key the other night because he had some books to pick up, and, planning ahead, you knew you would not be home when he came. (That finishes that *redundant pathway*, as engineers call it.)

Well, it is getting late, but there is always the neighbor's car. The neighbor is a nice old gent who drives his car about once a month and keeps it in good condition. You knock on the door, your tale ready. But he tells you that it just so happened that the generator went out last week and the man is coming this afternoon to pick it up and fix it. Another "backup" system has failed you, this time through no connection with your behavior at all (*uncoupled* or independent events, in this case, since the key and the generator are rarely connected). Well, there is always the bus. But not always. The nice old gent has been listening to the radio and tells you the threatened lock-out of the drivers by the bus company has indeed occurred. The drivers refuse to drive what they claim are unsafe buses, and incidentally want more money as well. (A safety system has foiled you, of all things.) You call a cab from your neighbor's apartment, but none can be had because of the bus strike. (These two events, the bus strike and the lack of cabs, are tightly connected, dependent events, or *tightly coupled* events, as we shall call them, since one triggers the other.)

You call the interviewer's secretary and say, "It's just too crazy to try to explain, but all sorts of things happened this morning and I can't make the interview with Mrs. Thompson. Can we reschedule it?" And you say to yourself, next week I am going to line up two cars and a cab and make the morning coffee myself. The secretary answers "Sure," but says to himself, "This person is obviously unreliable; now this after pushing for weeks for an interview with Thompson." He makes a note to that effect on the record and searches for the most inconvenient time imaginable for next week, one that Mrs. Thompson might have to cancel.

Now I would like you to answer a brief questionnaire about this event. Which was the primary cause of this "accident" or foul-up?

Introduction

1. Human error (such as leaving the heat on under the coffee, or forgetting the keys in the rush)? Yes_____ No_____ Unsure_____
2. Mechanical failure (the generator on the neighbor's car)? Yes_____ No_____ Unsure_____
3. The environment (bus strike and taxi overload)? Yes_____ No_____ Unsure_____
4. Design of the system (in which you can lock yourself out of the apartment rather than having to use a door key to set the lock; a lack of emergency capacity in the taxi fleet)? Yes_____ No_____ Unsure_____
5. Procedures used (such as warming up coffee in a glass pot; allowing only normal time to get out on this morning)? Yes_____ No_____ Unsure_____

If you answered "not sure" or "no" to all of the above, I am with you. If you answered "yes" to the first, human error, you are taking a stand on multiple failure accidents that resembles that of the President's Commission to Investigate the Accident at Three Mile Island. The Commission blamed everyone, but primarily the operators.¹ The builders of the equipment, Babcock and Wilcox, blamed *only* the operators. If you answered "yes" to the second choice, mechanical error, you can join the Metropolitan Edison officials who run the Three Mile Island plant. They said the accident was caused by the faulty valve, and then sued the vendor, Babcock and Wilcox. If you answered "yes" to the fourth, design of the system, you can join the experts of the Essex Corporation, who did a study for the Nuclear Regulatory Commission of the control room.²

The best answer is not "all of the above" or any one of the choices, but rather "none of the above." (Of course I did not give you this as an option.) The cause of the accident is to be found in the complexity of the system. That is, each of the failures—design, equipment, operators, procedures, or environment—was trivial by itself. Such failures are expected to occur since nothing is perfect, and we normally take little notice of them. The bus strike would not affect you if you had your car key or the neighbor's car. The neighbor's generator failure would be of little consequence if taxis were available. If it were not an important appointment, the absence of cars, buses, and taxis would not matter. On any other morning the broken coffeepot would have been an annoyance (an *incident*, we will call it), but would not have added to your anxiety and caused you to dash out without your keys.

Though the failures were trivial in themselves, and each one had a backup system, or redundant path to tread if the main one were blocked, the failures became serious when they interacted. It is the *interaction* of the multiple failures that explains the accident. We expect bus strikes

occasionally, we expect to forget our keys with that kind of apartment lock (why else hide a redundant key?), we occasionally loan the extra key to someone rather than disclose its hiding place. What we don't expect is for all of these events to come together at once. That is why we told the secretary that it was a crazy morning, too complex to explain, and invoked Murphy's law to ourselves (if anything can go wrong, it will).

That accident had its cause in the interactive nature of the world for us that morning and in its tight coupling—not in the discrete failures, which are to be expected and which are guarded against with backup systems. Most of the time we don't notice the inherent coupling in our world, because most of the time there are no failures, or the failures that occur do not interact. But all of a sudden, things that we did not realize could be linked (buses and generators, coffee and a loaned key) became linked. The system is suddenly more tightly coupled than we had realized. When we have interactive systems that are also tightly coupled, it is "normal" for them to have this kind of an accident, even though it is infrequent. It is normal not in the sense of being frequent or being expected—indeed, neither is true, which is why we were so baffled by what went wrong. It is normal in the sense that it is an inherent property of the system to occasionally experience this interaction. Three Mile Island was such a normal or system accident, and so were countless others that we shall examine in this book. We have such accidents because we have built an industrial society that has some parts, like industrial plants or military adventures, that have highly interactive and tightly coupled units. Unfortunately, some of these have high potential for catastrophic accidents.

Our "day in the life" example introduced some useful terms. Accidents can be the result of *multiple failures*. Our example illustrated failures in five components: in design, equipment, procedures, operators, and environment. To apply this concept to accidents in general, we will need to add a sixth area—supplies and materials. All six will be abbreviated as the DEPOSE components (for design, equipment, procedures, operators, supplies and materials, and environment). The example showed how different parts of the system can be quite dependent upon one another, as when the bus strike created a shortage of taxis. This dependence is known as *tight coupling*. On the other hand, events in a system can occur independently as we noted with the failure of the generator and forgetting the keys. These are *loosely coupled* events, because although at this time they were both involved in the same production sequence, one was not caused by the other.

One final point which our example cannot illustrate. It isn't the best case of a normal accident or system accident, as we shall use these terms,

because the interdependence of the events was comprehensible for the person or "operator." She or he could not do much about the events singly or in their interdependence, but she or he could understand the interactions. In complex industrial, space, and military systems, the normal accident generally (not always) means that the interactions are not only unexpected, but are *incomprehensible* for some critical period of time. In part this is because in these human-machine systems the interactions literally cannot be seen. In part it is because, even if they are seen, they are not believed. As we shall find out and as Robert Jervis and Karl Weick have noted,³ seeing is not necessarily believing; sometimes, we must believe before we can see.

Variations on the Theme

While basically simple, the idea that guides this book has some quite radical ramifications. For example, virtually every system we will examine places "operator error" high on its list of causal factors—generally about 60 to 80 percent of accidents are attributed to this factor. But if, as we shall see time and time again, the operator is confronted by unexpected and usually mysterious interactions among failures, saying that he or she should have zigged instead of zagged is possible only after the fact. Before the accident no one could know what was going on and what should have been done. Sometimes the errors are bizarre. We will encounter "noncollision course collisions," for example, where ships that were about to pass in the night suddenly turn and ram each other. But careful inquiry suggests that the mariners had quite reasonable explanations for their actions; it is just that the interaction of small failures led them to construct quite erroneous worlds in their minds, and in this case these conflicting images led to collision.

Another ramification is that great events have small beginnings. Running through the book are accidents that start with trivial kitchen mishaps; we will find them on aircraft and ships and in nuclear plants, having to do with making coffee or washing up. Small failures abound in big systems; accidents are not often caused by massive pipe breaks, wings coming off, or motors running amok. Patient accident reconstruction reveals the banality and triviality behind most catastrophes.

Small beginnings all too often cause great events when the system uses a "transformation" process rather than an additive or fabricating one.

Where chemical reactions, high temperature and pressure, or air, vapor, or water turbulence is involved, we cannot see what is going on or even, at times, understand the principles. In many transformation systems we generally know what works, but sometimes do not know why. These systems are particularly vulnerable to small failures that "propagate" unexpectedly, because of complexity and tight coupling. We will examine other systems where there is less transformation and more fabrication or assembly, systems that process raw materials rather than change them. Here there is an opportunity to learn from accidents and greatly reduce complexity and coupling. These systems can still have accidents—all systems can. But they are more likely to stem from major failures whose dynamics are obvious, rather than the trivial ones that are hidden from understanding.

Another ramification is the role of organizations and management in preventing failures—or causing them. Organizations are at the center of our inquiry, even though we will often talk about hardware and pressure and temperature and the like. High-risk systems have a double penalty: because normal accidents stem from the mysterious interaction of failures, those closest to the system, the operators, have to be able to take independent and sometimes quite creative action. But because these systems are so tightly coupled, control of operators must be centralized because there is little time to check everything out and be aware of what another part of the system is doing. An operator can't just do her own thing; tight coupling means tightly prescribed steps and invariant sequences that cannot be changed. But systems cannot be both decentralized and centralized at the same time; they are organizational Pushmepullyous, straight out of Dr. Doolittle stories, trying to go in opposite directions at once. So we must add organizational contradictions to our list of problems.

Even aside from these inherent contradictions, the role of organizations is important in other respects for our story. Time and time again warnings are ignored, unnecessary risks taken, sloppy work done, deception and downright lying practiced. As an organizational theorist I am reasonably unshaken by this; it occurs in all organizations, and it is a part of the human condition. But when it comes to systems with radioactive, toxic, or explosive materials, or those operating in an unforgiving, hostile environment in the air, at sea, or under the ground, these routine sins of organizations have very nonroutine consequences. Our ability to organize does not match the inherent hazards of some of our organized activities. Better organization will always help any endeavor. But the best is not good enough for some that we have decided to pursue.

Nor can better technology always do the job. Besides being a book about organizations (but painlessly, without the jargon and the sacred texts), this is a book about technology. You will probably learn more than you ever wanted to about condensate polishers, buffet boundaries, reboilers, and slat retraction systems. But that is in passing (and even while passing you are allowed a considerable measure of incomprehension). What is not in passing but is essential here is an evaluation of technology and its "fixes." As the saying goes, man's reach has always exceeded his grasp (and of course that goes for women too). It should be so. But we might begin to learn that of all the glorious possibilities out there to reach for, some are going to be beyond our grasp in catastrophic ways. There is no technological imperative that says we *must* have power or weapons from nuclear fission or fusion, or that we *must* create and loose upon the earth organisms that will devour our oil spills. We could reach for, and grasp, solar power or safe coal-fired plants, and the safe ship designs and industry controls that would virtually eliminate oil spills. No catastrophic potential flows from these.

It is particularly important to evaluate technological fixes in the systems that we cannot or will not do without. Fixes, including safety devices, sometimes create new accidents, and quite often merely allow those in charge to run the system faster, or in worse weather, or with bigger explosives. Some technological fixes are error-reducing—the jet engine is simpler and safer than the piston engine; fathometers are better than lead lines; three engines are better than two on an airplane; computers are more reliable than pneumatic controls. But other technological fixes are excuses for poor organization or an attempt to compensate for poor system design. The attention of authorities in some of these systems, unfortunately, is hard to get when safety is involved.

When we add complexity and coupling to catastrophe, we have something that is fairly new in the world. Catastrophes have always been with us. In the distant past, the natural ones easily exceeded the human-made ones. Human-made catastrophes appear to have increased with industrialization as we built devices that could crash, sink, burn, or explode. In the last fifty years, however, and particularly in the last twenty-five, to the usual cause of accidents—some component failure, which could be prevented in the future—was added a new cause: interactive complexity in the presence of tight coupling, producing a system accident. We have produced designs so complicated that we cannot anticipate all the possible interactions of the inevitable failures; we add safety devices that are deceived or avoided or defeated by hidden paths in the systems. The systems have become more complicated because either they are dealing

with more deadly substances, or we demand they function in ever more hostile environments or with ever greater speed and volume. And still new systems keep appearing, such as gene splicing, and others grow ever more complex and tightly tied together. In the past, designers could learn from the collapse of a medieval cathedral under construction, or the explosion of boilers on steamboats, or the collision of railroad trains on a single track. But we seem to be unable to learn from chemical plant explosions or nuclear plant accidents. We may have reached a plateau where our learning curve is nearly flat. It is true that I should be wary of that supposition. Reviewing the wearisome Cassandras in history who prophesied that we had reached our limit with the reciprocating steam engine or the coal-fired railroad engine reminds us that predicting the course of technology in history is perilous. Some well-placed warnings will not harm us, however.

One last warning before outlining the chapters to come. The new risks have produced a new breed of shamans, called risk assessors. As with the shamans and the physicians of old, it might be more dangerous to go to them for advice than to suffer unattended. In our last chapter we will examine the dangers of this new alchemy where body counting replaces social and cultural values and excludes us from participating in decisions about the risks that a few have decided the many cannot do without. The issue is not risk, but power.

Fast Forward

Chapter 1 will examine the accident at Three Mile Island (TMI) where there were four independent failures, all small, none of which the operators could be aware of. The system caused that accident, not the operators. Chapter 2 raises the question of why, if these plants are so complex and tightly coupled, we have not had more TMIs. A review of the nuclear power industry and some of its trivial and its serious accidents will suggest that we have not given large plants of the size of TMI time to express themselves. The record of the industry and the Nuclear Regulatory Commission is frightening, but not because it is all that different from the records of other industries and regulatory agencies. It isn't. It is frightening because of the catastrophic potential of this industry; it has to have a perfect performance record, and it is far from achieving that.

We can go a fair distance with some loosely defined concepts such as

complexity, coupling, and catastrophe, but in order to venture further into the world of high-risk systems we need better definitions, and a better model of systems and accidents and their consequences. This is the work of Chapter 3, where terms are defined and amply illustrated with still more accident stories. In this chapter we explore the advantages of loose coupling, map the industrial, service, and voluntary organizational world according to complexity and coupling, and add a definition of types of catastrophes. Chapter 4 applies our complexity, coupling, and catastrophe theories to the chemical industry. I wish to make it clear that normal accidents or, as we will generally call them, system accidents, are not limited to the nuclear industry. Some of the most interesting and bizarre examples of the unanticipated interaction of failures appear in this chapter—and we are now talking about a quite well-run industry with ample riches to spend on safety, training, and high-technology solutions.

Yet chemical plants mostly just sit there, though occasionally they will send a several hundred pound missile a mile away into a community or incinerate a low flying airplane. In Chapter 5 we move out into the environment and examine aircraft and flying, and air traffic control and the airports and airways. Flying is in part a transformation system, but largely just very complex and tightly coupled. Technological fixes are made continuously here, but designers and airlines just keep pushing up against the limits with each new advance. Flying is risky, and always will be. With the airways system, on the other hand, we will examine the actual reduction of complexity and coupling through organizational changes and technological developments; this system has become very safe, as safety goes in inherently risky systems. An examination of the John Wayne International Airport in Orange County, California, will remind us of the inherent risks.

With marine transport, in Chapter 6, the opposite problem is identified. No reduction in complexity or coupling has been achieved. Horrendous tales are told, three of which we will detail, about the needless perils of this system. We will analyze it as one that induces errors through its very structure, examining insurance, shipbuilders, shippers, captains and crews, collision avoidance systems, and the international anarchy that prevents effective regulation and encourages cowboys and hot rodders at sea. One would not think that ships could pile up as if they were on the Long Island Expressway, but they do.

Chapter 7 might seem to be a diversion since dams, lakes, and mines are not prone to system accidents. But it will support our point because they are also linear, rather than complex systems, and the accidents there

are foreseeable and avoidable. However, when we move away from the individual dam or mine and take into account the larger system in which they exist, we find the "eco-system accident," an interaction of systems that were thought to be independent but are not because of the larger ecology. Once we realize this we can prevent future accidents of this type; in linear systems we can learn from our mistakes. Dams, lakes, and mines also simply provide tales worth telling. Do dams sink or float when they fail? Could we forestall a colossal earthquake in California by a series of mammoth chiropractic spinal adjustments? How could we lose a whole lake and barges and tugs in a matter of hours? (By inadvertently creating an eco-system accident.)

Chapter 8 deals with far more esoteric systems. Space missions are very complex and tightly coupled, but the catastrophic potential was small and now is smaller. More important, this system allows us to examine the role of the operator (in this case, extraordinarily well-trained astronauts) whom the omniscient designers and managers tried to treat like chimpanzees. It is a cautionary tale for all high-technology systems. Accidents with nuclear weapons, from dropping them to firing them by mistake, will illustrate a system so complicated and error-prone that the fate of the earth may be decided more by inadvertence than anger. The prospects are, I am afraid, terrifying. Equally frightening is the section in this chapter on gene splicing, or recombinant DNA. In this case, in the unseemly haste for prizes and profits, we have abandoned even the most elementary safeguards, and may loose upon the world a rude beast whose time need not have come.

In the last chapter we shall examine the new shamans, the risk assessors, and their inadvertent allies, the cognitive psychologists. Naturally as a sociologist, I will have a few sharp words to say about the latter, but point out that their research has really provided the grounds for a public role in high-risk decision making, one the risk assessors do not envisage. Finally, we will add up the credits and deficits of the systems we examined, and I will make a few modest suggestions for complicating the lives of some systems—and shutting others down completely.

Normal Accident at Three Mile Island

Our first example of the accident potential of complex systems is the accident at the Three Mile Island Unit 2 nuclear plant near Harrisburg, Pennsylvania, on March 28, 1979. I have simplified the technical details a great deal and have not tried to define all of the terms. It is not necessary to understand the technology in any depth. What I wish to convey is the interconnectedness of the system, and the occasions for baffling interactions. This will be the most demanding technological account in the book, but even a general sense of the complexity will suffice if one wishes to merely follow the drama rather than the technical evolution of the accident.*

TMI is clearly our most serious nuclear power plant accident to date. The high drama of the event gripped the nation for a fortnight, as reassurance gave way to near panic, and we learned of a massive hydrogen bubble and releases that sent pregnant women and others fleeing the area. The President of the United States toured the plant while two feeble pumps, designed for quite other duties, labored to keep the core from

*This account draws from many sources, and I have not cited each point individually. See the references from the first part of the bibliography.

melting further. (One of them soon failed, but fortunately by the time the second pump failed the system had cooled sufficiently to allow for natural circulation.) The subsequent investigations and law suits disclosed a seemingly endless story of incompetence, dishonesty, and cover-ups before, during, and after the event; indeed, new disclosures were appearing as this book went to press. Yet, as we shall see in chapter 2 when we examine other accidents, the performance of all concerned—utility, manufacturer, regulatory agency, and industry—was about average. Rather sizeable bits and pieces of the TMI disaster can be found elsewhere in the industry; they had just never been put together so dramatically before.

Unit 2 at Three Mile Island (TMI) had a hard time getting underway at the end of 1978. Nuclear plants are always plagued with start-up problems because the system is so complex, and the technology so new. Many processes are still not well understood, and the tolerances are frightfully small for some components. A nuclear plant is also a hybrid creation—the reactor itself being complex and new and carefully engineered by one company, while the system for drawing off the heat and using it to turn turbines is a rather conventional, old, and comparatively unsophisticated system built by another company. Unit 2 may have had more than the usual problems. The maintenance force was overworked at the time of the accident and had been reduced in size during an economizing drive. There were many shutdowns, and a variety of things turned out, in retrospect, to be out of order. But one suspects that it was not all that different from other plants; after a plant sustains an accident, a thorough investigation will turn up numerous problems that would have gone unnoticed or undocumented had the accident been avoided. Indeed, in the 1982 court case where the utility, Metropolitan Edison, sued the builder of the reactor, Babcock and Wilcox, the utility charged the builder with an embarrassing number of errors and failures, and the vendor returned the favor by charging that the utility was incompetent to run their machine.¹ But Metropolitan Edison runs other machines, and Babcock and Wilson have built many reactors that have not had such a serious accident. We know so much about the problems of Unit 2 only because the accident at Three Mile Island made it a subject for intense study; it is probably the most well-documented examination of organizational performance in the public record. At last count I found ten published technical volumes or books on the accident alone, perhaps one hundred articles, and many volumes of testimony.

The accident started in the cooling system. There are two cooling systems. The primary cooling system contains water under high pressure

and at high temperature that circulates through the core where the nuclear reaction is taking place. This water goes into a steam generator, where it bathes small tubes circulating water in a quite separate system, the secondary cooling system, and heats this water in the secondary system. This transfer of heat from the primary to the secondary system keeps the core from overheating, and uses the heat to make steam. Water in the secondary system is also under high pressure until it is called upon to turn into steam, which drives the turbines that generate the electric power. The accident started in the secondary cooling system.

The water in the secondary system is not radioactive (as is the water in the primary system), but it must be very pure because its steam drives the finely precisioned turbine blades. Resins get into the water and have to be removed by the condensate polisher system, which removes particles that are precipitated out.

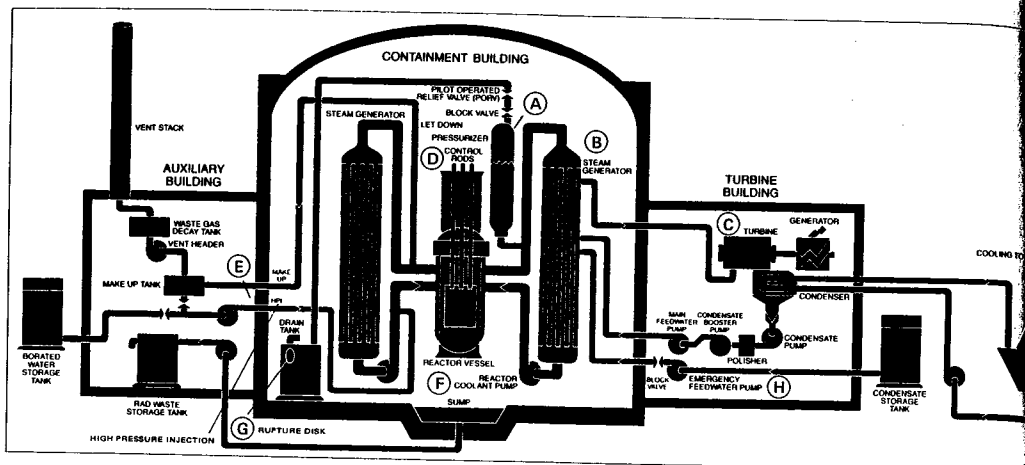
The polisher is a balky system, and it had failed three times in the few months the new unit had been in operation. After about eleven hours of work on the system, at 4:00 A.M. on March 28, 1979, the turbine tripped (stopped). Though the operators did not know why at the time, it is believed that some water leaked out of the polisher system—perhaps a cupful—through a leaky seal.

Seals are always in danger of leaking, but normally it is not a problem. In this case, however, the moisture got into the instrument air system of the plant. This is a pneumatic system that drives some of the instruments. The moisture interrupted the air pressure applied to two valves on two feedwater pumps. This interruption “told” the pumps that something was amiss (though it wasn’t) and that they should stop. They did. Without the pumps, the cold water was no longer flowing into the steam generator, where the heat of the primary system could be transferred to the cool water in the secondary system. When this flow is interrupted, the turbine shuts down, automatically—an automatic safety device, or ASD.

But stopping the turbine is not enough to render the plant safe. Somehow, the heat in the core, which makes the primary cooling system water so hot, has to be removed. If you take a whistling tea kettle off the stove and plug its opening, the heat in the metal and water will continue to produce steam, and if it cannot get out, it may explode. Therefore, the emergency feedwater pumps came on (they are at H in Figure 1.1; the regular feedwater pumps which just stopped are above them in the figure). They are designed to pull water from an emergency storage tank and run it through the secondary cooling system, compensating for the water in that system that will boil off now that it is not circulating. (It is

NORMAL ACCIDENTS

FIGURE 1.1



TMI Unit 2 March 28, 1978

Failure #1	Clogged condensate polisher line	ASD	Reactor coolant pumps come on Primary coolant pressure down, temperature up Steam voids form in coolant pipes and core, restricting flow forced by coolant pumps, creating uneven pressures in system
	Moisture in instrument air line		
	False signal to turbine		
ASD*	Turbine stops		
ASD	Feedwater pumps stop		
ASD	Emergency feedwater pumps start		
Failure #2	Flow blocked; valves closed instead of open	ASD	Hi Pressure Injection (HPI) starts, to reduce temperature Pressurizer fills with coolant as it seeks outlet through PORV
	No heat removal from primary coolant		
	Rise in core temperature and pressure		
ASD	Reactor scrams	"Operator error"	Operators reduce HPI to save pressurizer, per procedures Temperature and pressure in core continue to rise because of lack of heat removal, decay heat generation, steam voids, hydrogen generation from the zirconium-water reaction, and uncovering of core. Reactor coolant pumps cavitate and must be shut off, further restricting circulation.
	Reactor continues to heat, "decay heat"		
	Pressure and temperature rise		
ASD	Pilot Operated Relief Valve (PORV) opens		
ASD	PORV told to close		
Failure #3	PORV sticks open		
Failure #4	PORV position indicator signifies it has shut		

*ASD (automatic safety device)

SOURCE: Kemeny, John, et al. Report of the President's Commission on the Accident at Three Mile Island. Washington, D.C.: Government Printing Office, 1979.

Normal Accident at Three Mile Island

like pouring cold water over your plugged tea kettle.) However, these two pipes were unfortunately blocked; a valve in each pipe had been accidentally left in a closed position after maintenance two days before. The pumps came on and the operator verified that they did, but he did not know that they were pumping water into a closed pipe.

The President's Commission on the Accident at Three Mile Island (the Kemeny Commission) spent a lot of time trying to find out just who was responsible for leaving the valves closed, but they were unsuccessful. Three operators testified that it was a mystery to them how the valves had gotten closed, because they distinctly remembered opening them after the testing. You probably have had the same problem with closing the freezer door or locking the front door; you are sure you did, because you have done it many times. Operators testified at the Commission's hearings that with hundreds of valves being opened or closed in a nuclear plant, it is not unusual to find some in the wrong position—even when locks are put on them and a "lock sheet" is maintained so the operators can make an entry every time a special valve is opened or closed.

Accidents often involve such mysteries. A safety hatch on a Mercury spacecraft prematurely blew open (it had an explosive charge for opening it) as the recovery helicopter was about to pick it up out of the water after splashdown. Gus Grissom, the astronaut, insisted afterwards that he hadn't fired it prematurely or hit it accidentally. It just blew by itself. (He almost drowned.) It is the old war between operators and the equipment others have designed and built. The operators say it wasn't their fault; the designers say it wasn't the fault of the equipment or design. Ironically, the astronauts had insisted upon the escape hatch being put in as a safety device in case they had to exit rapidly; it is not the only example we shall uncover of safety devices increasing the chances of accidents. The Three Mile Island operators finally had to concede reluctantly that large valves do not close by themselves, so someone must have goofed.

There were two indicators on TMI's gigantic control panel that showed that the valves were closed instead of open. One was obscured by a repair tag hanging on the switch above it. But at this point the operators were unaware of any problem with emergency feedwater and had no occasion to make sure those valves, which are always open except during tests, were indeed open. Eight minutes later, when they were baffled by the performance of the plant, they discovered it. By then much of the initial damage had been done. Apparently our knowledge of these plants is quite incomplete, for while some experts thought the closed valves constituted an important operator error, other experts held that it did not make much difference whether the valves were closed or not, since the

supply of emergency feedwater is limited and worse problems were appearing anyway.

With no circulation of coolant in the secondary system, a number of complications were bound to occur. The steam generator boiled dry. Since no heat was being removed from the core, the reactor "scrammed." In a scram the graphite control rods, 80 percent silver, drop into the core and absorb the neutrons, stopping the chain reaction. (In the first experiments with chain reactions, the procedure was the same—"drop the rods and scram"; thus the graphic term *scram* for stopping the chain reaction.) But that isn't enough. The decaying radioactive materials still produce some heat, enough to generate electricity for 18,000 homes. The "decay heat" in this 40-foot-high stainless steel vessel, taller than a three-story building, builds up enormous temperature and pressure. Normally there are thousands of gallons of water in the primary and secondary cooling systems to draw off the intense heat of the reactor core. In a few days this cooling system should cool down the core. But the cooling system was not working.

There are, of course, ASDs to handle the problem. The first ASD is the pilot-operated relief valve (PORV), which will relieve the pressure in the core by channeling the water from the core through a big vessel called a pressurizer, and out the top of it into a drain pipe (called the "hot leg"), and down into a sump. It is radioactive water and is very hot, so the valve is a nuisance. Also, it should only be open long enough to relieve the pressure; if too much water comes through it, the pressure will drop so much that the water can flash into steam, creating bubbles of steam, called steam voids, in the core and the primary cooling pipes. These bubbles will restrict the flow of coolant, and allow certain spots to get much hotter than others—in particular, spots by the uranium rods, allowing them to start fissioning again.

The PORV is also known by its Dresser Industries' trade name of "electromatic relief valve." (Dresser Industries is the firm that sponsored ads shortly after the accident saying that actress Jane Fonda was more dangerous than nuclear plants. She was starring in the *China Syndrome*, a popular movie playing at the time that depicted a near meltdown in a nuclear plant.) It is expected to fail once in every fifty usages, but on the other hand, it is seldom needed. The President's Commission turned up at least eleven instances of it failing in other nuclear plants (to the surprise of the Nuclear Regulatory Commission and the builder of the reactor, Babcock and Wilcox, who only knew of four) and there had been two earlier failures in the short life of TMI-Unit 2. Unfortunately, it just so happened that this time, with the block valves closed and one indicator

hidden, and with the condensate pumps out of order, the PORV failed to reseal, or close, after the core had relieved itself sufficiently of pressure.

This meant that the reactor core, where the heat was building up because the coolant was not moving, had a sizeable hole in it—the stuck-open relief valve. The coolant in the core, the primary coolant system, was under high pressure, and was ejecting out through the stuck valve into a long curved pipe, the "hot leg," which went down to a drain tank. Thirty-two thousand gallons, one third of the capacity of the core, would eventually stream out. This was no small pipe break someplace as the operators originally thought; the thing was simply uncorked, relieving itself when it shouldn't.

Since there had been problems with this relief valve before (and it is a difficult engineering job to make a highly reliable valve under the conditions in which it must operate), an indicator had recently been added to the valve to warn operators if it did not reseal. The watchword is "safety" in nuclear plants. But, since nothing is perfect, it just so happened that this time the indicator itself failed, probably because of a faulty solenoid, a kind of electromagnetic toggle switch. Actually, it wasn't much of an indicator, and the utility and supplier would have been better off to have had none at all. Safety systems, such as warning lights, are necessary, but they have the potential for deception. If there had been no light assuring them the valve had closed, the operators would have taken other steps to check the status of the valve, as operators did in a similar accident at another plant a year and a half before. But if you can't believe the lights on your control panel, an army of operators would be necessary to check every part of the system that might be relevant. And one of the lessons of complex systems and TMI is that *any* part of the system might be interacting with other parts in unanticipated ways.

The indicator sent a signal to the control board that the valve had received the impulse to shut down. (It was not an indication that the valve had actually shut down; that would be much harder to provide.) So the operators noted that all was fine with the PORV, and waited for reactor pressure to rise again, since it had dropped quickly when the valve opened for a second. The cork stayed off the vessel for two hours and twenty minutes before a new shift supervisor, taking a fresh look at the problems, discovered it.

We are now, incredibly enough, only thirteen seconds into the "transient," as engineers call it. (It is not a perversely optimistic term meaning something quite temporary or transient, but rather it means a rapid change in some parameter, in this case, temperature.) In these few seconds there was a false signal causing the condensate pumps to fail, two

NORMAL ACCIDENTS

valves for emergency cooling out of position and the indicator obscured, a PORV that failed to reseal, and a failed indicator of its position. *The operators could have been aware of none of these.*

Moreover, while all these parts are highly interdependent, so that one affects the other, they are *not* in direct operational sequence. Direct operational sequence is a sequence of stages as in a production line, or an engineered safety sequence. The operator knows that a block in the condensate line will cause the condensate pump to trip, which will stop water from going to the steam generator and then going to the turbine as steam to drive it, so the turbine will shut down because it will have no source of power to turn it. This is quite comprehensible. But connected to this sequence, although not a part of its production role, is another system, the primary cooling system, which regulates the amount of water in the core. The water level in the core was judged to have fallen, which it had, because of the drop in the pressure and temperature in the primary cooling system. But for the operators there was no obvious connection between this drop and a turbine "trip" (shutdown). Unknown to them, there was an intimate connection because of the interactive complexity of the system. The connection is through the PORV, but that also has no production sequence or safety sequence connection to the trip of the turbines, or to the failure of the condensate polisher system, even had the operators been able to ascertain that this was the cause of the turbine trip. The PORV is expected to operate on the basis of core pressure, regardless of the functioning of the turbine, the secondary cooling system (feedwater to the steam generators and turbine), or the emergency core cooling pumps.

Even if there is a part of the system that is in direct operational sequence, an information failure in any part of that sequence can render the connection opaque, if not invisible. For example, the PORV is connected in a direct sequence to a drain pipe, then to a drain tank, and when that overflows, to a sump. A couple of readings of excessive radioactive water will appear along the way. But for the operators, this was water from an "unknown origin," since they were assured, by the signal light, that the PORV was closed. Since they assumed a pipe break somewhere and since the piping system in the plant is so complex that a member of the Presidential Commission had to use a magnifying glass to try to follow it on the drawings, there was reason to believe that the water could have come from any number of places. Indeed, later in the accident, they found that radioactive water was not traveling to the tank they intended, but because of complex flow and pressure interactions, was

Normal Accident at Three Mile Island

going to a different, wrong tank, which also overflowed, this time in the auxiliary building.

Here we have the essence of the normal accident: the interaction of multiple failures that are not in a direct operational sequence. You could underline this definition, but there is one other ingredient we have not explored in detail—incomprehensibility. In contrast to our appointment-car-key accident, which was quite comprehensible, most normal accidents have a significant degree of incomprehensibility. Let us go back to the TMI story to examine this incomprehensibility, which is the main reason why answer number one to the quiz, operator error, is so wrong for normal accidents.

The PORV was now open and would be for two hours and twenty minutes, and coolant from the core of the reactor was squirting out at a great rate down the hot leg to the drain tank, so pressure in the reactor dropped. This is dangerous unless the temperature is also going down rapidly, because without pressure on the superheated water (over 2,000° F), it will become steam, which does not cool as well and creates bubbles that block the flow of coolant. So one of two reactor coolant pumps (another emergency system) started up automatically and another was started by the operators (thirteen seconds into the accident; check it out on your watch). For two or three minutes things looked fine; the coolant in the core appeared stable. But it wasn't. For a variety of reasons that can only be matters of conjecture, it appears that voids or steam bubbles formed in such a way as to give the appearance of stabilization after the two reactor coolant pumps came on. The operators were not aware that the steam generators were not getting water. When they boiled dry, the reactor coolant heated up again because the secondary coolant system was not removing heat from the primary one, which removes it from the core. Since the core was losing water, pressure in the coolant system dropped sharply.

At this point, two minutes into the accident, another emergency device came on—high-pressure injection, or HPI, which forces water into the core at a rapid rate. Now came the high drama, the action that has been called the major source of the accident and the key operator error. After letting HPI run full tilt for about two minutes, they reduced it drastically, thus not replacing the water that was boiling out through the PORV. This meant that the core was steadily being uncovered—the most fearsome danger in a nuclear plant, for it will then melt the vessel and perhaps loose radiation on the world.

Probing this action by the operators, investigating committees were led

NORMAL ACCIDENTS

back to an earlier accident at an Ohio plant, memos from a TVA engineer, memos in the files of Babcock and Wilcox (the firm that built the reactor), and an accident in Belgium in a Westinghouse reactor. All of these warnings occurred well before TMI. A bureaucratic tale worthy of Franz Kafka came out of the investigation of TMI and the warnings, which we shall forego telling so we can stick with the villains of the piece, according to most reports: the hapless operators.

High-pressure injection involves the injection of cold water at a very high pressure into the reactor core in order to lower reactor temperatures. It goes in at about 1,000 gallons a minute, and could fill a swimming pool in twenty minutes. It is a risky business. The cold water may "shock" the core, producing hairline cracks in equipment in the core, or conceivably in the vessel itself (but probably only if it had been in operation for several years). The high pressure may also cause damage as the core fills up, putting a pressure strain on it. Most experts discount these dangers, but not all. As an indication of how little we understand nuclear systems, I should note that shortly after the accident, some even argued that it was fortunate that the operators cut back on HPI, although this was not the majority view.

Two years later, however, the Nuclear Regulatory Commission issued a report which gave substance to this danger. It disclosed that thirteen reactors, some of them only three to four years old, showed degrees of core vessel brittleness, because of the intense radioactive bombardment that were greater than predicted.² This raised serious safety concerns. Certainly the high-pressure injection of cold water into a brittle vessel could crack the vessel, leading to a meltdown and all its consequences. Fortunately, the TMI core had only been in operation at full power for about forty days.

Another problem with HPI is a matter of lively dispute. It may increase the pressure in the pressurizer by flooding it with water. The pressurizer is a kind of huge shock absorber and stabilizer. It is a large tank with, under normal circumstances, 800 cubic feet of water in the bottom and 700 cubic feet of steam above it. By using heaters in the tank, the pressure of the steam at the top can be raised or lowered, and this controls the pressure of the water cooling the core. If HPI sends too much water into the core, it will flood the pressurizer. (This is called "going solid"—solid water and no steam.) If there is a substantial pressure surge in the core, the cushion provided by the steam in the pressurizer would be lost and the coolant pipes could burst (one source of a LOCA, or loss of coolant accident), perhaps causing a meltdown. Even if the safety

Normal Accident at Three Mile Island

valves prevented a pipe burst, a full pressurizer still presents a serious situation. It is a first-line emergency safety device (ESD), and should not be disabled.

Operators were assiduously trained to avoid going solid in the pressurizer by both the vendor, Babcock and Wilcox, and the user, Metropolitan Edison, which operates TMI. There was no hint in the training manual or the procedure manual that under some circumstances it might be preferable to go solid in the pressurizer rather than cut back on HPI. Such a directive was considered after an earlier accident at another plant, but was rejected by Babcock and Wilcox. At this point, some two minutes into the accident, there was a circumstance in which HPI was needed more than a resilient pressurizer. The core was about to be uncovered.

After HPI came on, the operators were looking primarily at two dials, close to one another. One indicated that the pressure in the reactor was still falling, which was mysterious because the other indicated that pressure in the pressurizer was rising—indeed, it was dangerously high. But they should move together, and always had. They are connected by pipes, and the pressurizer is supposed to control the pressure in the coolant system; that is what it is there for. If pressure is up in the pressurizer, and it is connected to the core, it should be up in the core.

Perhaps the dials were wrong. It sometimes happens. But which one? If the reactor dial was correct, and pressure was falling in the reactor, there must be some large anomaly, because there was plenty of water going into the core through the reactor coolant pumps, which were still running, and through the high-pressure injection that had just started. Even if there were a small pipe break somewhere, the reactor coolant pumps would ensure that the core would remain covered even without HPI. With all this water going in, how could the pressure fall? On the other hand, since the operators knew that the emergency feedwater pumps came on (but not that they had nothing to pump because of the closed valves), they thought that the secondary cooling system should be cooling the core, so pressure in the core would be falling. But if it were, why did HPI come on? Perhaps the reactor pressure dial was wrong.

The other dial was a serious source of concern. The high pressure in the pressurizer eliminated a safety margin, and all instructions said the pressurizer should not be flooded. It stood between the operators and the possibility of a loss of coolant accident, a LOCA; because if there were no steam at the top, a pressure surge could lead to a pipe break. They could see the connection between HPI and the high pressure reading in the pressurizer. High-pressure injection was flooding the core and sending

NORMAL ACCIDENTS

water up to flood the pressurizer. So they cut back on it drastically (they "throttled back on the makeup valves"). Pressure in the pressurizer, sure enough, came down, relieving the danger of going solid.

What they didn't know, and couldn't know, was that with the PORV open and the two feedwater valves blocked, preventing the removal of residual heat, they already had a LOCA, but not from a pipe break. The rise in pressure in the pressurizer was probably due to the steam voids rapidly forming because the core was close to becoming uncovered. They thought they were avoiding a LOCA when they were *in* one and were making it worse. With the PORV stuck open, the danger of going solid in the pressurizer was reduced because the open valve would provide some relief. But no one knew it was open.

The Kemeny Commission thought the operators should have known, and berated them in its report—they were "oblivious" to the danger; two readings "should have clearly alerted" them to the LOCA; "the major cause of the accident was due to inappropriate actions by those who were operating the plant," they said in their final report.³ Babcock and Wilcox agreed; this was the sole cause of the accident, they argued in a press conference. The British Secretary of State for Energy was less diplomatic—the accident was caused by "stupid errors," he said.⁴

Actually, there were three readings that should have indicated a LOCA to the operator, and it is a lesson in the fate of warnings to examine them. First, we should note that a LOCA is the most feared of the probable accidents in a plant, for it means the core can melt, and in what are called worst-case analyses could cause a steam explosion and rupture the vessel, spewing radioactivity. Even without a steam explosion, the extreme heat of uncontrolled fissioning could breach containment. LOCA will occur when the water level drops below the level of the fuel rods, and they overheat. But there is no direct measure of water level in the core in the Babcock and Wilcox reactors. One could be put on, said a Babcock and Wilcox official during a press conference, but it would be hard to provide and would create other complications.⁵ One hesitates to penetrate the core more than needed, and it would be hard to measure surging water under high pressure, about to flash into steam. So, let's examine the indirect measures.

One device measured drain tank pressures. But it is not considered a particularly vital indicator by the designers, and is located on the back side of a 7-foot high control panel, near the bottom. Not suspecting they were in a LOCA, no one bothered to examine it (though the record is vague on this question). Another indicator showed the temperature of the drain tank; with hundreds of gallons of hot coolant spewing out and

Normal Accident at Three Mile Island

going to the drain tank, that temperature reading should be way up. It was indeed up. But they had been having trouble with a leaky PORV for some weeks, meaning that there was always some coolant going through it, so it was usual for it to be higher than normal. It did shoot up at one point, they noted, but that was shortly after the PORV opened, and when it didn't come down fast that was comprehensible, because the pipe heats up and stays hot. "That hot?" a commissioner interrogating an operator asked, in effect. The operator replied, in effect, "Yes; if it were a LOCA I would expect it to be much higher." It was not the LOCA they were trained for on the simulators that are used for training sessions, since it had some coolant coming in through an emergency system, and some coming in through HPI, which was only throttled back, not stopped. Their training never imagined a multiple accident with a stuck PORV, and blocked valves. Well, what about the drop in pressure in the core itself; surely this would indicate that the coolant was getting out somehow. But the operators discounted that indicator as erroneous or simply mysterious because it contradicted the one next to it, the pressurizer indicator, which was rising. A supervisor testified:

I think we knew we were experiencing something different, but I think each time we made a decision it was based on something we knew about. For instance: pressure was low, but they had opened the feed valves quickly in the steam generator, and they thought that might have been "shrink." There was logic at that time for most of the actions, even though today you can look back and say, well, that wasn't the cause of that, or, that shouldn't have been that long.⁶

We will encounter this man's dilemma a few more times in this book; it goes to the core of a common organizational problem. In the face of uncertainty, we must, of course, make a judgment, even if only a tentative and temporary one. Making a judgment means we create a "mental model" or an expected universe.

Suppose you get an ambiguous order from your boss. You don't know if you should do A or B because the order could mean either. Alternative A would be correct if something were terribly wrong or if the situation were quite unusual. B would be correct if it were a situation that had occurred a few times before and was not all that serious. You decide she must have meant B. This alternative has been used before, and is easy to carry out. To do it you perform steps 1, 2, and 3. Still uncertain, you check the consequences of each. After step 1, certain things should happen, and they do. The same with steps 2 and 3. Despite the fact that this is no proper test of the appropriateness of alternative B rather than A, it

NORMAL ACCIDENTS

serves to "confirm" your decision. In so believing, you are actually creating a world that is congruent with your interpretation, even though it may be the wrong world. It may be too late before you find that out.

The operators at TMI were faced with this dilemma. Alternative A, believing in the core pressure indicator, would mean that the core was being uncovered. Uncovering the core is unheard of; it had never happened in a large (over 750 megawatts) commercial light water reactor in the 380 or so "reactor years" of large commercial light water reactor operation. (A reactor year measure adds together the number of years each reactor has operated. For reactors of around 1,000 Mws, the more appropriate comparison, there were only about thirty-five reactor years of operating experience). Believing the B gauge rather than the A one (or attributing A to some temporary phenomenon) was soon confirmed—pressure dropped in the pressurizer after HPI was cut back. The other anomalies were accounted for in rapid fashion. Since the light showed the PORV had shut, the pressure decline in the core could be due to "cold shock" (from the two-minute burst of HPI fluid), or it could be a faulty reading. There had been faulty readings in the past; the drain tank temperature was one example.

Besides, about this time—just four or five minutes into the accident—another more pressing problem arose. The reactor coolant pumps that had turned on started thumping and shaking. They could be heard and felt from far away in the control room. Would they withstand the violence they were exposed to? Or should they be shut off? A hasty conference was called, and they were shut off. (It could have been, perhaps should have been, a sign that there were further dangers ahead, since they were "cavitating"—not getting enough emergency coolant going through them to function properly.)

In the control room there were three audible alarms sounding, and many of the 1,600 lights (on-off lights and rectangular displays with some code numbers and letters on them) were on or blinking. The operators did not turn off the main audible alarm because it would cancel some of the annunciator lights. The computer was beginning to run far behind schedule; in fact it took some hours before its message that something might be wrong with the PORV finally got its chance to be printed. Radiation alarms were coming on. The control room was filling with experts; later in the day there were about forty people there. The phones were ringing constantly, demanding information the operators did not have.

Two hours and twenty minutes after the start of the accident, a new shift came on. The record is unclear, but either the new shift supervisor

Normal Accident at Three Mile Island

decided to check the PORV, or an expert talking with a supervisor over the telephone questioned its status, and the operators discovered the stuck valve, and closed a block valve to shut off the flow to the PORV. The operator testified at the Kemeny Commission hearings that it was more of an act of desperation to shut the block valve than an act of understanding. After all, he said, you do not casually block off a safety system. It was fortunate that it occurred when it did; incredible damage had been done, with substantial parts of the core melting, but had it remained open for another thirty minutes or so, and HPI remained throttled back, there would probably have been a complete meltdown, with the fissioning material threatening to breach containment.

But the accident was far from over. New dangers appeared every few hours. Thirty-three hours into the accident another unexpected and mysterious interaction occurred. Confusion still reigned when the first sign of the famous hydrogen bubble appeared; the bubble threatened the integrity of the plant for the next few days. Again we have a lesson in the meaning of warnings, and in the difficulty that even experts have in understanding such a complex human-made system as a nuclear plant. Here is the background:

The fuel rods—36,816 of them—contain enriched uranium in little pills, all stacked within a thin liner, like the cigarette paper around tobacco, only about 12 feet long. Water circulates through the stacks of rods and cools the cladding so it won't melt. When they get too hot, though, the liner, or "cladding," can react with the water in a zirconium-water reaction. This consumes oxygen, thus freeing hydrogen, making hydrogen bubbles, which then can make pockets of hydrogen gas if there is room for them, and a dandy explosion if there is also a bit of oxygen and a spark.

It is not a well-understood aspect of nuclear engineering, I take it. Three years before the accident, when a nuclear physicist from the University of Pittsburgh mentioned the danger in the *Bulletin of the Atomic Scientist*,⁷ a nuclear physicist from Pennsylvania State University wrote a scoffing rebuttal, saying the matter had been well studied and there was no danger.⁸ We might put this quarrel down to the traditional rivalry between these two universities and treat it as insignificant, except that the latter, scoffing, scientist turned out to be the advisor on nuclear power production to Governor Thornburg of Pennsylvania, and was in the thick of the expert advice at TMI. After TMI President Reagan appointed him Chairman of the Nuclear Regulatory Commission. Contrary to industry pronouncements, there is still a good bit of mystery about atomic power plants, and this was an unfortunate case, since it was hours or days

NORMAL ACCIDENTS

(depending upon whose testimony you wish to believe) before the bubble was conceived by the experts. Thus, the operators might be forgiven for ignoring yet another signal that something was drastically wrong, the "spike."

Here is how the warning occurred. At 1:00 P.M. Wednesday, thirty-three hours into the accident, there was a soft but distinct bang heard in the control room. This is not the kind of thing you expect or like to hear. A quick glance showed that the reading of the amount of pressure in the containment building—the building that holds the core vessel itself, and the pressurizer, drain tank, sump pump, pipes, electrical connections, et cetera—had jumped suddenly. In fact, the pressure spike had reached half the design limit for the building (if the pressure had been twice as high, the building might have cracked). Here the story gets murky. The operator, interviewed by the President's Commission, said, "We kind of wrote it off at the time as possibly instrument malfunction of some sort." This was not an unreasonable conclusion, since instruments were malfunctioning. "We did not have a firm conclusion regarding the spike," he went on, "since it appeared and went away with such rapidity."⁹

But another story has it that someone on the floor—there were perhaps twenty people there—knew that there had been a hydrogen explosion. Fearing another pocket of gas might appear and be ignited by a spark, he asked another operator not to restart a failed pump. The operator replied, "I already have." (Pumps have motors; they are big and make sparks.) That means, the first fellow said, that we don't have more hydrogen.¹⁰ That is, he knew there had already been one hydrogen "burn." If this story is true, a lot of people went through the rest of the day ignorant of a vital piece of information.

Why worry? Because with more hydrogen being produced, the gas might find other ways to be vented from the core—whose condition was unknown to the personnel—and collect in the containment building. With pumps starting and stopping and other activity, a spark could easily be available, and the containment building had oxygen in it. If the hydrogen managed to collect in a spot near a lot of equipment and explode there, the pressure force could send missiles flying. (Indeed, three years later they found the huge crane required to lift off the top of the reactor vessel had been damaged by missiles from the explosion; two engineers protested the crane was not safe enough to use and were fired.¹¹) Even a small explosion might pierce a cable or two and cause a short circuit, shutting off the emergency cooling, or rupture a pipe, causing a more rapid LOCA, and so on, though the design does take into account the possibility of "guillotine" accidents where pipes enter the containment

Normal Accident at Three Mile Island

wall. Even after the PORV was closed, the build-up of hydrogen in the core vessel itself is extremely dangerous, because its bubble can prevent the flow needed for cooling. The hydrogen will not explode there, but it need not explode to be dangerous.

Of such complexities is the normal accident made. For all but one operator, presumably, and for all the experts, the pressure spike and the hydrogen bubble were incomprehensible. To understand the accident, they would have had to know that the core was seriously uncovered, and that a zirconium-water reaction was likely (a possibility disputed by an expert), and would have had to recall that the PORV had been open, allowing the hydrogen to get out of the core vessel into the building that contained it. These are not expected sequences in a production or safety system; they are multiple failures that interacted in an incomprehensible manner—for all but at least one person, who, incredibly enough, wasn't talking, or didn't examine the implications of his hunch. A warning such as the spike is only effective if it fits into our mental model of what is going on. As with the "warnings" of Pearl Harbor, it can get swamped by the multitude of signals that fit our expectations, and thus be discounted as "noise" in the system.

That's enough on the accident for now. We will return to Harrisburg a few more times. But first we should pose a question that may have been bothering you: If this is typical of a nuclear plant, why have we had only one TMI? Or is this just a bad apple in the nuclear barrel? In the next chapter I will try to show that TMI isn't unusual, and yet indicate why there has been only one TMI. In Chapter 3, we will have to examine our language and define major terms such as complexity, coupling, and catastrophes. Thus equipped, we will be ready to journey through other systems in subsequent chapters, exploring ways to prevent such threatening accidents as the near meltdown at TMI. For example, wouldn't better organization help, or more money and resources for better people and equipment? Not much, I shall argue.

Nuclear Power as a High-Risk System: Why We Have Not Had More TMIs—But Will Soon

Why haven't we had more Three Mile Islands? If nuclear power is so risky, why has no one been killed by radiation exposure as a result of a nuclear power plant accident? If the safety systems have worked so far, nearly twenty years into the nuclear power age, why call this a high-risk system? One answer is that the "defense in depth" safety systems have worked, limiting the course of accidents. We shall examine these safety systems briefly. But a more accurate and less reassuring answer is that we simply have not given the nuclear power system a reasonable amount of time to disclose its potential. We do not really have twenty years of

experience, but very little—too little, by most industrial standards, to make a reasonable assessment of the risks.

The nuclear industry does not agree that it lacks experience. Therefore, we must journey into the heart of industry experience, taking a close look at some serious accidents, some trivial ones, problems of reliability and management, and above all, the special characteristics of the nuclear power system. This will give us the necessary tools, in the form of ideas or concepts, to enter, in later chapters, the world of other high-risk systems that someone has decided we cannot live without.

Operating Experience

We have not given the nuclear power generation system enough time to express itself; and we are only just beginning to uncover the potential dangers that make any prediction of risk very uncertain. We are about twenty years into the era of commercial plant operation, but our experience is not all with one type or size of plant. Indeed, the oldest plant in operation in 1982 was a 430 megawatt (Mw) reactor operating more or less continuously since 1967. We do not build this size any more, so its sixteen years of operating experience is of somewhat limited value.

The small plants of around 400 Mws are different in many respects from the larger ones of around 1,000 Mws; changes in scale produce surprising results. For example, the larger plants appear to be less reliable; there is more downtime after the first two or three years. In addition to size, there are two different types of U.S. reactors, the pressurized water reactors (PWR) and the boiling water reactors (BWR). Experience accumulated in one does not necessarily enable us to judge the reliability of the other; some aspects are similar, some different. In addition to size and type, there are four different U.S. manufacturers. General Electric builds only BWRs, while Westinghouse, Babcock and Wilcox, and Combustion Engineering all build PWRs. The designs differ, of course, limiting the accumulation of experience to some degree.

Thus, to say, as proponents of nuclear power often do, that we have 500 "reactor years" of experience with commercial plants (summing up the number of plants times the number of years each has been operating) is quite misleading. There is no consensus on what would be adequate experience for such a complex and novel transformation process as con-

trolled nuclear fission creating steam that drives turbines; there are thousands of years operating experience with large turbines, but very little with nuclear fission. The condensate polisher problem on the turbine side of the plant at TMI would have been trivial in a coal fired plant but was not in a nuclear plant. We have been building large pressure vessels since the late nineteenth century, but are only beginning to learn the problems with welded stainless steel vessels 40 feet high that are bombarded with neutrons. Every few months new problems appear in nuclear plants, including the failure of supposedly failure-proof emergency scram systems. At the time of TMI we had only thirty-five years experience with reactors the size of Unit 2; that is infancy for a system of this size and complexity.

The first order for a commercial plant that was not in part a demonstration project was placed in 1963. Before that plant was even operating, the boom was on. By the end of 1967 there were seventy-five plants on order; in 1966-67 alone, forty-nine firm orders were placed.¹ More important, by 1968 the utilities were taking orders for plants six times larger than the largest in operation. This extrapolation, from the size of a plant one has some experience with to another six times larger, is very unusual for large, complex installations.

Bupp and Derain, who give the history of commercial reactor development, note that "electric power generation was an industry which had previously operated on the belief that extrapolations of two to one over operating experience were at the outer boundary of acceptable risk."² By 1967, cumulative operating capacity, a measure of experience, was only 3.5 percent of ordered capacity, rather than two to one. In short, no one knew if the seventy-five plants on order would ever work. They also did not know what the capital costs of building them would be. The plants completed in 1975 were about three times the cost per kilowatt produced, in constant dollars, of those completed only five years earlier. "The learning that usually lowers initial costs has not generally occurred in the nuclear power business. Contrary to the industry's own oft-repeated claims that reactor costs were 'soon going to stabilize' and that 'learning by doing' would soon produce cost decreases, just the opposite happened."³

The technical learning curve with these plants (sometimes called "light water" plants) also failed to materialize, according to the study just quoted. "After more than a decade of experience with large light water nuclear power plants, important engineering and design changes were still being made. This is contrary to experience with most other complex industrial products."⁴ After a decade the major problems of well-de-

signed systems should be far behind, but not in this case. The reason, the authors believe, is the haste with which untried designs were ordered, and the stubborn refusal "to face up to the sheer technical complexity of the job that remained even after the first prototype nuclear power plants had been built in the mid- and late 1950s."⁵ Nor was this due to a stodgy industry "boiler business" mentality. The utility industry had been one of the great growth areas in the postwar American economy. Energy production was doubling every nine or ten years, and operating costs were declining steadily, largely as a result of technological progress. Generating costs declined as fossil plant sizes increased and as improvements in operating efficiency continued. It was not a technologically stagnant industry. But it was unprepared for the technological complexities of controlled fission. New complexities are now being realized (and publicized) almost monthly. In time, it seems, the problems for which TMI was an early precursor will unfold in more TMIs.

For example, steam generators are a problem with all power plants; the pipes rust. Special care and materials are used in nuclear plants, but in 1981 it appeared that seventeen reactors, some only five or six years old, had serious rusting problems. The repairs on two plants owned by the Virginia Electric Power Company cost a total of \$112 million. Rusting is a special problem in nuclear plants since the thin tubes in the generators are immersed in water continuously, and leaks will allow radioactive water to get in the secondary (nonradioactive) cooling system. Various steps were taken to reduce the rust, but apparently without success in some plants.⁶ The point is, in a nuclear plant leaks in the generator are failures that can interact with other failures, and thus be a source of system accidents; repairs to such a system can be enormously expensive (in contrast to a conventionally fueled power plant); and there was no way to anticipate these problems in a new technology with such large design and construction lead times.

More serious is the problem of core embrittlement. The bombardment of the containment vessel by the nuclear reaction going on within it has had a greater impact than anticipated. The 40-foot stainless steel vessel is designed to last 40 years, but there are already potential brittleness problems in forty-seven plants, the Nuclear Regulatory Commission (NRC) announced in 1981, and of these, thirteen have serious problems. One of these is only three years old; three others, four years old. The problem is that the core is very hot—about 550° F.—and if you have an emergency and must force thousands of gallons of cold water into the core, the inside of the 8-inch-thick vessel will shrink faster than the outside, creating cracks. In an accident, the pressure must be kept high, further strain-

NORMAL ACCIDENTS

ing the core. These problems apply to PWR systems only, but PWRs account for two-thirds of the operating reactors in the U.S.⁷

These are technical reasons why we have not had sufficient time to have a truly serious nuclear accident—the system is quite new and has not been given a chance to reveal its full potential for danger. Unknown potential cannot be corrected, except by running the plants and taking the risks; without experience, we cannot be sure of the potential for damage inherent in the system's characteristics.

The Construction Problem

There are other problems that are not so directly related to the technological nature of the system, but rather to the nature of the utility and construction industries. Several weeks after TMI, the NRC reported on a continuing study of earthquake protection measures at operating plants. At that point they had identified thirty-five plants with "significant differences" between the way they were designed and the way they were built. This raised questions about "the whole procedure for checking plants," one NRC official said.⁸ Since there is only about one engineer from the NRC to watch over each plant in construction, there is "almost complete reliance on the utility and its contractor to monitor themselves and report on deviations from acceptable standards," said a General Accounting Office report of the previous year.⁹ One would think that reliance on the utility would be adequate, since it is the utility that owns the plant, not the government. But enough stories have appeared to question whether it is possible to rely on anyone to build safe nuclear plants.

For example, at the Marble Hill nuclear project in Madison, Indiana, it took affidavits filed by workers and former workers to alert the NRC to the fact that, as John Emshwiller puts it in his *Wall Street Journal* article, "the builders can't seem to get the hang of pouring concrete."¹⁰ So far, 500 voids (some up to 180 cubic feet in size!), had been found in the concrete structures. Workers were ordered to do cosmetic patching jobs in order to get them past inspection. At another plant, the Brown and Root construction firm was accused of intimidating federal inspectors, in one case putting the inspector into the hospital for two days. On the other hand, engineers have resigned in protest from the NRC, charging coverups and intimidation by the NRC itself. The NRC was informed of falsified documents regarding the inspection of a safety system at one

Nuclear Power as a High-Risk System

midwestern plant, but, according to the NRC administrator, he ignored them. Three months later two employees went public with the documents, and the NRC promised to investigate.¹¹

Perhaps the most striking testimony on unsafe construction in this business is the Diablo Canyon case. Diablo Canyon, in central California, has been waiting for several years to be allowed to operate. After construction was underway, an earthquake fault was discovered a short distance from the site and extensive earthquake protection was required. A little more than a week before the plant was scheduled to open (after some dramatic protests from anti-nuclear groups and local residents and 1,600 arrests), a diffident, 25-year-old engineer for Pacific Gas and Electric Co., owners of Diablo Canyon, was staring at some drawings of a part of the plant. The drawings divided the floor of the containment building into five segments, and showed the location of some heavy equipment (fan coolers). Something about the drawings bothered him. "Just out of curiosity, I pulled some detailed cutaway engineering drawings out of the file—drawings that showed the actual placing of those coolers, and the two diagrams didn't match. It didn't make sense."¹² He insisted that he was not looking for flaws; his discovery was accidental.

What he found was that in 1977 the utility had mistakenly sent the wrong set of diagrams to its seismic engineering consultants, who were to provide seismic shock calculations to be used in strengthening the vulnerable parts of the plant. Instead they sent the diagrams for a second reactor, still under construction, which was the mirror image of the Diablo Canyon reactor about to be retrofitted. The work was performed, and many parts were needlessly reinforced, while others, which should have been strengthened, were left untouched.¹³ Subsequent investigations turned up no fewer than 111 other flaws in the construction of this \$2.5 billion reactor, and by the end of 1982 it was still not operating.

Shoddy construction and inadvertent errors, intimidation and actual deception—these are part and parcel of industrial life. No industry is without these problems, just as no valve can be made failure-proof. Normally, the consequences are not catastrophic. They may be, however, if you build systems with catastrophic potential. No less an authority than former reactor designer and former Dean of the Engineering College at Pennsylvania State University, Nunzio J. Palladino, appointed Chairman of the Nuclear Regulatory Commission in 1981, remarked in December of that year:

During my first five months as NRC chairman a number of deficiencies at some plants have come to my attention which show a surprising lack of profes-

sionalism in the construction and preparation for operation of nuclear facilities. The responsibility for such deficiencies rests squarely on the shoulders of management. . . . There have been lapses of many kinds—in design analyses resulting in built-in errors, in poor construction, in harassment of quality control personnel and inadequate training of reactor operators.¹⁴

Safer Designs?

If the plants are not built well, and we do not have enough operating experience to assure us the design and equipment are safe, could we turn to other, safer designs? *Are* there safer designs? Apparently there are, though it is well beyond my capacities and the argument of this book to be confident about this. The Canadian reactor, the CANDU, is said to be slower, more “forgiving,” and less tightly coupled than our PWRs and BWRs. Operators have more time to take action, and can take more actions. This has not prevented Canada from having some nuclear accidents, but I gather they are less serious than those we have suffered. But the Canadian plants are also smaller and less efficient than ours.

Some engineers believe we missed the boat in not investing more heavily in the gas-cooled reactor, considered to be safer. A small commercial one was built, but has been shut down for some time, though the utility—Pennsylvania Electric—indicates it still wants to keep alive the possibility of developing gas-cooled reactors because of their increased margin for safety. A second, larger one recently began operation in another utility. A sodium-cooled breeder reactor is operating in France and a much larger one is being built there. These produce more fuel than they use, which is useful since the world supply of uranium is quite limited. But the technology of sodium-cooled breeder reactors is very new and some feel the dangers of radioactive sodium far exceed the dangers from light water reactors—PWRs and BWRs. We shall encounter this later in the chapter when we examine our experience with the Fermi breeder reactor. There are other designs, but there is no evidence that any nuclear reactor designs are significantly less complex and interactive, or significantly less tightly coupled than the light water ones we have been concerned with.

There is a good reason why our dominant design, the pressurized light water reactor (PWR), was adopted, even though heavy water reactors (CANDU), gas-cooled reactors, and perhaps other designs might be better. In the 1950s the U.S. government was very anxious to find peaceful

uses for atomic energy, and, in particular, to develop atomic power production. The reasons for the government's haste are in dispute at this writing, but it certainly was not an expected power shortage or increase in energy costs. In fact, cheap oil and gas was driving out small hydroelectric dams in the northeast and the popular solar hot water heaters found in the south. The government had to offer large incentives to private utilities, and when that did not work, to threaten them with the prospect of socialized power—federal atomic power plants on the TVA model across the country—before the utilities would build them. The government had on hand a design for a reactor; it was being built for submarines. Such a reactor is very compact, very responsive, and can easily be refueled once a year when the submarine returns to port and does not need the power.

None of these characteristics were appropriate to utility installations; indeed, for these, the size, responsiveness, and refueling cycle of submarine reactors are counterproductive. A company does not want to have to shut down its plant each year for refueling, because replacement power has to be bought, and since it generally comes from the least efficient generating sources that are maintained only for peak loads (gas and oil plants with small output), it is very expensive. Compactness is not a requirement at a plant site. Responsiveness is not necessary since these are “base-load” installations, designed to handle the bulk of demand on a steady basis, rather than requiring fluctuations, and they do not need to come up to power or cool down quickly. Nevertheless, the firms that built and sold nuclear plants took over the designs for the submarine systems and modified and greatly enlarged them. There appears to have been a rush to get into the business. Indeed, the first “turnkey” plants—the vendor builds it and “turns” the key over to the utility—were sold at substantial losses in order to get established in the industry. It is a good example of a technological “push” rather than a demand “pull.” This unseemly haste has left us with a particularly complex and tightly coupled design, and a design that was assumed to be capable of being scaled up in size without any serious complications.

Even if there were a technological breakthrough, and a much safer design were available, it is very unlikely that one would be built in the United States in the next decade or two. We have about seventy operating reactors now, and perhaps fifty more that might begin operating in the next few years (unless the rate at which they are being cancelled increases), according to an NRC commissioner.¹⁵ Even the most enthusiastic proponents do not anticipate more than 120 reactors operating in the next five or so years. A new design would not attract much interest in

the financial community; utilities generally find themselves with excess capacity because the rise in demand for electricity, for decades a stable 7 percent, has dropped steadily since 1974 to 1.7 percent in 1981. Further, it would take over ten years to design and build a new facility, even if it were significantly less complex than those we have now. Thus, we will have to live with the plants we have, safe or not; new, dramatically safer ones do not appear to be in the offing, and probably will not be built for a long time to come. Note that I am not saying there could never be a nuclear plant that was *not* highly interactive and tightly coupled (though I suspect the nature of the transformation process involved in this kind of energy production makes that impossible) but only that we shall not see one for many years. And we shall continue to see our existing and nearly ready plants for a longer time—perhaps forty years, if they live up to industry predictions.

Defense in Depth

There is yet a quite different answer to the initial question posed in this chapter: Why have there not been more accidents resembling TMI if these systems are all that dangerous? So far I have argued that we have not given them time. The design and construction flaws will not appear immediately nor in every reactor. But is it not possible that the “defense in depth” is working—that containment buildings do contain; that emergency core cooling systems do cool; that even if some unanticipated radioactivity escapes, the plants are sufficiently far from highly populated centers to reduce the risk to negligible proportions? Yes, but the situation, while reassuring, is not wholly so, because the possibilities for system accidents that evade these defenses still exist. Let us look at each of the defenses.

We can be glad that we have containment buildings. These are concrete shells that cover the reactor vessel and other key pieces of equipment, and are maintained at negative pressures—that is, at a lower air pressure than the atmosphere outside of them—so that if a leak occurs, clean air will flow in rather than radioactive air flowing out. The Soviet Union, which did not begin a large nuclear generating program until about 1970, is far less concerned about the chance of large accidents, so they did not build containment structures for their early reactors, nor do they yet require emergency core cooling systems. Had the accident at

Three Mile Island taken place in one of the plants near Moscow, it would have exposed the operators to potentially lethal doses, and irradiated a large population.

At TMI, the hydrogen explosion (or “burn”) that took place in the containment building generated a pressure surge equal to one-half that which the building was designed to handle. The building was built this strong only because the state of Pennsylvania insisted that it meet the criterion of being able to withstand a direct hit from a jet airliner (it is close to the Harrisburg airport). The initial plans did not call for this. Even if the building were not reinforced, it is unlikely, I am told, that the hydrogen burn would have breached containment and allowed the radioactive particles to escape. However, such a disaster might occur in a plant with all those flaws in the concrete we heard of; the explosion might have taken place thirty minutes later when there would have been much more hydrogen to burn; and it could have happened in a part of the building where more missiles would have been created, which could have ruptured the many penetrations required in the building for controls and pipes. While containment is absolutely necessary, it may not be sufficient. It *can* be ruptured.

We almost had a good test of the ability of the concrete containment structure to withstand an airplane crash in 1971. A B-52 bomber was flying a routine practice flight near Charlevoix, Michigan, on the shores of Lake Michigan. Bombers and fighter-bombers from a nearby Strategic Air Command base routinely flew low-level (1,000 foot) sorties directly over the plant, despite Air Force instructions to stay clear. This time the plane was heading directly toward the reactor when it crashed, skipping off the surface of the water, and raising a fireball 200 to 600 feet in the air. A Grumman aerospace official suggested that it might have flown into radioactive gases from the plant’s stack, which could interfere with the plane’s electronics. The plane was two miles, or about twenty seconds, short of crashing into the plant and testing containment.¹⁶

Fortunately, we tend to build our plants in sparsely populated areas, though they are generally near big cities. The ideal spot for a nuclear plant cannot exist. It should be far from any population concentration in case of an accident, but close to one because of transmission economies; it has to be near a large supply of water, but that is also where people like to live; it should be far from any earthquake faults, but these tend to be near coastlines or rivers or other desirable features; it should be far from agricultural activities, but that also puts it far from the places that need its power. The result has been that most of our plants are near population concentrations, but in farming or resort areas just outside of them.

NORMAL ACCIDENTS

The Indian Point nuclear stations, for example, are on the Hudson River, but just thirty-five miles upwind of Manhattan. The owner of one of the plants there, Consolidated Edison, once proposed building a nuclear plant in the middle of Queens—truly one of the most densely populated areas in the United States. Some plants are built on earthquake-prone coastlines, others on rivers that supply fresh water for large cities and for irrigation. Some people have suggested isolated reactor parks, where several nuclear plants will be built, with long transmission lines to populated areas. But an accident in one of the plants might require the abandonment of the adjacent plants in the park (and thus possible additional accidents).

Despite all these problems, semi-remote siting has no doubt increased the safety of nuclear plants. Many have had small emissions of radioactive materials as a result of accidents. Were they located in Queens, the long-term dangers would be higher. Furthermore, though it is said to be minuscule by almost all experts, the plants do release radioactive materials to the environment in the course of normal operation. The farther you are from that, the better.

Finally, there is the Emergency Core Cooling System (ECCS). Should there be a danger of a core melt, this system will flood the core with water, cooling it. It is in the nature of the beast that we cannot use full-scale testing to see how effectively ECCS will work. In a series of tests with a 9-inch model reactor core, all tests failed.¹⁷ Some critics, such as the Union of Concerned Scientists, believe that as presently constituted, ECCS is an inadequate safeguard. At the Browns Ferry nuclear station, the fire that shut down two reactors and burned out of control for several hours rendered the ECCS system inoperative. Fortunately, other means were used to prevent massive fuel melting. The assessment of the ECCS made in the most ambitious safety study commissioned and carried out by the Atomic Energy Committee (forerunner of the NRC), the *Reactor Safety Study* (RSS or WASH 1400, or Rasmussen Report as it is variously referred to), failed to consider that anything else might be wrong in a plant when there was an emergency that required ECCS. That is, the study ignored the possibility that there could be a variety of failures that in themselves would defeat this safety device. For example, steam generators are a continuous problem with nuclear plants; should many of the tubes in them fail in an accident, so would ECCS. There are also problems with other major subsystems. The integrity of the reactor vessel itself has been questioned, drawing upon industrial experience with vessels in nonnuclear systems.¹⁸ Finally, in 1981 half the Browns Ferry control rods failed to drop on command, and in 1983 the automatic shutdown

Nuclear Power as a High-Risk System

system at the Salem plant in South Jersey failed twice. Both events were assumed to have extremely low probability; both could easily defeat the ECCS.

It is true we can be glad that containment, siting, and major emergency systems exist to reduce the dangers. No doubt there would have been more severe accidents without them. But they are unlikely to prevent all future disasters. Siting is not remote enough; containment is vulnerable to hydrogen explosions, missiles, and faulty construction; and the “defense in depth” major emergency systems such as ECCS are defenses with perhaps not that much depth.

Trivial Events in Nontrivial Systems

Nothing is perfect; every part of every system, industrial or not, is liable to failure. Common, run-of-the-mill industrial plants have a steady run of unremarked failures. The more complicated, highly engineered continuous processing plants, such as chemical, pharmaceutical, and some steel processing plants, are no exception. The more complicated or tightly coupled the plant, the more attention is paid to reducing the occasion for failures, but as I shall argue in the next chapter, this can never be enough. If we add catastrophic potential, as we must with nuclear plants, the everyday failures should not go unremarked. They now become significant. What I will be reporting in this section would not even make a news story in the plant paper, let alone the *New York Times* and the like, if it did not occur in a nuclear plant. In fact, not until after Three Mile Island would most of these incidents even be picked up by the daily paper.

Utilities are quite sensitive to this unwanted and “unjustified” scrutiny, but we *should* be sensitive to trivial events in nontrivial systems. I will start with some trivia, to show the course of consequences in these expensive systems, and then proceed to a few of the famous accidents. Keep in mind that these types of mishaps go on all the time in most organizations; we are being unrealistic if we are surprised that they go on in nuclear plants.

Let's start with a trivial event like the ones that plague us all. In 1980 a worker in the North Anna Number 1 plant of the Virginia Electric and Power Company (VEPCO) was cleaning the floor in an auxiliary building. His shirt caught on a 3-inch handle of a circuit breaker protruding

from a wall. He pulled it free, and apparently was unaware that in doing so he activated the breaker. This shut off the current to the control rod mechanism, and the reactor scrammed (shut off) automatically. This trivial event caused a four-day shutdown, which cost consumers several hundred thousand dollars. Fortunately, the weather was mild, so demand was low. The executive vice president of VEPCO termed the accident embarrassing, but suggested there was a fortunate lesson for us all: The incident "clearly demonstrates the sensitivity of nuclear station systems to the slightest deviation from normal and the ability of these systems to perform safely as designed in immediately stopping the unit."¹⁹ Shutting off current to a major safety system is hardly a slight deviation from normal, and that it can be done so casually suggests an undue degree of sensitivity.

Piping is always a problem in any plant. In a nuclear plant this problem is a bit more severe. During the TMI accident, operators sent radioactive water to the wrong places because the plumbing was so complex and pressures could cause reverse flows. At one plant a small error sent radioactive waste water into the drinking water system that went to the fountains!

Clams are another problem. The filters used on cooling water intake systems from rivers and bays do not keep out the clam larvae, which then lodge in the cooling pipes in the plants and begin reproducing. Eventually, the pipes become clogged with thousands of clams. A report on one plant in Arkansas suggested a week-long shutdown to remove them. Clams foul non-nuclear plants too, but stopping and starting them is not as dangerous.

Even changing light bulbs has its dangers in these highly engineered, complex systems. In 1978 a worker changing a light bulb in a control panel at the Rancho Secco 1 reactor in Clay Station, California, dropped the bulb. It created a short circuit in some sensors and controls. Fortunately, the reactor scram controls were not among those affected, and the reactor automatically scrammed. But the loss of some sensors meant the operators could not determine the condition of the plant, and there was a rapid cooling of the core. As we have already noted, normally the inside temperature of the reactor vessel is at 550° F. Within an hour it had dropped to 280°. The colder, internal walls tried to shrink but the hotter, external ones would not allow shrinkage. This put strong internal stresses on the core. Meanwhile, to prevent a meltdown of the fuel rods, the internal pressure must remain high—2,200 pounds per square inch—while the temperature must drop. At the lower temperature of 280°, the

strength of the vessel is reduced, but the pressure remains high. This rapid cooling, which can occur with high pressure injection, or with a loss of instrumentation and control, did not in this case damage the core. But this is probably only because the plant had been operating at full power for less than three years. A spokesman for the NRC said: "If it had been 10 to 15 full power years, instead of two to three, which it was, that vessel might have cracked."²⁰ A cracked vessel would result in a loss of coolant and a meltdown; no emergency system would be available to cool the core.

Knowledge of such problems, after Three Mile Island, should lead to extra surveillance; we should learn from experience. The record has not been encouraging, however. The Indian Point Number 2 nuclear plant, thirty-five miles upwind of New York City, run by Consolidated Edison (Con Ed), had been having problems with leaks in the fan cooling unit service for some time. The leaks occurred in the containment building. Early in October, 1980, a light went on, warning of high water in sumps in the building, and remained on for several days. The indicator light itself was apparently considered to be malfunctioning. But water was actually leaking into the building from the fan cooling unit; eventually 100,000 gallons would collect, covering the first 9 feet of the reactor vessel in salty, brackish, cold Hudson River water. A safety device, involving two moisture-level indicators, failed to detect the water, because the indicators were designed to detect hot, not cold water.

The leak might have gone undetected for hours or days more were it not for an operator error. A warning signal came on, indicating a fluctuation in reactor power. It presumably was not related to the water leak. Operators reduced power and checked; nothing seemed wrong, so they supposed it was a faulty signal (it possibly was; they are common). But to go up to full power again, an adjustment was necessary on a governor. It was made too quickly (the operator error), and the entire reactor shut down automatically. The technicians had to enter the containment building before starting up again. They then discovered areas flooded with over 9 feet of water. The two sump pumps, which should have removed the water, were both inoperative. In one the fuses were blown, in the other the float mechanism was stuck.

We should not be aghast; these are just the routine problems of industrial equipment. But this case occurred in a building that is inconvenient to enter, and is not visited except for maintenance, unless there is trouble. The supervisor then restarted the reactor twice, without considering whether having the bottom 9 feet immersed in cold water for hours or

NORMAL ACCIDENTS

days might have led to thermal cracking or other problems. Fortunately, another supervisor, who just happened by on his day off, recognized the danger and shut the reactor down.

All this took place on a Friday, October 17, 1981, at 11:00 A.M. Contrary to an agreement with federal and state officials to notify them immediately of any trouble at the problem-plagued plant, nothing was done until 3:20 that afternoon. A plant official called the NRC's resident inspector on the site. But it was Friday afternoon and he was away. The plant official did leave a message on the answering system, but it was a message simply to call him, not one saying, "We have just found 100,000 gallons of water in the containment building." Nor did the Con Ed official call the emergency number that is given on the answering service tape. The resident inspector came back to work on Monday to find the plant shut down. He waited until 4:20 that afternoon to inform the NRC regional office of the problem. Con Ed waited another day—five all told—before it informed local officials and the public of the leak.

So it goes with organizational safeguards designed to save us from technological failures. The NRC proposed a fine of \$210,000 for the utility, which Con Ed of course protested. To replace the power during the long shutdown (they eventually concluded there was no damage to the reactor vessel), expensive oil-fired power plants had to be used, at a cost of \$800,000 a day to Con Ed's customers.²¹ Stuck floats, light bulbs, and shirt tails are just a few of the trivialities to which this system is vulnerable.

Learning from Our Mistakes

Those were simple failures or shutdowns. It is time to get more deeply into run-of-the-mill accidents. The NRC puts out a journal called *Nuclear Safety*. One of its regular features is a compilation of safety-related occurrences, selected by the editor and briefly described. Though technical, they provide endless, numbing fascination as they describe all the things that can go wrong in these awesome plants. Here is one brief account, not particularly remarkable, but it will give you the flavor. Don't try to follow it too closely; just note the failures of the equipment, operators, and design, before we turn to the journal's editorial comment that the incident shows how the industry has managed to achieve its excellent safety record.

Nuclear Power as a High-Risk System

A small, early BWR reactor at Humboldt Bay, California, (Pacific Gas and Electric) lost its offsite power source on July 17, 1970, and scrambled, as designed. The emergency power supply came on, but it was not designed to provide power to the particular sensors that turned out to be needed. Reactor pressure rose, but the emergency condenser, which would reduce it, did not come on because the gate on the switch stuck in the guides, probably as a result of a poor setting on a valve. The operators knew the emergency condenser did not operate, but assumed that a safety valve had opened to reduce pressure. Instead, a different safety valve opened, and, due to coolant shrink from its discharge, a low-water level signal came on. This, combined with the loss of feedwater and an increase in dry-well pressure, opened the reactor vent system. Meanwhile, a pipe joint ruptured in the safety valve discharge line. The vent valves were open for four minutes before the operators discovered them. There was no indication of a rupture, so they closed them. Then the fire pumps started automatically, indicating excessive pressure in the reactor, low water level, high pressure in the dry well, and loss of power to some safety systems. The accident was successfully contained, but the pressure in the reactor had exceeded safety levels; 24,000 pounds of reactor water was "blown down" (forced out of the core), indicating that the top of the fuel rods in the core were in danger of being uncovered. This was not a particularly remarkable accident; many are far worse. What is interesting is the comment that precedes it, which I quote:

The nuclear industry is not vastly different from other industries. Things do go wrong, as is attested to by these safety-related occurrences which are reported in each issue of *Nuclear Safety*. Even so, the nuclear industry has an excellent safety record. The items chosen for this article demonstrate how this record has been attained. For example, safety systems are designed with backups that take into account the possibility of failure, operations people watch for anomalies and investigate them quickly, and routine checks are run to assure that all is proceeding as planned.²²

It is hard to believe the cheerful author read his own account of the Humboldt Bay accident—or any other accounts in *Nuclear Safety*. In the previous issue of the journal, a fuel meltdown was graphically described (though in a plant in France); in the issue with the above quote we find among others a report of another plant in which, even after a seven-month shutdown for repair of primary coolant piping, an important motor broke, and sixty-three valves malfunctioned—35 percent of those tested prior to start-up. Reassuringly, we are told that "the frequency of valve testing will be increased, and a better method of cleaning the air used for some valve operations will be studied."²³

NORMAL ACCIDENTS

In the next issue of *Nuclear Safety*, after a discussion of some fires and other problems, we find the following: "A core-spray injection valve failed to close, and then it was discovered that the injection valves for the other core-spray system would not work either. Also, the valves on the low-pressure coolant-injection systems would not operate properly. While the problem of these valves was being pondered, one of four control valves on the main turbine unexpectedly closed completely with the reactor at full power."²⁴ And so *Nuclear Safety* goes on, issue after issue, editorializing how the "excellent safety record . . . has been maintained," in spite of the accounts presented in its pages.

But the following year a more plaintive note is struck. "Two-thirds of the problems discussed in this issue are strikingly similar to ones previously reported in *Nuclear Safety* in the hope and expectation that we will all be able to learn from the experience of others. . . . Operators should take particular note of these occurrences so that they can more readily avoid similar happenings in their own plants."²⁵

The following is an unusually literate and straightforward account of an incident that has many parallels in marine tankers and chemical plants. Since gases are invisible, and the subtle interactions of pressure, temperature, and operator actions cannot be fully anticipated, these events are unavoidable in highly interactive systems. In this case, after the event, two additional valves and some additional procedures were added to the system to prevent its happening again, but then no one thought it could happen in the first place; it can also happen in a slightly different way in another location of the plant.

During a shutdown, service personnel requested that demineralized water be made available in the containment in order to fill pails to be used for cleaning. The shift supervisor informed them that a valve lineup would have to be made by operations people before water would be available. When the service personnel entered the containment building, they tried the faucet to see if there was water yet. There was none, and so they closed the valve and waited a while. Then once again they tried the faucet, leaving it partially open while waiting for water, and called the main control room on the plant intercom to ask when the water would be available. Operations personnel said that a man was on his way into the containment building to align the demineralized water system. The service people closed the valve.

Shortly thereafter the radiation monitors in the containment alarmed, and the control-room operators ordered an evacuation of the containment building. The increase in radiation levels in the containment building was traced to the gas escaping from another tank, the collecting tank, during the brief period the faucet was opened. Inasmuch as several tanks are interconnected by the same demineralized water supply, the valving in the system has to be properly aligned to prevent undesirable interaction. The premature manipulation of the

Nuclear Power as a High-Risk System

sink valve before operations personnel could align the system resulted in the venting of the quench tank back through the primary water header to the open faucet. However, the radiation levels were low, and there were no overexposures.²⁶

Dresden 2 is a nuclear plant outside of Chicago that is hardly a household name for most people, but for me it holds the distinction of providing the quintessential example of a system accident. It is owned and operated by Commonwealth Edison, reputedly one of the top two utilities in the country in terms of organization and management. This is important, for it indicates what can happen even in a well-run utility. The following description is much simplified, though you would hardly believe so in reading it. Do not try to understand the complex interactions, but let yourself be overwhelmed by the operators' frequent, uncomprehending attempts to cope with multiple equipment failures, false signals, and bewildering interactions.

The steam valve began to malfunction and then closed. Fortunately, the reactor SCRAMmed automatically. The power dropped to the afterheat level, reducing the size of the steam bubbles in the core. This caused the water level in the reactor to drop, which caused the feedwater pumps to increase coolant flow into the reactor to avoid uncovering the core. As the water level rose, the operator noticed that the level indicator was reading a low level. Actually, however, the indicator was stuck and giving a false low-water-level reading. The operator reacted by manually increasing the feedwater flow still further, so that the water then filled the reactor and spilled over into the steam line. The feedwater-flow error was uncovered and corrected; but then the pressure began to rise, and two safety systems designed to cope with the problem and cool down the reactor were found inoperative. The operator then reduced pressure by opening a relief valve momentarily. At this point, water hammer occurred, produced by the water spill-over into the steam line, and this popped safety valves (pressure relief valves), which stuck open due to a design error. The relief valves then discharged reactor steam to the reactor containment atmosphere, which began to pressurize the containment. The loss of coolant through the stuck relief valves should have caused the ECCS to activate to inject replacement coolant; but one system was found inoperative, and the operators blocked the operation of the other system on the assumption that the loss-of-coolant problem was minor. However, they did not know the cause (stuck valve) and could not make a sound judgment (it could have been a leaky coolant pipe about to completely rupture). Meanwhile, the pressure in the containment rose beyond the range of the pressure gauge (5 psig). The containment is equipped with water sprays to quench the steam pressure whenever two psig pressure is exceeded, but the operators blocked this safety action because that would have cold-shocked some equipment and thereby damaged it. They did not, however, have sufficient knowledge of the events to justify their action. The containment reached 20 psig compared to 60 psig design pressure before the plant was finally brought under control.²⁷

Fermi

Our final example is not, strictly speaking, of a system accident, but of a component failure accident, though the recovery effort involved some of the typical complexities of the system accident. (The distinction between system accidents and component failure accidents is fully developed in Chapter 3.) The account, based largely on a book by John Fuller, of the Fermi core meltdown will serve to illustrate dramatically the complexity for these systems, the pressures on operators, and the tremendous problems of clean-up. It also shows that attempts to make the system safer are sometimes ill-conceived and add danger; that completely novel accidents make the question of operator error irrelevant; and that the industry, instead of worrying about the disaster potential, only draws strength from the fact that it was not worse. The accident occurred in a demonstration reactor on Lake Erie in the small community of Lagoona Beach, near Monroe, Michigan—which is very near Detroit. A report by the Atomic Energy Committee completed before the accident (and promptly classified) predicted that, given a severe accident at Fermi with unfavorable wind conditions, 133,000 people would receive high doses of radiation, and one-half would quickly die. Another 181,000 could receive 150 rads.²⁸ As Fuller's account makes clear, it was a close call.²⁹

The reactor was a sodium-cooled breeder reactor, large enough to produce substantial power as well as plutonium that could be used to fuel other conventional reactors. It was the first and only U.S. breeder reactor, and thus an untried design near Detroit's millions. In October of 1966 the operators were trying to achieve the first stage of a high power goal set by the company, slowly and carefully bringing up the temperature of the reactor. Delays and problems had been numerous in the past and continued this time. One of the steam generator valves malfunctioned, and six hours were spent correcting it. Then a boiler feedwater pump failed, but was quickly corrected. The operators once again increased the fissioning in the reactor. But the engineer on duty noticed some erratic changes in the neutron activity of the fission process, which could have been merely the electronic system picking up some "noise" or static. They paused, it disappeared, and they continued. Next, the engineer noted that for the amount of power the reactor was producing, the control rods, which shut off the fissioning when fully inserted, should have been raised only 6 inches, but they were 9 inches out of the core, and the neutron activity signal was again erratic. The reactor was put on hold, and the engineer went to check the instruments on the individual

subassemblies of the fuel rods—some 30 feet away from the control board. The results were puzzling. The outlet temperature of one of the subassemblies was clearly too high, but they had been having trouble with that one. Indeed, since it appeared that the instrument was faulty, they had moved the instrument to a different part of the fuel bundle. But now a second subassembly showed high temperatures too, but none of the ones that were nearby and also instrumented were abnormal. Unfortunately, only one of every four subassemblies were instrumented, but if one overheated those near it that also had instruments should show overheating.

Then radiation alarms went off, the air horn began blasting twice every three seconds, and the public address system came on with a laconic, "How hear this. Now hear this. The containment building and the fission product detector building have been secured. There are high radiation readings, and they are sealed off. Do not attempt to enter. Stay out. Both buildings are isolated. This is a Class I emergency. Stand by for further instructions. Stand by for further instructions."³⁰ The operators first counted the crew to make sure that no one had been sealed in the containment building, with its high radiation readings. Everyone was safe. Next they had to pull down the power in the reactor. They were reluctant to scram the reactor immediately because of thermal shock from a sudden change in the temperature of the sodium coolant. One hypothesis that was quickly ruled out was that the radiation alarms were false; an engineer had been working on the fission product monitor and thought that he might have triggered a false alarm. But the temperature readings on the subassemblies indicated something real was going on.

Eleven minutes after they started to cool down the reactor, they decided to scram it manually. There was no way of knowing whether this was too late, or too soon, with this type of reactor. Indeed, there was no way of knowing what had happened inside the core. According to the design of the safety features, if there were any fuel melting, the reactor should scram automatically. It obviously hadn't scrambled, which suggested that the problem was not a fuel melt. Yet perhaps it was a fuel melt which just was not indicated. On the other hand, it might be an instrument problem; there had been problems with the instruments before. Still, they could not be sure if a fuel melt had been avoided or not. It was essential to find this out, because a fuel melt would block the flow of sodium coolant, and could lead to heat build-up and more melting, and thus a secondary accident. The assistant general manager took charge of the immediate efforts and announced, "We will go at this very, very slowly."³¹ Fortunately, the Fermi engineers had time, since the core tem-

perature continued to slowly decline. Since there were no procedures for such an emergency, they had to write ones out, and check them very carefully. They feared stirring up trouble inside the core. They soon found there had been both fuel melting and fuel redistribution. The redistribution could cause blockage and further fissioning.

The fuel melting conclusion was no doubt reluctantly arrived at for another reason—expert advice. Nobel Laureate physicist and nuclear power advocate Hans Bethe had confidently predicted a core meltdown could not happen with this reactor. Another expert was less certain, but he had predicted that at worst only one subassembly could melt. The evidence now was that two or more had melted. The second expert had also stated that the automatic safety devices would shut the reactor down if there were any melting; the devices did nothing of the sort. The Fermi engineers now talked of “hair-raising decisions” and “terrifying thoughts”; they were sitting on top of a volcano next to Detroit. They could not walk away and leave it there; they could not be sure there would not be a secondary accident, and in any case the melted uranium would eventually eat through the core and the concrete base of the building.

For a month the reactor sat there while the company let it cool and planned the next step. Then the engineers very carefully removed the top and hoped that none of the fuel subassemblies were stuck together in such a way as to produce “criticality” (the conditions for fissioning). If they could pull out the damaged subassemblies, it would be safe. It took three months to learn that four were damaged, and two stuck together. It took five more months to remove them. Special equipment was built; the deadly sodium had to be drained, and there was no provision for this in the reactor design. Almost a year from the accident, they were able to lower a periscope 40 feet down to the bottom of the core, where there was a conical flow guide—a safety device similar to a huge inverted ice-cream cone that was meant to widely distribute any uranium that might inconceivably melt and drop to the bottom of the vessel. Here they spied a crumpled bit of metal, for all the world looking like a crushed beer can, which could have blocked the flow of sodium coolant.

It wasn't a beer can, but the operators could not see clearly enough to identify it. The periscope had fifteen optical relay lenses, would cloud up and take a day to clean, was very hard to maneuver, and had to be operated from specially-built, locked-air chambers to avoid radiation. To turn the metal over to examine it required the use of another complex, snake-like tool operated 35 feet from the base of the reactor. The operators managed to get a grip on the metal, and after an hour and a half it was removed.

The crumpled bit of metal turned out to be one of five triangular pieces of zirconium that had been installed as a safety device at the insistence of the Advisory Reactor Safety Committee, a prestigious group of nuclear experts who advise the NRC. It wasn't even on the blueprints. The flow of sodium coolant had ripped it loose. Moving about, it soon took a position that blocked the flow of coolant, causing the melting of the fuel bundles.

During this time, and for many months afterwards, the reactor had to be constantly bathed in argon gas or nitrogen to make sure that the extremely volatile sodium coolant did not come into contact with any air or water; if it did, it would explode and could rupture the core. It was constantly monitored with Geiger counters by health physicists. Even loud noises had to be avoided. Though the reactor was subcritical, there was still a chance of a reactivity accident. Slowly the fuel assemblies were removed and cut into three pieces so they could be shipped out of the plant for burial. But first they had to be cooled off for months in spent-fuel pools—huge swimming pools of water, where the rods of uranium could not be placed too close to each other. Then they were placed in cylinders 9 feet in diameter weighing 18 tons each. These were designed to withstand a 30-foot fall and a 30-minute fire, so dangerous is the spent fuel. Leakage from the casks could kill children a half a mile away. It took three years to remove the poisonous materials from the plant and to seal the radioactive sodium up in steel drums for storage at the site (none of the six burial grounds in the country would take it) where it will have to be monitored for generations. The plant, incredibly enough, was re-commissioned some years later and operated at low power for a short time. It was finally permanently decommissioned after more troubles.

This account clearly illustrates some of the principles investigated in this book, as can be seen below:

1. The problem originated with a safety device. Indeed, installation of the device was prompted by concerns of a prestigious committee made up of nuclear scientists and engineers, many from the elite universities, responsible for advising the NRC on safety matters. They were worried about a fuel drop, and the sheets were part of the response.³²

2. Poor design and negligent construction led to the accident. Though it did not start with diverse failures, it is hardly reassuring that the sheets were poorly secured and the force of the surging coolant not anticipated, and the addition left off the final drawings.

3. As in other accidents, some parties were to suggest operator error, when in fact there was no clear procedure to follow; nothing like this had been anticipated. R. L. Scott, in his account of the accident for *Nuclear*

Safety, hints that one of the major problems was a failure of the operators to scram the reactor immediately. But some of the technical papers that he cites in his article point out that there was insufficient information available for the operators to know what the danger was and what was going on.³³

4. Finally, we should note once again that those attached to high-risk systems can be uncommonly cheerful about these system failures. Scott is pleased to point out in the NRC journal that the melted fuel resolidified only a short distance from the hot spot, and did not cause the melting of adjacent subassemblies. This should give us more confidence, presumably, in breeder reactors. He next tells us, "Much additional benefit was derived from the recovery operations . . . not the least of these was the experience gained by the personnel directly involved." We may be very happy that these personnel had their experience increased, but unhappy that most of Detroit had to be at risk to secure the gain. He goes on, "Many innovations are required to cope with the new and different problems that presented themselves."³⁴ As an example of positive thinking about this, he lists the number of changes subsequently made in the system, such as a provision for draining the radioactive sodium from the reactor vessel. One would hope that a serious accident would not be required to bring this matter to the attention of designers. Finally, he cheerfully concludes that "the Fermi fuel melting incident [sic] has been quite instructive, emphasizing the need for design provisions for inservice inspection and the desirability for a simple, rapid presentation of critical operating information to the operator, together with adequate procedures and precise criteria for operator action."

In our accident is our salvation.

The Fuel Cycle as a System

We have treated the plant as the unit of analysis, and generally will continue to do so as we investigate other high-risk systems. But nuclear power involves the whole "fuel cycle"—the sequence from mining uranium ore, processing it into fuel, burning it in reactors to boil water, and the disposing of the many kinds of wastes. All of these involve serious hazards. Indeed, while we will not discuss the waste problem in this book, it probably has a greater long-run catastrophic potential (if we include military wastes) than nuclear plant operation. Mining is probably

responsible for more radiation-induced deaths than any other part of the cycle to date (for the waste problem will take much longer to reveal itself), although these deaths are generally not the result of system accidents. But system accidents do occur in the fuel processing stage. A quick look at this stage will suggest that the processing of dangerous materials is, as a rule, associated with system accidents. These accounts are included for another reason: to point out the trivial details that can have large consequences, and the lack of understanding still evident in a production process that is well beyond the research stage.

Thirteen accidents involving fabrication of fuel are described in a *Nuclear Safety* article.³⁵ Some appear to be due to carelessness or inadequate technology. For example, there is the spontaneous ignition of contaminated wastes that are unaccountably stored in cardboard cartons in a waste storage room. Part of the plutonium released from this fire was washed from the building by the fire hoses, contaminating the surrounding ground. In another case, plutonium-casting residues were placed in a plastic bag, and burned through. In another, a five-year-old filter "heavily loaded with plutonium dust" caught fire from the sparks of a welding torch.

Cleanup is difficult when radioactive materials are involved. In an explosion at the Oak Ridge National Laboratory on November 19, 1959, "buildings and nearby streets were contaminated by the air flow through open pipes and other cell wall penetrations." The streets had to be scraped up. But the author of the *Nuclear Safety* article is reassuring. He concludes that "in all plutonium incidents to date, only a small fraction of the plutonium involved was released."³⁶ That is like saying that in a war, only a small fraction of the bullets kill anyone.

A bit more revealing is another discussion of seven "criticality" accidents. If plutonium, which is exceedingly volatile and hard to machine or handle, experiences the proper conditions, it can attain a self-sustaining fission chain reaction. Criticality depends upon the quantity of the plutonium, the size, shape, and material of the vessel that holds it, the nature of any solvents or dilutants, and even adjacent material, which may reflect neutrons back into the plutonium. It is apparently hard to know when these conditions might be just right. In the seven critical accidents that occurred between 1958 and 1970, fifteen workers were reported as receiving significant degrees of irradiation (an average of 140 rems, while the current legal yearly maximum for nuclear personnel is 5 rems) and two more died within two days of an accident.³⁷

The accidents reveal the highly interactive nature of the systems. In one case, two poorly working pumps were involved, along with a line

that may have been plugged. In the attempt to free the line, a bubble of high pressure air was created, though no one knew it. This forced 40 liters of a solution up a 5-inch-diameter storage pipe and out into another vessel that just happened to have the proper dimensions for criticality, given this particular solution and its volume. In another case, a plug of uranium nitrate crystals was found in a line. Operators dissolved it with steam, but the liquid was then drained into some available bottles, which just happened to be identical with those used to store a much safer liquid. One of the bottles, now containing U-235, was poured into a make-up tank. After stirring was commenced, it blew up, knocking the operator to the floor. He managed to escape the building but died forty-nine horrible hours later. Two operators went in to drain the solution into safe containers, but in turning off the stirrer, apparently (and who can know with this technology) the change in geometry added enough reactivity to again produce criticality. The operators did not know this had happened because the alarm that would indicate the danger was still sounding from the first "excursion." These two men received dosages of from 60 to 100 rads. (A total of 50 rads is the exposure level needed to double the risks of genetic defects, and is the legal maximum accumulated dosage for nuclear workers over twenty-seven years of age.)

A government report, WASH-1192, documents 111 accidents involving unplanned release of radioactivity that exposed 317 people to excess radiation from a few to as many as 80,000 rads. These occurred between 1959 and 1970.³⁸ The average dose of workers at the West Valley Reprocessing Plant (now closed), near Buffalo, New York, run by the Getty Oil Company, was 6.7 rads in 1971 and 7.1 rads in 1972, well beyond the legal minimum dose.³⁹ These assaults upon personnel, most of which will not reveal their damage for two decades, are not considered in the statistics that show that nuclear power is "safe."

As we have just seen, the power generation phase of the nuclear cycle is not the only one prone to system accidents; the fuel processing and reprocessing systems are at risk also. Transformation processes in the nuclear fuel cycle seem to have an inherent degree of unpredictability. But two questions still remain. How frequent are system accidents, and could not management learn to at least prevent the minor failures that can occasionally come together to create a system accident? We can only guess at the first, but there is unmistakable evidence regarding the latter.

Can We Handle It?

I hope I have convinced you of the frequency of serious accidents or near accidents in nuclear plants, and the existence of system accidents in the above examples. How many system accidents there are is impossible to tell; the reports in *Nuclear Safety* are often not detailed enough to judge. One serious attempt to analyze accidents in terms of the multiplicity of failures and the variety of component failures supports the argument of this book. Morris and Engelken examined eight Loss of Coolant Accidents (LOCAs) in BWRs in a two-year period when there were only twenty-nine plants operating. They occurred in six different BWRs. The authors estimate there will be a LOCA for each two reactor years of operation.

They conclude that "No two of the incidents were initiated by a common system or component malfunction. . . . The reactor primary coolant was released during these transients through safety and relief valves that either operated prematurely or that operated correctly but failed to close."⁴⁰ So each accident was unique, and each involved, among other things, the failure of a key safety device. These could easily be system accidents. In their summary of the eight occurrences, they identify eight categories of failures (such as, valves lifting below the set point at which they are supposed to lift; valves failing to reseal; flooding of steam lines; isolation valves closing too soon; condensor malfunctions; violation of operating procedures). Each accident involved from two to four of these failures. In half of the eight accidents there were violations of operating procedures, but they always occurred in conjunction with at least two and as many as five other failures. Failures not only were spread over the eight categories, they were widespread for all vendors and manufacturers. In one sample of valves, 15 percent were thinner than the design specified. Deficient valves were found at twenty plants owned by fifteen different utilities and supplied by ten different suppliers.⁴¹ It is from such minor failures that system accidents can grow; the uniqueness of the accidents and the multiplicity of failures in this survey of just one system suggests that system accidents are not all that rare.

Could not management prevent these failures? The authors' investigation of eight LOCAs led them to a broad indictment of management practices and to a further study. Frequently, "abnormal situations and incidents . . . have not been thoroughly investigated"; "minor abnormalities were often ignored or their implications not understood, and these sometimes led to more serious conditions." It is in minor abnormalities, we might interject, that the system accident is spawned. They note, "Of

course, there is always a very strong incentive to keep the plant on line, i.e., producing power." Finding all this, they conducted a management appraisal program in the plants, and found that "even though a deliberate effort was not made to look for violations," there were seventy-five in only seven appraisals. Of these, eighteen were failures to test "vital safety equipment." "In summary, there has been a surprising lack of knowledge, understanding, and effort by some utility executives to discharge their own responsibilities and those imposed by the specific requirements of an Atomic Energy Commission license."⁴²

Strong stuff. But that was in 1972, and the industry was young. Since then, there have been major accidents at Dresden, Browns Ferry, and TMI, and several critical reviews of performance. Yet in an NRC review of operating plants, conducted in 1980 as a result of the excoriating criticism of the NRC by the Kemeny Commission, little seems to have changed. In possibly the most dangerous industrial activity that humans have yet to engage in, the study described the twenty-one "below average" facilities in numbing, repetitious terms: inadequate technical staff, insufficient training, poor supervision, failure to follow procedures, radiation protection weaknesses, incomplete licensee event reports and failure to consider their implications, unmonitored and uncontrolled release of airborne radioactive material, noncompliance with quality assurance programs, inadequate control over liquid and solid radioactive waste, repetitive equipment problems, problems in management coordination and attention, inadequate fire protection, failure to meet commitments made to the NRC, "repetitive instances of system misalignments, impaired ECCS equipment operability and containment integrity," personnel overexposure, and longstanding and uncorrected design problems. Most of these items appeared several times.⁴³

We are not told anything about the average plants; the above list refers to the 29 percent found to be below average. Let's take a plant studied by the NRC from May 1979 to May 1980 and rated as average—San Onofre, owned by Southern California Edison. It is a small plant, 436 megawatts, which has been in operation for thirteen years. (Actually, in those thirteen years it has operated at full power for only 8.8 years, or 68 percent of the time, which is above the industry average.) In 1980, it tied for first place in the newly established NRC category of having "especially significant mishaps" (serious incidents). It is one of the eight plants listed by the NRC as having the most serious weakening of steel, which could cause the core vessel to crack. Some of the recent problems: In November 1979, a nest of field mice (a sign of poor housekeeping) caused an electrical fire that shut the plant down for a week and cost \$2 million.

From April 1980 to June 1981 it was shut down for steam generator repairs (a problem that plagues all power plants), costing \$68 million; the repairs are good, at best, for five years. During the overhaul, seventy-three workers were overexposed to radiation (the NRC fined the utility \$100,000 for that and \$50,000 for additional violations and exposure to workers). Fifty truckloads of radioactive sand had to be removed from the ocean beach in front of the plant in May 1981. A fire in an auxiliary diesel generator shut the plant down for four weeks July 1981, costing \$2.5 million in repairs. During this time there was an explosion in a radioactive gas holding tank with a release of 8.8 curies of radioactive Krypton gas to the atmosphere. In September of 1981 a failure in a voltage regulator was investigated and the company found inoperative valves in the ECCS—estimated, by the NRC, to have been inoperative since 1977. This resulted in a finding by the Commission of "deficiencies in management and procedure controls." The NRC estimated that the unit may become unsafe to operate by 1983 due to "embrittlement" problems that could crack the core vessel.

Much of this took place just after the evaluation that rated it as average, rather than before or during it. But even in 1980, while the evaluation was going on, there were thirty-seven safety-related failures they were required by law to report to the NRC and seven "especially significant mishaps."⁴⁴ If this is an average plant, those below average by the NRC's standards might give their neighbors cause for concern.

Well, if things have not improved much since the 1972 study of a few plants, judging from the NRC's 1980 evaluation, perhaps we shall do better with the plants being built and those about to come on-stream. The Diablo Canyon plant of Pacific Gas and Electric has been ready for a long time, and in 1980 the NRC also evaluated it, along with seventy-five others that were in various stages of construction. It rated Diablo Canyon as average (the highest rating given for those under construction). Yet the next year an engineer in the utility accidentally discovered that the required earthquake reinforcements of key equipment had been incorrectly installed, as we noted earlier, and then 111 other violations were found. Something similar had happened to a second unit being built at San Onofre, where the reactor was installed 180 degrees out of alignment, and it took Southern California Edison seven months to discover the error. To correct it they reversed the wiring in the control room; but it was not that simple at Diablo Canyon.

Conclusion

We have not had more serious accidents of the scope of Three Mile Island simply because we have not given them enough time to appear. But the ingredients for such accidents are there, and unless we are very lucky, one or more will appear in the next decade and breach containment. Large nuclear plants of 1,000 or so megawatts have not been operating very long—only about thirty-five to forty years of operating experience exists, and that constitutes “industrial infancy” for complicated, poorly understood transformation systems. There is ample evidence that problems abound in these large systems, and that they are different from the problems of the smaller units where we have a bit more experience. For all nuclear power plants, the steam generator and the core embrittlement problems are awesome. Small failures can interact and render inoperative the safety systems designed to prevent a steam generator failure from being catastrophic. Trivial events can place stress on the embrittled core in ways unimagined by designers. The sources of other errors and failures appear all too numerous, judging from the events covered in this chapter.

The catastrophic potential of nuclear plant accidents is acknowledged by all, but defense in depth is held by experts to reduce accident probabilities to nearly zero. Yet core containment, emergency cooling systems, and isolated siting all appear to be inadequate; all have been threatened. Nor can we have any confidence whatsoever that quality control in construction and maintenance is near the heroic levels necessary to make these dangerous systems safe. A long list of construction failures, cover-ups, threats, and sheer ineptitude plagues the industry. I have argued that construction problems are probably no worse than in most other industries, but that is no comfort; it has to be much better. Nor has the actual operation of nuclear plants appeared to be as far above normal industrial standards as would be required of such a dangerous undertaking. If anything, it is somewhat below industrial standards. These statements regarding construction, maintenance, and organizational management are based upon the reviews and statements of the Nuclear Regulatory Commission itself, including its chairman. Finally, a review of some of the serious accidents that have occurred reveals the complexity of the plants, the difficulty of recovery from minor accidents so that they will not become major ones, the unlikelihood that the industry will even learn from the accidents, and the sanguine and casual response of the industry and the NRC.

When the Kemeny Commission was writing its final report, its members debated two key issues at length: Are these plants different from other industrial plants, and thus need to be judged by different criteria; and if they are different, what kind of organization is required to run them safely? A group of pro-industry members of the panel argued first that the plants are not different, and the restrictions being considered were not necessary; they then argued that if there were unique dangers, the plants should be run on a paramilitary basis. This position frightened other Commission members and led them to ask if a peacetime economy really needed an authoritarian, dictatorial segment managing a system with such catastrophic potential. These are heady issues, going beyond shoddy construction, inept management, untried and hasty designs. We will discuss these issues at length after we have reviewed other high-risk systems such as nuclear weapons and DNA engineering.

Yet despite the glaring failures of the nuclear power industry, it is clear that its design, construction, and operating problems do not, in themselves, constitute the cause of system accidents. It is instead the potential for unexpected interactions of small failures in that system that makes it prone to the system accident. Some systems with catastrophic potential are not liable to these complex failures; their accidents have different, more mundane sources. Some highly interactive systems are without catastrophic potential. To tread our way through these complexities, we need some careful analysis and more precise terms and concepts. That is the task of the next chapter.