plant, but throughout the Soviet design, operating, and regulatory organizations that existed at the time." More important, at least in regard to the safety of the TVA reactors that are to be used for tritium production, is that study's emphasis on management's role in shaping the motivational environment: "Safety culture . . . requires total dedication, which at nuclear power plants is primarily generated by the attitudes of managers of organizations involved in their development and operation."[16]

In the U.S. nuclear establishment, such abstract concepts were at first no better received than Perrow's systems theories. The initial emphasis for the NRC and the nuclear industry after the TMI and Chernobyl accidents was to find practical ways to accommodate, without having to shut down any power plants, the new reality that severe accidents could really happen and were really dangerous.

The New Science of Severe Accidents

In hindsight, the fundamental flaw in the government's approach to reactor safety was the artificial distinction it made between "credible" and "incredible" accidents, particularly since it used a rather soft connotation of the word "credible." The formulators of the original licensing approach knew core melt accidents were possible: an accident at one of the experimental reactors had shown that.[17] But they felt that reactors could be made adequately safe by protecting against all the accident scenarios that were likely enough to be worth thinking about. Intuitively, they judged that core melt accidents didn't fall into that category. They looked at safety from a black-and-white perspective: some accidents one worries about, some are so unlikely that one doesn't. What they learned from TMI and its aftermath was that the measure of importance of a particular type of accident is not just its probability but the *product* of its probability times its consequences. On that basis, core melt accidents, unlikely as they are, dominate nuclear power's threat to public health.

A simple example may help. If a hang glider manufacturer hired a safety engineer to evaluate a new design for its safety characteristics, a fairly straightforward approach would suffice. The engineer would study the kinds of accidents that had happened with other hang gliders, analyze what went wrong in most cases, and modify the design so those things wouldn't go wrong with the new design. What is straightforward about this task is that the event to be avoided is a crash—any crash. In a sense, all crashes are equal, because they result in the same consequence: the injury or death of one person.

In contrast, a nuclear power plant accident has the potential to cause fatalities ranging in number from none or a few to tens of thousands. This wide range of outcomes changes how designers must think about safety. For example, if a particular type of event that occurs only once in 10,000 years causes 10,000 fatalities, its contribution to public risk is one fatality per year. This would be far more important than a different type of event that occurs every ten years, but on average causes a fatality only 1 percent of the time (its contribution to public risk is 0.1×0.01, or 0.001 fatalities per year).

Reactor safety analysts now know that nuclear power plants are a special type of complex system whose risk to the public is dominated by "low-probability, high-consequence events." For such systems, safety engineers are not focusing on the right issues if they arbitrarily ignore all low-probability events. Because of the huge range of accident consequences, it is necessary instead to think about the *risk profile* over a wide range of event probabilities.

Early in the history of the nuclear industry there were hints of this reality. A study done for the AEC in 1957 estimated that a severe accident could cause 3,400 deaths and cost $7 billion in property damages.[18] Eight years later the study was revised to account for the larger plant sizes being considered at that time. The new numbers were so shocking that the AEC suppressed publication of the report. The chairman at the time later said, "we didn't want to publish it because we thought it would be misunderstood by the public."[19] The truth appears to be that the AEC misunderstood it.

Nuclear reactors have some of the most reliable and redundant safety systems of any man-made facilities. But as TMI showed, systems can fail, people can err, and sometimes more than one thing can go wrong at a time. Despite clever design, careful manufacture, and dutiful maintenance, multiple failures of safety systems can and will occur.

Severe reactor accidents—those involving melting of the core—are expected to be quite rare, occurring in the United States less than once in a generation, perhaps. But if the health consequences are large enough, they are a more important threat to the public than more likely but more benign accidents. Ignoring core melt accidents in designing reactors was therefore a huge mistake.

Before TMI, there was recognition by some that core melt accidents dominate public risk. For example, in 1974, the AEC commissioned Norman Rasmussen of the Massachusetts Institute of Technology to lead a major study[20] that systematically evaluated public health impacts of nuclear reactors, focusing necessarily on the improbable but highly consequential "beyond–design basis" accidents. It was a seminal study that after TMI became the starting point for greatly expanded efforts. But the AEC's goal for the Rasmussen study was primarily to deal with growing public opposition to nuclear energy by comparing its risks quantitatively with other risks that society routinely accepts.[21] It was never intended to modify the government's approach to reactor design or regulation.

After TMI, everything changed. Now the question was, how could these risk-dominant events be factored into regulation? The new challenge to the NRC was to sort out what types of potential accidents were "risk-significant," having the property that their probability times their consequences was large, and to decide what should be done about them.

In the early 1980s the NRC embraced this new perspective with grim diligence. Great sums for new research were requested from Congress and granted. To establish a higher degree of credibility for the new research, the NRC engaged the services of the DOE's great nuclear weapons laboratories, such as Sandia and Oak Ridge. In parallel, the nuclear industry formed the Industry Degraded Core Rulemaking (IDCOR) consortium to supplement the government program and to ensure that if "rulemaking" (that is to say, new regulatory approaches) were to take place, the industry's interests would be well defended. Eventually severe-accident research became an international enterprise, as Japan and Europe established their own programs (primarily after the great scare they got from the horrendous Chernobyl accident).

The NRC's severe-accident research program had to find new ways of thinking about reactor safety and to develop entirely new fields of science and engineering. The probabilisitic methods pioneered in the Rasmussen study were expanded and computerized, leading to a new field of statistics called "probabilistic risk assessment." Specialists in the new field developed sophisticated computer programs that categorized various types of abnormal events and followed the causal connections from one event to another, using probabilistic treatments when it was not known which of several outcomes might occur. The new science of severe accidents spawned a new mathematics of probabilistic calculation that today is being applied to all sorts of other enterprises besides nuclear power, such as evaluating the safety of nuclear weapons.

In addition to the new probabilistic techniques, effective safety assessment required information about what can happen, physically, when the core begins to melt and relocate and when core debris breaks through the reactor vessel boundary and enters the containment. Since it is not possible to carry out experiments on severe accidents with real reactors, a variety of test facilities were needed that could simulate the extreme conditions of a severe accident. The NRC (and its counterparts all around the world) made large investments in such facilities, and the experiments performed in them yielded many new insights about what might occur in severe accidents.

A third element of the NRC's new research program was the development of computational simulation software. When low-probability events are considered, the possible combinations of conditions and configurations are overwhelming. It would be impossible to conduct experiments for all possibilities. For that reason, the NRC sponsored the development of sophisticated computer programs that could calculate how conditions in a hypothetical accident would evolve. These "codes" used theoretical models for the physical processes involved but were grounded in reality by checking how well they could predict the results of experiments in the severe accident test facilities. As time went on, hundreds of person-years were invested in these codes, and in a way they took on a life of their own, embodying hundreds of thousands of lines of computer code and becoming more complex than any one individual could comprehend. They had names like MELCOR, for modeling core melt progression, and CONTAIN, for studying events in the containment.

Within a few years, the NRC's research program had yielded abundant insights about the nature of severe accidents and nuclear power plant risk. If this knowledge had been obtained before all the U.S. plants had been designed and built, the plants could have been designed to be intrinsically resistant to core melt accidents. Instead, the agency in 1983 or so was left with few winning strategies. Certainly, shutting down all nuclear power plants would have been foolish: all risk studies showed (and still show) that nuclear energy is on a par with other energy sources as far as risk to public health is concerned. The agency's challenge was instead to steer an objective course, working with the industry to minimize wasted regulatory effort, while at the same time maintaining a high degree of vigilance and integrity in dealing with the issues that truly affect public health and safety.

It became apparent in the mid-1980s that studying severe accidents was something that could be funded, delegated, and executed readily enough, but doing something concrete with the results of such studies posed many dilemmas for a federal regulatory agency. Prioritizing actions was particularly challenging. Early regulatory responses to TMI were often ad hoc and inconsistent and met with hostility from the industry. It seemed that the only rational basis for decision making would be systematic, quantitative assessments of public risk based on the knowledge and tools developed in the research program.

The new probabilistic tools measured public risk in quantities such as the average number of fatalities per year of reactor operation. Balancing risk reduction against the cost of changes was an appealing decision strategy, but adequate studies of the risk posed by reactors were not available. The NRC's research organization therefore embarked on the extraordinary NUREG-1150 project, so named from the document number of its major report, *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants.*[22] This massive, coordinated assessment of quantitative risk and its uncertainty took many years and hundreds of specialists to complete, but it established a firm foundation for factoring severe-accident issues into regulatory practice.

The NUREG-1150 study analyzed five different plants, each representing a different containment concept and together spanning the entire fleet of U.S. nuclear power plants. The five types and the plants chosen to represent them were as follows:

- large dry containment (the Zion plant)
- large dry subatmospheric containment (the Surry plant, whose containment is held at negative pressure compared to the outside environment)

- Mark I suppression pool containment (the Peach Bottom plant)
- Mark III suppression pool containment (the Grand Gulf plant)
- ice condenser containment (the Sequoyah plant)

Although the study took far longer and cost far more than originally planned, the picture that emerged from the project had, by 1987 or so, established the battle lines for the emerging conflict with the nuclear industry over severe accidents.

Resistance Builds

The history of commercial nuclear power is notable for the contrast between the rosy picture motivating the early headlong plunge into a nuclear future and the bitter experience of the people who signed up for it. One after another, the assumptions that lay behind that rosy picture were found to be invalid or were made invalid by government action. The federal government did not, for example, find a solution to the spent-fuel disposal problem, and its decision to abandon the plutonium fuel cycle (see chapter 2) accelerated the nuclear waste crisis. Extraordinary levels of inflation in the 1970s and 1980s were punishing to the capital-intensive economics of nuclear power. In many cases local opposition to power plants near population centers forced more time, attention, and money to be spent on public safety than had been envisioned. Then, of course, TMI and Chernobyl occurred, with devastating impacts on public acceptance. In contrast to the land rush atmosphere that began around 1965 (199 plants ordered in ten years), new plant orders dried up completely by 1977, and construction of many previously ordered plants was canceled.

In the years following TMI, the industry gamely prepared for changes in regulations and even in the plants themselves, since the accident had revealed deep flaws in the safety philosophy of commercial nuclear power. But the early regulatory actions by the NRC were often poorly motivated and implemented with a heavy hand. The industry felt unnecessarily burdened with new regulations that appeared to add nothing to safety or reliability. By the time the NRC established, in the form of the NUREG-1150 reports, a solid scientific basis on which to ask for safety improvements, the atmosphere between industry and the regulators had become adversarial and bitter.

This is not to say that common ground was not found for some changes. Through their own research and that of the government, many utilities saw ways to modify their plants to reduce the likelihood of a severe accident in the first place. These insights came from the so-called front end of the probabilistic risk analyses: the analysis of how various failures of control systems and safety features interact to create a core melting situation. Even if no serious harm were done to the public (as was the case with TMI), a core melt accident would almost certainly end the life of the reactor as a producer of electricity. Hence utilities had a double incentive—public safety and capital preservation—to improve the front-end safety picture of their nuclear power plants. Consequently, a broad range of plant changes were made, often at the initiative of the plant owners, to reduce the likelihood of a severe accident.

When it came to the "back end," that is, those events that might occur after core melting, the industry's motivation to make changes was not as strong. The costs reactor owners would incur to modify the containment buildings to better protect the public would do nothing to restore the lost revenue stream from the ruined plant. Recognizing this reality, the government chose to organize its NUREG-1150 study around the five major containment types, assessing in detail both the front end and the back end of a representative plant for each type. In 1987 the stage was set for a series of bitter conflicts over containment improvements between the NRC and the industry it regulated.

Of the five containment types studied in NUREG-1150, two were notable for their poor performance in severe accidents. Those two were the "ingenious containments" first developed as integrated features of the power plant to save costs and shorten construction time: the Mark I pressure suppression containment and the ice condenser.

To translate research results into positive action, the NRC initiated the Containment Performance Improvement (CPI) program in the late 1980s. The idea of the program was to review from a generic (i.e., not plant-specific) standpoint what changes in each of the five containment types might be justified in the interest of public safety. Rather than taking on all five types at once, the program was to address one at a time, from the most problematic to the least.

The performance numbers for the Mark I plants and the ice condenser plants were fairly similar,[23] but the CPI program decided to take on the Mark I plants first. There was a good reason for giving these plants priority: twenty-four reactors in the United States used General Electric's Mark I design, whereas only eight used Westinghouse's ice condenser concept.[24]

The Mark I severe-accident problems arose from a variety of event scenarios, mostly due to the small size of the containment. One of the most troublesome scenarios involved the molten core, which analysts had come to call "corium," penetrating the bottom of the reactor vessel and pouring onto the containment floor. The corium would be so hot in such an event that it would flow essentially like water, and if enough flowed out of the vessel, it would form a large, shallow pool that would come into direct contact with the steel liner of the containment building. No steel can stand up to such temperatures, and numerous calculations indicated that the core material would soon melt through, creating a pathway from the highly radioactive containment atmosphere to the external environment.

Such hypothetical scenarios and the theoretical calculations surrounding them were the grounds for fierce arguments between the safety analysts of the NRC and those of the nuclear power industry. Years of bitter controversy passed, but in 1989 the NRC issued guidelines that recommended specific (and expensive) changes to the Mark I containment buildings. Legally, the NRC couldn't mandate the changes, but the language they used carried a sufficiently threatening tone[25] that most of the Mark I owners made at least some of the specified modifications.[26] But the battle had repercussions that went far beyond the technical arena.

For years, Congress had heard the nuclear industry's complaints that it was being regulated to death. The dire economic state of nuclear power certainly supported the claim of imminent demise, and the charge of overregulation resonated with the political atmosphere in Washington during the administration of President Ronald Reagan. The NRC's response to pressure from the industry, Congress, and the administration to address these complaints was to move out managers who seemed too enthusiastic about dictating industry's actions and move in people who agreed it was time to end the obsession with the TMI accident, at that point some ten years in the past.

In 1988, NRC management published new guidance for its research program, entitled *Integration Plan for Severe Accident Closure*

(SECY-88-174).[27] This plan introduced the novel concept that the goal of research on severe accidents should be to show that they were not a problem and that further research was unneeded. This "closure" philosophy conflicted sharply with the principles of scientific objectivity, but it nonetheless became the organizing principle for the NRC's approach to severe accidents from that point onward.

The CPI program collided head on with this new management philosophy. The CPI program's recommended Mark I improvements were duly endorsed by the new management, though in a form somewhat diluted from earlier conceptions. But the rest of the CPI program was summarily terminated within a year of the release of SECY-88-174. The CPI program's final report devoted two sentences to the ice condenser plants, suggesting that any problems that existed should be assessed and dealt with by the plant owners.[28]

At this point in the story of the NRC's response to TMI's rude awakening, a familiar pattern is emerging. The initial, vigorous program of research and reassessment, set against a background of resistance from vested interests, was running out of steam. A new attitude was emerging that reflected a new agenda for the agency, the need to soften the impact severe-accident research was having on the operations and image of the nuclear industry. Ten years had passed since the TMI accident. What the public knew was that there had been a lot of work on reactor safety, and electric utilities were cancelling, not ordering, nuclear power plants. A strident but small sector of the population remained energized about reactor safety, but the proponents of vigilance in the NRC did not have the kind of broad-based citizen support that other regulatory agencies, like the Environmental Protection Agency, enjoyed.

The growing obscurity of reactor safety allowed key people in the NRC to move forward with a plan to implement the new agenda regarding severe accidents. To them, the epitaphs were written, and it was time to put the corpses in the caskets. They were aided in this enterprise by the complexity and subtlety of severe-accident issues. But hidden agendas can be clearly revealed when specialists within the system choose not to cooperate. Then, if the energy behind the agenda is strong enough and an individual is stubborn enough, the conflict between the system and the individual generates enough heat and light to reveal what is really going on.

Just such illumination was provided when a prominent, respected scientist working under contract to the NRC refused to go along with a rush to judgment about a thorny severe-accident issue called "direct containment heating" (DCH), an accident scenario believed to be a potential problem for all containment types. It became one of the most visible and contentious severe-accident issues the NRC ever dealt with and persisted as a source of rancor well into the 1990s. The next section explores this debate and focuses on the story of how one conscientious researcher tried to maintain scientific integrity, even in the face of potent political pressure, and lost.

Shoot the Messenger

DCH eventually provoked one of the most intense and sustained conflicts between the NRC and the nuclear industry about containment vulnerabilities, but awareness of the scenario emerged slowly after the TMI event. One of the aspects of the TMI accident that caught engineers by surprise was that so much of the core melted even though the pressure in the reactor remained very high. This occurred because the steam leak into the containment was not due to a large pipe break, as envisioned in the design basis accident, but rather to a stuck-open valve, an accident in the "small pipe break" category, which, by design basis philosophy, should be less threatening to the containment than a large-break accident. But in reality it proved to be a greater threat.

Analysts studying the TMI event began to indulge in "what if?" exercises about this high-pressure situation. What if the corium at the bottom of the vessel had melted through? Wouldn't the melt then be sprayed with great force into the containment? Was that a good thing or a bad thing compared with what was expected for a low-pressure accident?

The irony that the movie *The China Syndrome* had been released just before the TMI accident was a source of great consternation on the part of nuclear power advocates. The title derived from a hypothetical type of accident in which a large quantity of molten fuel penetrates the steel reactor vessel and pours onto the concrete floor of the containment building. The relentless decay heat in the fuel would keep it so hot that the concrete would be slowly consumed as gravity drew the mass of molten corium toward the center of the earth.

The terminology "China syndrome" was a bit of gallows humor that caught on, but no one really envisioned the corium's penetrating deeply into the earth. The real danger in this scenario comes from the cauldron-like interactions between the corium and the concrete, producing prodigious quantities of gas and heat. The resulting high pressure has the potential to rupture the containment building, releasing massive quantities of radioactivity into the outside atmosphere.

Early speculation about TMI-like accidents suggested that there might be a silver lining for this kind of event, compared with the classic China syndrome variety. Instead of having a self-heating, red-hot pool of corium melting its way through concrete, you might have wide dispersal of the melt, driven hard by the pressurized steam coming out of the reactor vessel. The result might be that the corium would be more widely distributed throughout the building, and hence it might cool down more and be unable to melt concrete.

To explore these and other issues, the NRC sponsored a series of experiments at Sandia National Laboratories in Albuquerque, New Mexico. It was of course not possible to use real nuclear fuel (which is extremely dangerous stuff), but the Sandia scientists developed methods to melt mixtures of materials with similar properties. They also learned how to use pressurized gas to drive the molten mixture out of a hole in a simulated reactor vessel to see how thoroughly the fuel simulant would be dispersed. Initial tests, performed outdoors in the desert south of Albuquerque, demonstrated that the melt would indeed be efficiently swept out of the reactor cavity. Some of these tests were conducted at night, and the molten fuel simulant was ejected in a brilliant arc hundreds of feet into the air, like a giant Roman candle.

When the researchers went on to perform the same experiment inside a closed building, a surprise awaited them. The steel building was violently lifted off its concrete pad, shearing the stout bolts fastening it down. Thus was born the concern about DCH as a means of overpressurizing the containment building. Over the next fifteen years, DCH would become one of the most contentious issues in reactor safety.

What had happened in the indoor test was that the driving gas had broken up the melt into small droplets, not unlike the spray of droplets from a garden sprayer. These droplets were so small and so hot that they burned in the presence of the air in the building. The result was a rapid heating and pressurization, similar to accidents in grain elevators when finely dispersed dust ignites. A more vivid picture might be that of setting off a Roman candle inside the trunk of a car.

These first experiments were performed in the early 1980s, and as additional tests and theoretical analyses were performed, DCH became increasingly important in assessments of severe accidents in U.S. reactors. The NUREG-1150 study found that DCH was one of the largest contributors to public risk for several of the containment types, because if DCH were to cause a breach in the containment, the release of radioisotopes to the environment would be very large, since so much of the radioactive material would be suspended in the containment atmosphere at that time. As for most things, timing is everything for severe accidents, and the DCH scenario has poor timing: virtually simultaneous reactor vessel failure and containment failure. Many other accident scenarios involve a long delay between the two events, so the airborne radioactive material can slowly settle out onto containment surfaces; consequently, much less radioactivity would be released if the containment were finally to fail.

Throughout the 1980s the NRC invested heavily in research on DCH, and the nuclear industry responded with its own research and analysis (generally attempting to minimize the threat). But the issue was exceptionally resistant to a confident resolution.

The NUREG-1150 risk study was based on what was essentially a snapshot of scientific understanding of severe accidents as of about 1985. The study showed that some containment types were much more resistant to severe accidents than others and that, in particular, the large dry containments held up very well to most challenges. Only DCH appeared to be a potent threat to such containments. Since containments of this type represented 57 percent of the entire U.S. reactor fleet,[29] there was a strong incentive to determine how significant the threat from DCH was. So the NRC continued to invest research funds in more experiments, more modeling, and more probabilistic analysis.

This unprecedented, sustained investment in a single severe-accident scenario began to pay off in the late 1980s. The experiments and codes were showing that a number of mitigative effects would occur in a DCH event at the same time as the melt ejection, chemical reaction, and atmosphere heat up. The trend of the results was to reduce the peak pressure calculated in accident simulations, compared to earlier studies. But the analysts still found important scenarios that resulted in pressures above what large dry containment buildings could handle.

The slow pace of scientific progress on this issue was a source of frustration to NRC management, who were taking a no-nonsense approach to severe accident research, exemplified by the 1988 *Integration Plan for Severe Accident Closure* mentioned earlier. But even more potent was the message from Congress expressed in the form of reduced budgets for the NRC.

The nuclear industry was facing dire times in the late 1980s, and many of its leaders believed that the current difficulties were due in large part to image problems: if the public understood nuclear energy better, there would be more support for it. Then, they believed, the other economic problems the industry was facing would be more amenable to solution. For these nuclear cheerleaders, spending time discussing severe accidents was highly counterproductive to improving nuclear's public image and had no corresponding benefit. Many members of Congress were sympathetic to these arguments and saw the NRC as the quintessential self-focused government bureaucracy, standing in the way of the country's economic prosperity and greatness. Such attitudes even began to show up in the commission itself, whose members, unlike the NRC staff, were presidentially appointed and Senate-approved officials, attuned to the political atmosphere favoring deregulation.

The NRC's predicament worsened with the transition of its funding base from taxpayer funds to industry fees in 1991. Now, when nuclear industry lobbyists complained to Congress about overregulation, they were heard even more sympathetically as the ones paying the bills. Congress responded with steady decreases in NRC funding (which, despite being off-budget, was controlled by Congress) and in many cases provided guidance on how the budget was to be allocated among the NRC's functions. Reflecting these forces, the overall research budget for the NRC declined sharply after its peak shortly after the TMI accident, as figure 3 shows.

The NRC's continued frustration with slow progress on severe accident "closure" led to more assertive measures to control the outcome of research. In the 1990s, the NRC's research came more and more to consist of developing information that would justify the foregone con-

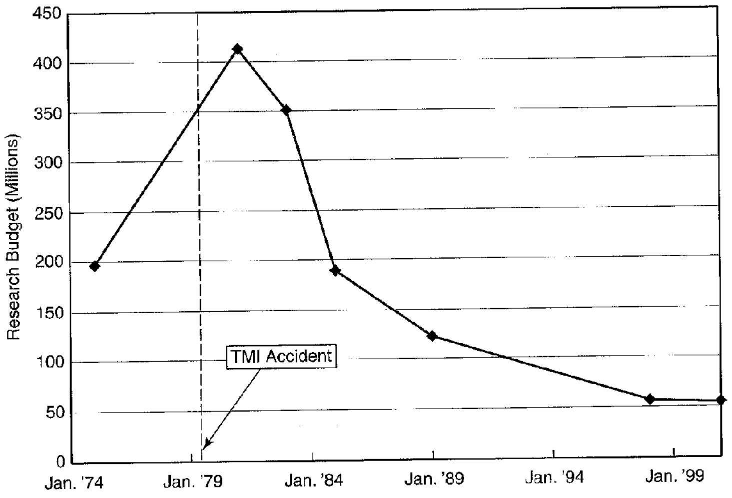

Figure 3
NRC reactor safety research funding before and after TMI (in year-2000 dollars). *Source:* Data for years prior to 1985 from letter from Carlyle Michelson, Chairman of ACRS to Kenneth Carr, chairman of NRC dated June 13, 1991 (available at NRC's Public Document Room). Data for later years from U.S. Nuclear Regulatory Commission, *Budget Estimates Fiscal Year ••*, NUREG-1100, volumes 1–17, published annually (•• = year) since 1985.

clusion that the top reactor "issues" were not a problem for any plant. This concept of closure for severe-accident issues was articulated in a 1992 report, the *Revised Severe Accident Research Plan.*[30] The following statement from a briefing to the commission by the NRC's director of research illuminates the new perspective: "[The *Revised Severe Accident Research Plan*] describes the progress and our understanding of important severe accident phenomena. It defines the research that would lead to closure of core melt progression, direct containment heating, which as I said before was an early containment failure mode, and fuel coolant interactions."[31] The commissioners warmly received the report's message that the continued research program was effectively a mopping-up campaign.

In the ensuing years, the conflict between the philosophy that the goal of severe accident research was to justify its own termination, on one hand, and the actual state of knowledge in the various subject areas, on the other, became a source of great frustration to the participants in the research program. The NRC's duty to maintain vigilance concerning reactor safety had come into direct conflict with the interests of the flagging U.S. nuclear industry.

To accelerate progress on closing the DCH issue, the *Revised Severe Accident Research Plan* laid out a more aggressive and more coordinated effort on DCH research. The NRC chose Sandia National Laboratories to lead a team of national laboratories and other contractors in this project, which was named the DCH Issue Resolution project.

Sandia's management was gratified by the confidence that the NRC displayed in choosing it as the lead laboratory for this program. It was a logical choice. Sandia's experimental facilities and expertise in the area were unparalleled; they were the developers of the CONTAIN computer code, which was the NRC's leading analysis tool for DCH; and finally, Sandia had been the NRC's principal contractor for its NUREG-1150 project.

Sandia is primarily a nuclear weapons laboratory for the Department of Energy, and as the cold war was drawing to an end, budgets for such laboratories were tight and uncertain. Sandia's director of nuclear energy technology, Nestor Ortiz, had been trying to get his managers and their staffs to apply customer-focused management concepts such as total quality management. Being awarded the contract to lead this important new NRC program was a vindication of his efforts. At the same time, he recognized that leading the DCH Issue Resolution Project would be a challenging assignment, since the project would be conducted in a fishbowl environment. The customer was already frustrated at the slow pace of progress and would no doubt pay close attention to see whether the new funding and streamlined organization involved in the project was paying off.

The DCH Issue Resolution Project was to proceed in phases. The initial phase would focus on only one reactor, the Zion nuclear power plant, located near Chicago. Later phases would systematically apply the same methods to the remaining plants. The Zion plant had a large dry containment that was quite strong. It was considered a good representative of the most robust class of containments in the country and had been thoroughly studied in the NUREG-1150 project.

Mathematical methods for accident analysis had reached a point by this time that the NRC could establish a quantitative criterion for resolving the DCH issue at Zion. It wanted Sandia to show that if the reactor vessel failed at high pressure, there would be only a 10 percent or less chance that the containment would leak. It was a way of expressing the philosophy of defense in depth in quantitative terms.

Most experts thought that Zion would have a containment failure probability much less than 10 percent, but the challenge for the project was to develop a method for doing the calculation that was applicable to all reactors and that found favor with a panel of external experts selected by the NRC. Once the method was developed and accepted, the project could proceed with the remaining U.S. plants until all had been studied. It was hoped that all plants could be covered by 1994.

The project ran into trouble in early 1993, when the first draft of the Zion report was submitted for internal peer review at Sandia. Many reviewers took exception to what they considered a lack of rigor in the analysis. Numerous assumptions were challenged as unjustified, and the analysis was criticized for ignoring important phenomena or treating them incorrectly. The treatment of uncertainty was a particular source of criticism: some felt that uncertainties were being grossly underrepresented.

In scientific work, it is normal for peer reviews to generate questions that technical authors resolve with either revisions to the initial report or explanations that satisfy the reviewers. In this case, however, the report authors had a good sense of what the NRC contract managers expected, and they found it difficult to find compromises that satisfied their peers at Sandia. Tensions within the program grew.

The internal arguments over Sandia's peer review soon spilled over to the authors. One of the report authors was from Sandia, and two were subcontractors to Sandia from the University of California at Santa Barbara (UCSB). In the course of trying to resolve the reviewer comments, bitter disagreements between the two organizations developed. This rift stymied progress. Although the UCSB authors technically reported to Sandia, the lead professor maintained separate lines of communication with the NRC's research director, the man who had already gone on record to the commission about the expected outcome of the DCH program.

Internal review comments were still not resolved in June 1993 when the NRC directed that the external committee of experts it had selected review the report.[32] When the comments of those reviewers came in, many of the criticisms they expressed paralleled those of the Sandia reviewers. They questioned such things as how the authors decided what the conditions were in the reactor vessel at the time it failed, that is, the quantity of melt at the bottom of the vessel, its temperature, and its metallic content. Some, like the Sandia reviewers, questioned the use of simplistic models for pressurization of the containment, when more detailed treatments were available. And some of them felt that the treatment of uncertainties was just too superficial, again consistent with the earlier internal review.

As this crisis worsened, Ortiz got personally involved and attempted to apply total quality management methods to this contentious situation. He appointed a total quality committee of his managers to facilitate technical resolution of the disagreements between the authors and the reviewers. The committee's work served to clarify many of the concerns and to organize the disagreements into clear categories but also concluded that the impasse had not been overcome.[33] Ortiz reported this unfortunate state of affairs officially to the NRC in January 1994, a full year after the Zion report was first submitted to reviewers.

The NRC was not sympathetic: this was not the kind of progress it had expected when it entrusted the DCH Issue Resolution Project to Sandia. Brian Sheron, the director of NRC's severe accident work, wrote back to Ortiz, "[T]here are instances where it appears that the authors and the SNL [Sandia National Laboratories] internal review committee are no longer supporting the conclusion in the report, i.e., DCH issue is resolved for the Zion NPP [nuclear power plant]."[34] Later in the same letter, Sheron sharpened the point: "Therefore, when responses to the peer reviewers' comments appear to be unenthusiastic and do not stand behind the conclusion in the report that DCH is essentially resolved for the Zion NPP, it is very unlikely that we can resolve DCH for other PWRs [pressurized water reactors]." Sheron called on Sandia to reevaluate its responses to the peer review comments. Thus the Sandia total quality committee concluded that the authors had yielded too little to reviewers' criticisms, but the NRC customer felt they had yielded too much.

One of the most persistent critics of the Zion report was David C. Williams, the top DCH calculational analyst at Sandia. While the report authors struggled with both internal and external peer review comments, Williams had been looking carefully at some of the experimental test results and found convincing new evidence that the DCH methodology used for the study on Zion was seriously flawed. To his great surprise, though, no one in the program was interested in his findings.

Williams was an internationally respected expert in computer modeling of accident conditions in containment buildings. He had spent much of the previous ten years studying DCH and was generally recognized as the most knowledgeable analyst in the world on the subject. Over the years, he had guided the development of DCH models in the NRC's massive CONTAIN computer code. He followed the experiments on DCH closely and used insights gained from the tests to improve CONTAIN's treatment of physical reality.

In the spring of 1994, Williams was frustrated by the fact that no one was paying attention to the important new results he had obtained. Worse, he found he was excluded from the meetings Sandia had set up with the external review committee to work out the Zion report problems. Unable to convince his management that he belonged in these meetings, he offered to prepare material for others to present, but this offer was rejected. He complained about this in a memorandum that summarized the new information and cautioned, "An outside observer could conceivably interpret my exclusion from the Peer Review Meetings as implying that there is a desire by Sandia management to conceal the findings summarized above from the Peer Review group."[35] The meetings proceeded without Williams's involvement. He acquiesced and turned to other assignments, little knowing that the impediments to his continued professional work would become greater.

The grueling business of resolving the myriad technical disagreements over the Zion report continued over the summer, and Sandia finally found a path to bring the vexing Zion work to an end. In December 1994, two reports were published simultaneously. The first was the original report with only minor changes, authored by the UCSB and Sandia authors. The second did not include the UCSB authors; it was called a "supplement" to the first and contained commentary and (unresolved) criticism about the main report, without contradicting the conclusion that the DCH issue was resolved for Zion. This was an odd compromise, but it got Sandia past the roadblock. It was hardly a satisfactory outcome, since most of the substantive criticisms that Williams and other internal reviewers had raised were ignored.

To his great disappointment, Williams was not allowed to be involved in the follow-on work extending the analysis to other nuclear plants. He was a very capable scientist, however, and there was plenty of other work in Sandia's Nuclear Energy Technology Center to keep him busy.

He at least had the satisfaction in 1994 of completing a large report culminating years of improvement to the CONTAIN DCH models, showing in exhaustive detail how well (or not so well) these models agreed with the dozens and dozens of DCH experiments that had been performed over the previous ten years. This extremely technical treatise was for Williams a magnum opus, highlighting many years of his professional career. Normally, such technical reports are routinely submitted to the NRC for approval before they are published and distributed to technical organizations and libraries throughout the world. This approval step was not for ensuring the technical quality of the work—that was Sandia's job. Rather, the approval step was intended to prevent such things as premature release of safety policy information. In the case of Williams's CONTAIN assessment report, there should have been no problems of that nature: it was strictly a discussion of the science of accident modeling. This "hands off" policy regarding the technical content of contractor reports was conveyed in the formal procedures of the NRC's Office of Research (RES): "While the RES has responsibility for setting the scope, conduct or methodology, and objectives of research, there should be no interference with the contractor presenting his judgement as to the nature and interpretation of the research results."[36] By this time, however, the entire DCH issue had become so politicized that the NRC refused to give Sandia permission to publish the report,[37] even after Ortiz had offered to publish it as a Sandia report, with no NRC affiliation and at no cost to the government. Thus, years of Williams's work came to nothing from the all-important perspective of scientific publication. The NRC's action was clearly inconsistent with its own "hands-off" policy, but at that point Sandia management essentially gave up on Williams's cause. It had done what it could do.

Williams's sense of justice was badly frayed by this time, and he tried to find relief through Sandia's ethics organization. He was concerned with the damage to his career, of course, but he also had broader concerns. He wrote:

> For over 12 years, DCH has been recognized as a major nuclear safety issue. The DCH issue resolution project has offered sweeping conclusions that DCH presents little or no threat in all plants studied to date, which include the most common type of plant in the U.S. It would be surprising if general acceptance of this claim did not result in reduction of efforts to understand DCH better and/or reduction in precautions taken to minimize DCH threats. If, as I believe, there is substantially greater uncertainty concerning the conclusions of the DCH resolution work than is being acknowledged, general acceptance of those conclusions would result in overconfidence concerning DCH that could eventually degrade safety.[38]

Since Sandia's Ethics Office was established to deal with issues of a legal nature and had no real mechanisms to assist in scientific controversies such as this, it directed Williams back to his line organization. Ortiz again tried to implement additional total quality management processes to ameliorate the situation, but Williams was to get no satisfaction from them. In reality, a federal contractor like Sandia has few means other than persuasion to bring pressure on a customer in such a situation. In 1996, Williams had an opportunity to retire early from Sandia with an incentive package (part of a downsizing), and he decided to leave.

David Williams was not an antinuclear troublemaker. He was a brilliant scientist, holding a Ph.D. in nuclear chemistry from the Massachusetts Institute of Technology. He had started his career in nuclear energy as an enthusiastic promoter of nuclear power. Early on, he spent a great deal of his private time in promotional activities, such as speaking in favor of nuclear power at public meetings. He was a scientist who had achieved the highest level of technical achievement at Sandia, the prestigious rank of Distinguished Member of the Technical Staff. He had an international reputation as one of the most capable reactor safety analysts in the world. Yet he left Sandia feeling angry and defeated.

This has been the story of one severe-accident issue. It is not, unfortunately, an isolated case. There is no doubt that the NRC conveyed to Sandia and its other contractors the conclusions they expected from the contracted research. Facing the kind of pressure created by the NRC on the DCH Issue Resolution project, it is not easy for a contractor such as a national laboratory to maintain an open mind. The long tradition of technical integrity in the laboratories' research is certainly a bulwark against blatant distortion of research results. But management must also consider its commitment to the customer paying the bill. Ortiz had to think not only about Williams's concerns, but also about the livelihoods of the hundreds of other scientists and engineers in his center. Total quality management processes are not well suited for resolving such conflicts. (Williams, for his part, believes Sandia was at least as responsible for the shortcomings of the DCH Issue Resolution project as the NRC was.)[39]

The contrast is sharp between the forthright research atmosphere surrounding the NUREG-1150 project and the way things were handled with the DCH Issue Resolution Project ten years later. The earlier project recognized the importance of factoring in the opinions of many different experts, who might not see every issue the same way; at that time Williams was an important and respected player on the expert panels the NRC convened to wrestle with the uncertainties of hypothetical severe accidents. Ten years later, his expertise was of no value to the NRC, because he refused to conform his technical opinions to the positions adopted by NRC management.

Ice Condensers, the Little Containments that Don't

As noted above, the pattern of intense vigilance triggered by an external event, followed by a slide into complacency, seems to be characteristic of the federal government's attention to the hazards of nuclear technology. It appeared in the history of U.S. nonproliferation policy presented in chapter 2, and it can be seen in the context of the NRC's attitudes toward severe reactor accidents in the decades following TMI, as just described in the story of David Williams and the DCH Issue Resolution project. It may very well be true that the benefits of nuclear technology are worth the risks, but the risks are persistent and long-term and cannot be properly managed with such uneven attention by the agencies charged with protecting the public's interests.

Similarly, today's public indifference about nuclear war and proliferation allows the federal government's new tritium policy to move forward despite its grave shortcomings. But there is also a more direct connection between the DCH Issue Resolution work and the new tritium policy. It concerns the safety of the ice condenser plants that TVA will modify to be able to produce tritium.

After the Zion report was finally published, NRC and Sandia proceeded to evaluate the DCH threat for one group of reactors after another, eventually declaring the issue "resolved" for all large dry containment plants without a significant hitch. Then, around 1996, they decided to take on one of the problem containments, the ice condenser.

The last time ice condensers had been in the NRC's spotlight was in the late 1980s, when the CPI program had been poised to assess systematically the performance of ice condensers in severe accidents. Then the program was abruptly terminated, as discussed earlier in this chapter. It seemed then that ice condenser containments had dodged a bullet, avoiding the intense scrutiny that led NRC to call for containment changes at the more numerous Mark I plants. The DCH Issue Resolution program put the ice condenser containments, with their remarkable weaknesses, right back in the spotlight.

For the ice condenser work, Sandia was the only contractor involved, since NRC's research budgets had been squeezed dramatically in the years following the infamous Zion DCH report. But the starting point for Sandia's specialists was not auspicious. Given the NUREG-1150 results, it was not apparent to them how ice condensers could be found resistant to DCH.

In 1997 Sandia submitted to the NRC its draft report on DCH Issue Resolution for ice condensers. It concluded that, unlike all the containment types studied previously, ice condensers could not meet the "success criterion"—containment failure probability less than 10 percent—for resolution of DCH. The report also pointed out that DCH was not the only problem ice condensers had. In many sequences, if DCH were hypothesized by fiat not to cause the containment to fail, some other challenge would.

Characteristically, the NRC project managers did not publish this report and instead contemplated ways to salvage the commitment the NRC's management had made to the commission to close the DCH issue. The broad vulnerability of ice condensers to a variety of severe accident challenges, not just DCH, led the NRC to commission Sandia to do a systematic evaluation of ice condenser response to the entire threat spectrum. This is exactly what the CPI program was going to do nine years earlier, but there was a difference. In those days, the goal was to decide if changes in the containments were called for. Now, the purpose of the work was to demonstrate there was no longer a need for research on DCH.

Sandia carried out the integrated study of ice condenser vulnerabilities in 1998 and 1999. The report it submitted to the NRC showed quantitatively what most analysts expected: that these containments provide far less defense in depth than conventional containments.

The problem with these reactors is that the volume of the containment building is small compared to large dry containments, but the reactor core is about the same size. The ice suspended in the massive banks of wire baskets (see figure 4) would certainly be effective in absorbing the steam from a pipe break accident, but severe accidents tend to overwhelm the ice. Buildup of hydrogen is a particular problem, since removing steam from a steam-hydrogen-air mixture increases the concentration of hydrogen, making the gas more combustible, even detonable.

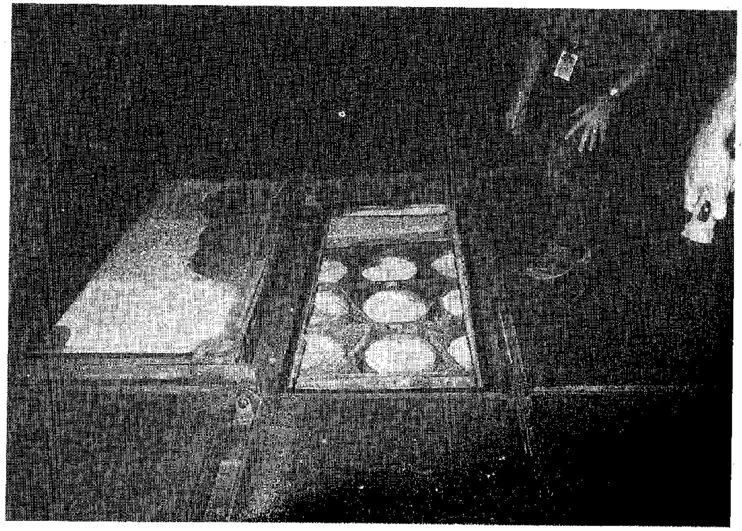

Figure 4
View of ice baskets from top deck of ice chest. Each wire basket is a forty-eight-foot-high cylinder. The entire ice condenser containment holds more than 1,000 tons of ice.

One of the safety features in these containments is an array of strategically located hydrogen igniters, essentially spark plugs, that would burn the hydrogen off before the concentration got too high. But there are some conditions in which the igniters do not work, such as in the "station blackout" scenario, in which all electrical power is lost.

Some of the difficulties of the Mark I containments also derived from their small volume, but the hydrogen combustion danger had been eliminated by the use of an inerted (that is to say, nitrogen-filled) containment building. Without oxygen in the atmosphere, the hydrogen cannot burn. In contrast, the atmosphere in the ice condenser containments is ordinary air, and hydrogen combustion is a major severe-accident problem in such containments even without DCH. The Sandia analysts found that DCH and hydrogen combustion could occur simultaneously and cause containment failure in ice condenser containments for accidents that large dry containments could easily survive.

The revised Sandia report concluded, as had its original report, that the ice condensers could not meet the 10-percent containment failure criterion for DCH. But because an integrated analysis had been done, the NRC had another way to declare victory. Since there were so many other ways for the containment to be compromised, it turned out that DCH was a minor contributor, on a probabilistic basis, to the overall probability of containment failure. This was partly because most severe accidents were found not to result in high-pressure ejection of corium into the containment. For those accidents, the containment's resistance to DCH was a moot point. In other cases, the containment failed because of hydrogen combustion before melt ejection occurred, again making the containment's resistance to DCH irrelevant.

Since the focus of the project was DCH, not overall protection provided by the containment, the NRC chose to view the Sandia results in a positive light, and in June 2000 Ashok Thadani, the NRC's director of research, announced that the DCH issue had been resolved for ice condensers.[40] His logic? Since it was not possible to declare victory on the basis of the original success criterion, the NRC used a new criterion, that despite being vulnerable to DCH, these containments were equally vulnerable to a wide variety of more likely events. The logic was equivalent to saying, "why worry about automobile gas tanks exploding due to collisions from the rear if most collisions are from the side or the front?"

From the narrow perspective of the DCH Issue Resolution project, the NRC decision may have been logical. But from the perspective of how well the public is protected from reactor accidents, the message of the Sandia study was bleak. The calculations showed that ice condensers are very susceptible to failure under a wide variety of severe accident conditions, not just DCH. If, for example, electric power to the igniter system were lost, hydrogen could build up to the point that either random sparks or restoration of power to the igniters would trigger an overwhelming hydrogen burn that would cause the containment to fail. The containment could fail if melt ejected from the reactor vessel built up on the containment shell and melted through. The containment could fail if there were a large "spike" of steam created by rapid mixing of molten debris with a pool of water beneath the vessel.

Station blackouts posed a particular problem. Because many of the containment's safety features require AC power (rather than battery, or DC, power that other plants use), ice condenser containments are more or less "sitting ducks" in such an accident. Unless power is restored, failure of the reactor vessel, release of corium into the containment, and containment failure are all virtually inevitable. The Sandia study found, for example, that in the event of a station blackout, the Sequoyah's probability of containment failure was about 97 percent.[41]

Thadani's memorandum announcing the success of the DCH Issue Resolution effort recognized this darker picture:

> The recent ice condenser study . . . concluded that the ice condenser plants are more vulnerable to early containment failure than large dry containments, but that this vulnerability is not due to DCH. In fact, early containment failure in ice condensers was dominated by non-DCH hydrogen combustion events rather than by DCH, and was seen largely to depend on plant specific probabilities for station blackout (ice condenser igniter systems are not operable during station blackout events).

Despite this recognition the NRC, still gun-shy about the politics of severe accidents, has not vigorously pursued the issue of ice condenser vulnerability.[42] As described at the beginning of this chapter, in the early days of nuclear power the AEC required effective containment systems as a trade-off to balance the increase in public risk due to locating nuclear power plants near population centers. But for many of the most important accidents, ice condenser plants pose essentially the same risk to the public as the same plants would with no containment at all.

What has all this to do with the DOE's selection of the Sequoyah and Watts Bar plants to be the first commercial power plants to produce nuclear weapons materials? As far as DOE has been concerned, nothing. The safety of the plants was irrelevant to the DOE selection. What mattered was the fact that, of all the utilities operating nuclear power plants in the United States, only the TVA could be persuaded to cooperate with the nuclear weapons program by producing tritium in its reactors.

Chapter 5 will explain how this occurred. But before that discussion, it is time to switch attention from the danger of nuclear reactor accidents to the other downside of the government's decision to produce tritium in commercial reactors: the danger that it will encourage proliferation of nuclear weapons throughout the world.

regardless of isotope. By "physics," I mean the physics of cross-sections, neutron flux, nuclear fission, and chain reactions, not details of the neutron spectrum or the hydrodynamics of inertial confinement.

2. TVA's power program is self-financed, but its much smaller nonpower programs, such as regional economic development, are federally financed.

Chapter 2

1. Richard Rhodes, *The Making of the Atomic Bomb* (New York: Touchstone, 1986), p. 745.

2. Ibid, pp. 734, 740.

3. Paul Kennedy, *The Rise and Fall of the Great Powers* (New York: Vintage Books, 1987), p. 343.

4. Richard G. Hewlett and Jack M. Holl, *Atoms for Peace and War* (Berkeley and Los Angeles: University of California Press, 1989), p. 72.

5. David Lilienthal, *Change, Hope and the Bomb* (Princeton: Princeton University Press, 1963), p. 103.

6. The liability limits specified by the Price-Anderson Act have changed since its implementation in 1957. Today, each reactor must be covered by a $200 million liability insurance policy for an accident involving that reactor, but if an accident at *any* reactor exceeds that amount, *all* reactor owners must contribute to the excess damages, up to a maximum of $88 million per reactor owned. The government assumes responsibility for any costs incurred above that point. Since there are currently 104 operating reactors, the government would to have to step in only if the costs of a reactor accident exceeded $9.4 billion ($200 million plus $88 million times 104 reactors). Web site of the American Nuclear Insurers (http://www.amnucins.com).

7. Hewlett and Holl, *Atoms for Peace and War*, p. 516.

8. Elizabeth S. Rolph, *Nuclear Power and the Public Safety: A Study in Regulation* (Lexington, MA: Lexington Books, 1979), p. 56.

9. Ibid., p. 93.

10. Michael J. Brenner, *Nuclear Power and Non-proliferation* (New York: Cambridge University Press, 1981), p. 19.

11. Ibid., p. 30.

12. Brenner (ibid.) makes this point forcefully: "*The crisis that flared in June 1974 with the blocking of enrichment orders was in fact planned*" (p. 37, italics in original). He then provides ample evidence to back up the statement.

13. Ibid., p. 36.

14. Ibid., p. 52.

15. Ibid., p. 57.

16. Ibid., p. 61.

17. Ibid., p. 14.

18. Ibid., p. 70.

19. Energy Secretary Hazel O'Leary announced this fact at a press conference on June 27, 1994, according to Institute of Energy and Environmental Research, *Physical, Nuclear and Chemical Properties of Plutonium*, available on the institute's Web page at <http://www.ieer.org/fctsheet/pu-props.htm>.

20. President Gerald Ford, *Statement on Nuclear Policy*, October 28, 1976, quoted in Brenner, *Nuclear Power and Non-proliferation*, p. 115.

21. Ford Foundation Nuclear Energy Policy Study Group and Mitre Corporation, *Nuclear Power Issues and Choices* (Cambridge, MA: Ballinger, 1977).

22. Harald Müller, David Fischer, and Wolfgang Kötter, *Nuclear Nonproliferation and Global Order* (New York: Oxford University Press, 1994), p. 52.

23. Four countries today continue efforts toward plutonium-based reactor operations: France, Great Britain, Russia, and Japan. These programs have had virtually no impact, however, on the commercial electricity sector.

Chapter 3

1. This point is made clear in Rolph's discussion of standards: "And even by 1966 only six commercial plants were in operation, all 265 MW or less. . . . In spite of the obvious obstacles, for political reasons, the AEC had little option but to try, all the while arguing the standards should be general and flexible" (*Nuclear Power and the Public Safety*, p. 65).

2. Ibid., p. 61.

3. George T. Masuzin and J. Samuel Walker, *Controlling the Atom* (Berkeley and Los Angeles: University of California Press, 1984), pp. 219 ff.

4. Rolph, *Nuclear Power and the Public Safety*, table B-1.

5. Mazuzan and Walker, *Controlling the Atom*, p. 198.

6. "Nonetheless, the AEC devised no specific standard incorporating the relationship between distance and containment, and it continued to judge the site for each power reactor application on a case-by-case basis" (ibid., p. 219).

7. Rolph, *Nuclear Power and the Public Safety*, p. 68.

8. Donald W. Stever, Jr., *Seabrook and the Nuclear Regulatory Commission*, (Hanover, NH: University Press of New England, 1980), p. 45.

9. This discussion focuses primarily on the design criteria for containment strength. Actually, at least one other design criterion implicitly assumed core melting. The criterion for containment leak-tightness started by postulating a certain amount of radioactivity in the containment and then required the calculated dose to a human at the site boundary to be less than a certain value. The amount of radioactivity postulated for use in this calculation was so high that the only way it could occur would involve core uncovery. It is probable that despite its inconsistency with respect to other design criteria, this approach was chosen because using a design basis accident to determine the amount of radioactivity in the containment would have resulted in allowing leak rates to be so large

that they would have been in intuitive conflict with the idea of a "leak-tight" building.

10. Charles Perrow, *Normal Accidents* (Princeton, N.J.: Princeton University Press, 1999), p. 30 (first published by Basic Books, New York, 1984).

11. "Three Mile Island: A Chronology," *Washington Post*, March 28, 1989, available at <http://www.washingtonpost.com/wp-srv/national/longterm/tmi/stories/chrono032889.htm>.

12. Nuclear Regulatory Commission, *Issue 155: Generic Concerns Arising from TMI-2 Cleanup*, NUREG-0933, Rev. 2 (Washington, DC, 1992).

13. Perrow, *Normal Accidents.*

14. Memorandum adopted by the participants of "Lessons of Chernobyl" conference, April 22, 1996, Kiev, Ukraine, available at the Web site of the Nuclear Information and Resource Service, Washington, DC (http://www.nirs.org/confstat.htm).

15. United Nations Scientific Committee on the Effects of Atomic Radiation, *Sources and Effects of Ionizing Radiation: UNSCEAR 2000 Report to the General Assembly, with Scientific Annexes, Volume 2: Effects* (New York: United Nations, 2000).

16. International Nuclear Safety Advisory Group (INSAG), *The Chernobyl Accident: Updating of INSAG-1*, (INSAG-7) (Vienna International Atomic Energy Agency, 1992), p. 24.

17. In 1955, for example, a government test reactor called EBR-I underwent a partial core melt accident. Mazuzan and Walker, *Controlling the Atom*, p. 127.

18. Atomic Energy Commission, *Theoretical Possibilities and Consequences of Major Accidents in Large Nuclear Power Plants*, AEC/DOE WASH-740 (Washington, DC, March 1957).

19. Glenn Seaborg, quoted in Rolph, *Nuclear Power and the Public Safety*, p. 49.

20. Nuclear Regulatory Commission, *Reactor Safety Study—An Assessment of Accident Risk in U.S. Commercial Power Plants*, WASH-1400, NUREG 75/014, 1975.

21. Rolph, *Nuclear Power and the Public Safety*, p. 149.

22. Nuclear Regulatory Commission, *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants*, NUREG-1150 (Washington, DC, 1989).

23. The early conditional containment failure probability (CCFP) measure for Mark Is was lower than that for ice condensers, but the corresponding core damage frequency (CDF) had the opposite ranking. The product of the two is the probability of an early release, and the ice condenser had the largest value of the five plants studied.

24. These counts are based on reactors operating as of 1988 and are taken from Nuclear Regulatory Commission, *Nuclear Information Digest*, NUREG-1350, vol. 9, appendix A. With the start-up of Watts Bar Unit 1 in 1996 the total number of ice condenser plants operating in the United States was brought to nine.

25. Nuclear Regulatory Commission, *Mark I Containment Performance Improvement Program*, SECY-89-017 (Washington, DC, January 23, 1989); also "Installation of Hardened Wetwell Vent," Generic Letter 89-16, Washington, DC, September 1, 1989.

26. Nuclear Regulatory Commission, *Individual Plant Examination Program: Perspectives on Reactor Safety and Plant Performance, Part 1: Final Summary Report*, NUREG-1560 (Washington, DC, 1997), vol. 1, pp. 4–13.

27. Nuclear Regulatory Commission, *Integration Plan for Closure of Severe Accident Issues*, SECY-88-147, May 25, 1988.

28. Nuclear Regulatory Commission, "Completion of the Containment Performance Improvement Program and Forwarding of Insights for Use in the Individual Plant Examination for Severe Accident Vulnerabilities," Generic Letter 88-20, Supplement 3, Washington, DC, July 6, 1990.

29. Based on the NRC's list of licensed plants. Nuclear Regulatory Commission, *NRC Information Digest*, NUREG-1350 (Washington, DC, 1997), vol. 9, appendix A.

30. Nuclear Regulatory Commission, *Revised Severe Accident Research Plan*, NUREG 1365, Rev. 1 (Washington, DC, December 1992).

31. Transcript from "Briefing on Severe Accident Research Program," Tuesday, October 26, 1993 (available at NRC's Public Document Room).

32. This and subsequent elements of the Zion report chronology are from Dave Williams's "Chronology," in his internal memorandum to N. R. Ortiz, July 8, 1994.

33. Internal memorandum from D. A. Powers et al. to N. R. Ortiz, December 17, 1993.

34. Letter from Brian Sheron, U.S. Nuclear Regulatory Commission, to N. R. Ortiz, Sandia National Laboratories, February 4, 1994 (available at NRC's Public Document Room).

35. Internal memorandum from Williams to J. E. Kelly, Sandia National Laboratories, May 16, 1994.

36. Nuclear Regulatory Commission, "Implementing Research Projects," RES Office Letter No. 6, Rev. 1, Washington, DC, November 26, 1993.

37. Letter from C. J. Tinkler, U.S. Nuclear Regulatory Commission, to K. D. Bergeron, Sandia National Laboratories, February 7, 1996 (available at NRC's Public Document Room).

38. Internal memorandum from D. C. Williams to C. M. Tapp, both Sandia National Laboratories, September 19, 1996.

39. In a detailed e-mail dated April 21, 2001, Williams explained to me the reasons for his belief that Sandia was at least as responsible for the problems as was the NRC's

It was Ortiz's responsibility to insist that technical differences between the lead Sandia author of the study and other staff, including Williams, be resolved internally, and he failed to do so. It was Ortiz who ignored Williams' repeated pleas that the conflict between his results using the CONTAIN code and the simplistic Issue Resolution models be brought to the attention of the NRC and the external peer reviewers at an early stage, when corrective action could have been taken without greatly disrupting the program. It was Ortiz who decided Williams would not be allowed to present his results to the external peer reviewers. It was Ortiz who made the decision that Williams' report on the CONTAIN analysis of the DCH experiments would not be published: Sandia's contract with the NRC explicitly allowed Sandia to proceed with publication at Sandia's expense without the sponsor's endorsement. (Indeed, this contract provision had been included to cover precisely this situation, i.e., to allow Sandia to preserve its integrity by publishing research that might yield results displeasing to the sponsor.)

40. "DCH Issue Resolution for Ice Condenser Plants," memorandum from Ashok Thadani, Director of the Office of Nuclear Regulatory Research, to Samuel J. Collins, Director of the Office of Nuclear Reactor Regulation, June 22, 2000 available at NRC's Public Document Room, ADAMS Accession Number ML003725995.

41. M. M. Pilch, K. D. Bergeron, and J. J. Gregory, *Assessment of the DCH Issue for Plants with Ice Condenser Containments,* U.S. Nuclear Regulatory Commission, NUREG/CR 6427 (Washington, DC, April 2000), p. 109.

42. However, the NRC has not entirely ignored the vulnerabilities of ice condensers as revealed in NUREG/CR 6427. In 2001 NRC established a "Generic Issue," GI-189, regarding the vulnerability of ice condensers and GE's Mark III containments to failure from hydrogen combustion during station blackouts. The costs and benefits of improved igniter reliability for these containment types is a subject of ongoing research at the NRC. NRC, "Staff Plans for Proceeding with the Risk-Informed Alternative to the Standards for Combustible Gas Control Systems in Light-Water-Cooled Power Reactors in 10 CFR 50.44," SECY-01-0162, August 23, 2001.

Chapter 4

1. Harvard Nuclear Study Group, *Living with Nuclear Weapons* (Cambridge: Harvard University Press, 1983), p. 215.

2. Rodney W. Jones and Mark G. McDonough, *Tracking Nuclear Proliferation* (Washington, DC: Carnegie Endowment for International Peace, 1998).

3. United Nations, "Review Conference of Parties to NPT Opens at Headquarters; Much Disarmament Machinery Has 'Started to Rust,' Secretary-General Warns," Press Release DC/2692, United Nations, New York, April 24, 2000.

4. There is a second, overlapping export control group called the Zangger Committee with a somewhat less comprehensive program of control.

5. Federation of American Scientists, "Membership of Zangger Committee and Nuclear Suppliers Group," available at the federation's Web site (http://sun00781.dn.net/nuke/control/nsg/member.html). Note that China, though not in the NSG, is a member of the parallel Zangger Committee but does not agree to international safeguards on its nuclear-related exports. This is believed to be because of its sponsorship of Pakistan's nuclear program, according to Jones and McDonough, *Tracking Nuclear Proliferation*, p. 309.

6. International Atomic Energy Agency, "IAEA Information Circular INFCIRC/254/Rev.3/Part 2," Washington, DC, February 24, 1998. The annex of this circular contains the list of controlled dual-use exports.

7. Mitchell Reiss and Robert S. Litwak, eds., *Nuclear Proliferation after the Cold War* (Baltimore: Johns Hopkins University Press, 1994), p. 344.

8. Gary Samore, in Reiss and Litwak, *Nuclear Proliferation after the Cold War*, p. 16.

9. Müller et al., *Nuclear Non-proliferation and Global Order*, p. 144.

10. Tim Weiner, "U.S. and China Helped Pakistan Build Its Bomb," *New York Times*, June 1, 1998, p. A6.

11. Peter R. Lavoy, Scott D. Sagan, and James J. Wirtz, *Planning the Unthinkable* (Ithaca, NY: Cornell University Press, 2000), pp. 125 (India), 166 (Pakistan). Some of these tests are known to have involved the use of thermonuclear materials, for example, deuterium and tritium, but it is not known what degree of boost was achieved.

12. John Tagliabue, "A Warning from an Official about an Increased Possibility of Nuclear Terror," *New York Times*, November 2, 2001, p. B1.

13. The most powerful hydrogen bombs have two stages: a primary, which is a boosted fission weapon like that illustrated in figure 1, and adjacent to it, a secondary, which uses lithium deuteride and uranium as fuel and achieves a much higher fraction of its energy from fusion reactions than the boosted primary does. In a two-stage device, the primary ignites the secondary.

14. Jones and McDonough, *Tracking Nuclear Proliferation*, p. 29.

15. Matthew Bunn, *The Next Wave: Urgently Needed Steps to Control Warheads and Fissile Material* (Washington, DC: Carnegie Endowment for International Peace and Cambridge: Harvard Project on Managing the Atom, April 2000), p. 10, available at <http://ksgnotes1.harvard.edu/BCSIA/Library.nsf/pubs/Nextwave>.

16. Jones and McDonough, *Tracking Nuclear Proliferation*, p. 5, provides the inventories and critical masses for enriched uranium and separated plutonium that lead to this figure.

17. Jones and McDonough, *Tracking Nuclear Proliferation*, p. 18.

18. Matthew Bunn, "Enabling a Significant Future for Nuclear Power," in *Proceedings of Global '99: Nuclear Technology—Bridging the Millennia*, Jackson Hole, WY (La Grange Park IL: American Nuclear Society, 30 August 1999), available at <http://ksgnotes1.harvard.edu/BCSIA/Library.nsf/pubs/.mb-nuclearfutr.htm>.

19. Bunn, *The Next Wave*, p. 18.