Available online at www.sciencedirect.com
SCIENCE DIRECT
Reliability Engineering and System Safety 89 (2005) 271–285
RELIABILITY ENGINEERING & SYSTEM SAFETY
www.elsevier.com/locate/ress

A historical overview of probabilistic risk assessment development and its use in the nuclear power industry: a tribute to the late Professor Norman Carl Rasmussen

William Keller, Mohammad Modarres*

Center for Technology Risk Studies, 2100 Marie Mount Hall, University of Maryland, College Park, MA 20742 7531, USA

Received 27 July 2004; accepted 31 August 2004
Available online 11 November 2004

Abstract

This paper reviews the historical development of the probabilistic risk assessment (PRA) methods and applications in the nuclear industry. A review of nuclear safety and regulatory developments in the early days of nuclear power in the United States has been presented. It is argued that due to technical difficulties for measuring and characterizing uncertainties and concerns over legal challenges, safety design and regulation of nuclear power plants has primarily relied upon conservative safety assessment methods derived based on a set of design and safety principles. Further, it is noted that the conservatism adopted in safety and design assessments has allowed the use of deterministic performance assessment methods. This approach worked successfully in the early years of nuclear power epoch as the reactor design proved to be safe enough. However, it has been observed that as the conservative approach to design and safety criteria proved arbitrary, and yielded inconsistencies in the degree to which different safety measures in nuclear power plants protect safety and public heath, the urge for a more consistent assessment of safety became apparent in the late 1960s. In the early 1970s, as a result of public and political pressures, then the US Atomic Energy Commission initiated a new look at the safety of the nuclear power plants through a comprehensive study called 'Reactor Safety Study' (WASH-1400, or 'Rasmussen Study'—after its charismatic study leader Professor Norman Rasmussen of MIT) to demonstrate safety of the nuclear power plants. Completed in October 1975, this landmark study introduced a novel probabilistic, systematic and holistic approach to the assessment of safety, which ultimately resulted in a sweeping paradigm shift in safety design and regulation of nuclear power in the United States in the turn of the Century. Technical issues of historic significance and concerns raised by the subsequent reviews of the Rasmussen Study have been discussed. Effect of major events and developments such as the Three Mile Island accident and the Nuclear Regulatory Commission and the Nuclear Industry sponsored studies on the tools, techniques and applications of the PRA that culminated in the present day risk-informed initiatives has been discussed.

Keywords: PRA; History of PRA; Reactor safety study, RSS; Rasmussen study; Probabilistic risk assessment; Risk-informed regulation; Norman C. Rasmussen

"A new scientific truth triumphs not because its opponents become convinced and finally see the light, [but] rather, because they eventually die and a new generation is born which is familiar with the new concepts."—Max Planck [1]

1. Introduction and background

This paper has several purposes. The first is to analyze the history leading to the development, and past and present uses of probabilistic risk assessment (PRA) in the nuclear industry. A second purpose is to highlight the effects of the pioneering PRA study, the so-called Rasmussen Report [2], on the nuclear power industry. The third purpose is to characterize the current use of PRAs and related issues, along with a perspective for its future uses in the nuclear industry. The paper particularly emphasizes the role played

* Corresponding author. Tel.: +1 301 405 5225; fax: +1 301 314 9601.
E-mail address: modarres@umd.edu (M. Modarres).

0951-8320/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ress.2004.08.022

by the late Professor Norman Rasmussen of MIT in guiding, presenting and defending the study methodology and its results during the 1970s and into the early 1980s.

1.1. The origin of nuclear power

Nuclear regulation was the responsibility of the Atomic Energy Commission (AEC), a five-member Commission which Congress first established as part of the Atomic Energy Act of 1946 to maintain strict control over atomic technology and to exploit it further for military applications. The 1946 Act, passed while relations with the Soviet Union were strained with the start of the Cold War, tacitly acknowledged the potential peaceful benefits of atomic power. It highlighted the military aspects of nuclear energy and the need for secrecy. The 1946 law excluded commercial applications of atomic energy and rested the ownership of the nuclear knowledge with the government.

Congress later replaced the 1946 Act with the Atomic Energy Act of 1954, which made the commercial development of nuclear power possible. The 1954 Act ended the government's monopoly on technical data and made the need for commercial nuclear power an urgent national goal to promote the peaceful uses of atomic energy provided a *reasonable assurance* exists that such uses would not result in undue risks to the health and safety of the public. The 1954 Act directed the AEC "to encourage widespread participation in the development and utilization of atomic energy for peaceful purposes". At the same time, it required the AEC to regulate the anticipated nuclear industry to protect public health and safety from radiation hazards. Thus, the 1954 act assigned the AEC three major roles: to continue its weapons program, to promote the private use of atomic energy for peaceful applications, and to protect public health and safety from the hazards of commercial nuclear power [3]. The urgency that led to the 1954 Act and to the commercial nuclear power program was largely came from the fear of falling behind other nations such as Great Britain and, possibly, the Soviet Union as well as from perceptions of the long-range need for new energy sources. Wary of the costs involved and the possible risks of nuclear power, the electric industry did not respond enthusiastically to the 1954 Act [3].

1.2. The origin of reactor safety regulation and probabilistic methods

As with most histories of nuclear power, the initial consideration of safety issues begins with the Manhattan Project during the World War II. The Manhattan Project included several separate disciplines: experimental and theoretical physics, chemical engineering, mechanical engineering, and electrical engineering. Each group brought different methods of design and construction to the project. The chemical engineers of the Du Pont Corporation led the effort to build the nuclear reactors at Hanford, WA. During the construction process, the chemical engineers clashed with the physicists over safety issues, especially with Eugene Wigner, who had led the design effort for the smaller, prototype reactors built in Oak Ridge. Using their background in chemical processes, the Du Pont engineers broke the reactor design into smaller, relatively independent sub-systems, whose design would be frozen early, so any dependent systems could be designed as well [4]. This created the concept of functional and structural independence and later gave rise to the concept of 'defense-in-depth', which promoted layers of independent 'barriers' to prevent release of radioactive substances into the environment. Because the Du Pont engineers lacked a track record with the nuclear technology, they incorporated several safety features to overcome the uncertainties in characterizing the performance and effectiveness of the 'barriers', including redundancy, large safety margins, and systems designed to limit the release of radioactive effluents which would contaminate the environment.

To measure the effectiveness and performance of the safety systems the AEC regulatory engineers for avoiding the need to calculate best estimate uncertainties, proposed using deterministic approaches through conservative assumptions and calculations. They devised the concept of 'design basis accidents' to measure the effectiveness of the 'barriers' and safety systems. Safety was therefore defined as the ability of the nuclear reactor to withstanding a fixed set of prescribed accident scenarios judged by the AEC experts as the most significant adverse events in a nuclear power plant. The premise was that if the plant can handle the design basis accidents, it can handle any other accidents—an attempt to eliminate the possibility of reactor failure from fundamental design flaws and worst possible accidents. As part of defense-in-depth the AEC required multiple back-up equipment and redundancies in safety design. The AEC believed that in general, accidents would be credible if their occurrence might be caused by one single equipment failure or operational error following certain initiating events, with some considerations of the probability of such accidents. However, consideration of 'incredible events' such as the catastrophic failure of the reactor pressure vessel or multiple independent failure events was excluded.

The concept of defense-in-depth originated in 1940s and dominated by the lack of a precise knowledge of design margins evolved into a set of regulatory design and safety principles namely:

1. Use of multiple active and/or passive engineered barriers to rule out any single failures leading to release of radioactive materials.
2. Incorporation of large design margins to overcome any lack of the precise knowledge (epistemic uncertainty) about capacity of barriers and magnitude of challenges imposed by normal or accident conditions.

3. Application of quality assurance in design and manufacture.
4. Operation within predetermined safe design limits.
5. Continuous testing, inspections, and maintenance to preserve original design margins.

Examples of the application of this principle are the calculations performed at Savannah River [5] where the time needed for a reactor shutdown was calculated conservatively through deterministic methods for every power level. The engineers then designed and built the reactor to account the calculated shutdown times to avoid core meltdowns. Another example of the deterministic approach was design of fuel slugs for the reactors. Du Pont conservatively assumed that some fuel slugs would fail, and designed the system accordingly, partly because the company (and everybody else) lacked experience and data in using the fuel slugs [5].

Concern with quantitative measures of risk and reliability of reactors was not a prime factor in the early design process. However, reliability began to appear with the education of nuclear engineers in the 1950s. North Carolina State University had the first nuclear engineering program starting in 1957 followed by MIT shortly thereafter. Ernst Frankel, a professor at MIT, wrote a textbook, *System Reliability and Risk Analysis*, published in the early 1960s, which provided both the mathematical framework and probabilistic methods for assessment of engineering systems. The generation of electrical and nuclear engineers who graduated from MIT in the 1960s studied reliability methods in which Frankel linked to the traditional deterministic approach. Frankel taught engineers how to estimate the failure probabilities of systems given uncertainties with certain operating parameters [6]. Another book by Green and Bourne [7] published in the early 1970s provided a strong theoretical basis for applications of reliability methods in risk assessment of complex engineering systems.

While formal consideration of risk and reliability was not a concern of the AEC, it acknowledged, however, that it could not eliminate all risks through its defense-in-depth principle and design basis accident methods. C. Rogers McCullough, chairman of the Advisory Committee on Reactor Safeguard (ACRS), informed the Joint (Congressional) Committee on Atomic Energy (JCAE) in 1956 that because of technical uncertainties and limited operating experience, 'the determination that the hazard is acceptably low is a matter of competent judgment'. In 1957, WASH-740, the first comprehensive look at the consequences of a large nuclear accident, was published by the AEC. The purpose was to help focus Congressional deliberation of the Price-Anderson Act on the potential harms from reactor accidents. The Price-Anderson Act was being considered to provide supplemental government insurance for private nuclear reactors. WASH-740 originally only looked at the 200 MW class of reactors then in operation and predicted potential damage due to an accident in the $7 billion range [8]. WASH-740 estimated the risk for a serious reactor accident as 10^{-6} per reactor year of operation, a value still within the range of probabilities being estimated today for an occurrence of a large early release of radiation due to reactor accidents. But, the Price-Anderson Act arbitrarily used a $500-million of insurance figure, above the commercial insurance of $60-million provided by private insurance companies. When it was revised in 1964–1965 because of the larger reactors being designed, the worst-case nuclear accident cost rose to $17-billion [9]. The WASH-740 study focused on the dangers of large Loss of Coolant Accidents (LOCAs) as the leading source of worst radiation release into the environment.

As the reactor safety systems continued to grow in size and complexity, a new method of analysis was needed to produce reasonably more accurate risk estimates. At the urging of the ACRS, which first troubled about the so-called China syndrome, the AEC established a special task force to investigate the core melt problem in 1966. The task force report published in 1967 offered assurances about the improbability of a core meltdown and the reliability of emergency core cooling designs, but it also acknowledged that a LOCA could cause a breach of containment if the Emergency Core Cooling System (ECCS) failed to operate. From this point on, the containment could no longer be regarded as an unbreakable final barrier of radioactivity. This represented a key milestone in reactor regulation, as it modified the fundamental approach to reactor safety. Once the AEC realized that under some circumstances the containment building fail, the key to protecting the health and safety of the public shifted to *preventing* accidents severe enough to threaten containment.

In the late 1960s, two papers were published that brought Probabilistic Risk Assessment (PRA) to the forefront of nuclear engineering thought. The first was a 1967 paper presented at International Atomic Energy Agency's Vienna conference by F.R. Farmer entitled 'Reactor Safety and Siting: A Proposed Risk Criterion'. [10] This paper included the now famous Farmer Curves, and concentrated on the effects of iodine-131. Another paper was a 1969 Science article 'Social Benefit versus Technological Risk' by Chauncey Starr which further elaborated on risk perception and many of Farmer's points [11]. In the meantime, in 1966, the AEC asked General Electric, the contractor at Hanford, and Du Pont, the contractor at Savannah River, to perform calculations on the safety of the plutonium production plants that they operated. Partly influenced by Farmer's paper, General Electric showed, using a very simplistic probabilistic model, that the N-Reactor had a one-in-a-million chance per year for a catastrophic failure because each of the three major subsystems would only fail once-in-one-hundred per years [12]. General Electric and Douglas United Nuclear, the subcontractor that assumed operational control of the N-Reactor in 1967, then claimed that such a low probability meant that for all practical purposes,

the chance of a catastrophic failure was zero, a conclusion that the AEC and others disputed [13].

By 1971, nuclear critics were expressing resentment to the AEC because of the licensing of several reactors under review and its conflicting mission of both regulating and promoting nuclear power. To address the critics' concerns members of the AEC regulatory staff, led by Malcolm Ernst, investigated nuclear reactor safety, licensing, and risk during 1972–1973 using the incident reports from the nuclear power plants. The Ernst report ('Task Force Report: Study of Reactor Licensing Process') concluded that the complexities associated with the design and operation of the reactors then operating exhibited so many technical challenges that a quantified risk assessment would be impossible to produce. One reason for this was that the probabilistic methodologies lacked sophistication and rigor, and the information required was not fully available [14]. However, in 1974, the AEC's regulatory programs came under such strong attacks that Congress decided to abolish the agency. Supporters and critics of nuclear power agreed that the promotional and regulatory duties of the AEC are in conflict and should be assigned to separate agencies. The Energy Reorganization Act of 1974 created the Nuclear Regulatory Commission (NRC) which assumed the responsibility for civilian nuclear power regulation and assuring the protection of the health and safety of the public.

1.3. Probabilistic methods in the aerospace industries

Another industry in which issues of safety and risk were of paramount importance was the aerospace industry. The Boeing Company, in conjunction with Bell Laboratories, pioneered the use of fault tree analysis during the design of the Minuteman missile for the Air Force during the 1960s to prevent inadvertent launches. In 1966, Pan Am Airlines placed an order with the Boeing Company to build the Boeing-747. Because the Boeing-747 would be the largest commercial jet in operation, Boeing engineers felt that it would be important to look at the safety systems of the plane in a different manner than they had in previous aircraft designs. The method they chose was fault tree analysis, which provided a deductive, systematic, and holistic assessment of the airplane, and highlighted among the faults modeled, the critical ones and effects of such faults on the plane. This allowed the designers to appreciate how and why the failure of one system or component would affect other systems [15].

Probabilistic analysis of aircraft grew during the 1970s. A 1979 crash of a DC-10 at O'Hare International Airport led to an assessment of aircraft safety led by George M Low (at that point the President of Renssealaer Polytechnic Institute and former head of NASA). Eventually, in 1982, the FAA required the use of fault trees on new aircraft designs to identify single point of failures and reduce the chances of such failures to less than one-in-a-billion flight hours. However, considering the number of single failures in an aircraft the actual rate for the whole aircraft is probably one-in-20-million flight hours or so, assuming perfect maintenance. Without proper maintenance, the rate of failure could rise dramatically. In 1988, the FAA updated the requirements and added to the breadth of complex systems required to follow the regulation [16].

NASA began to use probabilistic risk assessment methods in 1967, following the disastrous fire on Apollo 1. Engineers from the Boeing Company helped complete a fault tree analysis for the entire Apollo system. They relied on highly conservative measures and data and estimated failure probabilities for Apollo missions to range 0.1–0.8 per mission; a range that was higher than the actual experience, and subsequently led to a distrust of probabilistic risk assessment results. However, following the Challenger explosion in 1986, probabilistic risk assessment at NASA was revived, and the Columbia break-up in 2003 reiterated the need for such analyses. NASA used risk assessment and a combination of fault and event trees methods borrowed from the nuclear industry to model possible accident scenarios for the shuttle and International Space Station (ISS) programs. One risk study performed by the US Air Force in 1983 calculated the chances of a space shuttle solid rocket booster failing during operation to be about 1 in 35, a number disputed by NASA management [17].

While risk assessment methods in the nuclear industry benefited tremendously from the experience of the aerospace industry in the early 1970s, in the late 1980s, when the need for systematic safety assessment became apparent in the space industry, it was the aerospace industry that turned to and relied primarily on the experience of the nuclear industry.

2. Emergence and impact of WASH-1400

As the number of nuclear power plants either under construction or already completed grew and the size of the reactors rapidly increased during the late 1960s and early 1970s, reactor safety became an important public policy issue. Often bitter debates over the reliability of ECCS, reactor pressure vessel integrity, and the likelihood of large accidents consumed the AEC, Congress, the nuclear industry, environmentalists, and the media. In the late 1940s and most of the 1950s, public attitudes toward the technology were highly favorable, as the opinion polls on the subject revealed. In the late 1950s and early 1960s, however, the public became more aware and worried about the hazards of radiation, largely as more was learned about radioactive fallout from nuclear weapons testing. Throughout the 1960s there was a desire and interest among the public to know whether or not nuclear plants were safe. This desire for safety assurance became more urgent in late 1960s and early 1970s as

the organized opposition to nuclear power grew and characterized AEC's safety criteria used for licensing nuclear plants as inadequate and inconsistent with respect to the apparent safety significance of various systems, structures and components of plants.

In 1972, Senator John O. Pastore, then the Chairman of the JCAE, helped initiate the project that became known as the Reactor Safety Study (RSS), better known as the Rasmussen Report (WASH-1400), by sending a letter to the chairman of the AEC, James Schlesinger. At that time, the JCAE controlled the funding of almost all federal nuclear programs and Senator Pastore worked to maintain that control. The purpose of the study was to help with the upcoming extension of the Price-Anderson Act. Also, in the early 1970s, serious concerns were raised over performance of ECCS in the larger reactors coming on-line. The AEC performed a series of experiments at the Reactor Test Facility in Idaho using a small-scale reactor mockup. The Loss Of Fluid Tests (LOFT) suggested that the ECCS might not work as well as planned[1]-steam build-up could prevent injection of water into the core, and lead to core damage [18].

As a result of Senator Pastore's letter coupled with concerns over the LOFT results, the AEC approached Professor Mason Benedict, then a member of the ACRS and a former head of the nuclear engineering department at MIT, to seek his help to initiate and run a study of the assessment of safety of the nuclear plants in the United States. Professor Benedict, who recently had received the Enrico Fermi Award for contributions to nuclear energy, declined the position citing his already busy schedule, but recommended Professor Norman Rasmussen, who, with interest in probabilistic and statistical techniques which form the basis of any formal risk assessment study, could serve as an alternative choice. The AEC accepted this recommendation and started the RSS in 1972.

A slightly different version of the origin of RSS offered by Dr. Herbert Kouts, then at Brookhaven National Laboratory, who notes that Harold Price, then Director of Regulation at the AEC, asked him to personally lead a new study to estimate risks in nuclear power plants. However, Kouts did not think a study could estimate probabilities of accidents, but recommended Professor Norman Rasmussen as one who could try [18]. Kouts later became head of the newly created Division of Reactor Safety Research in 1973, which had the management lead for the RSS, and later served on the Risk Assessment Review Group. Saul Levine became Kouts' deputy and ran the day-to-day operation of the RSS for Professor Rasmussen. A possibly apocryphal version of the origin of the RSS states that in a private meeting between Professor Rasmussen and Saul Levine, a deputy director for nuclear plant research at the AEC, during the annual MIT summer session in 1971, the idea of a comprehensive safety study came up. Then using Mr Levine's contacts with the JCAE, they were able to get the study initiated [19,20].

The staff of the RSS consisted of about 40 scientists and engineers, drawn from industry, academia, and government service. In addition to Saul Levine, seven full-time participants were AEC employees. Many were experts in particular reactor safety systems; some were experts in risk assessment. Also, several outside experts worked in consequence analysis—modeling the release and health effects of radiation in the environment, in which Rasmussen showed a keen interest. The study cost just under $4 million over its three-year span which was significant by the early 1970s standards (~$15 million today after inflation adjustment). The RSS initially used fault trees as the basis for reactor risk calculations, a decision reached by Rasmussen and Levine in September 1972 at the urging of Dr William Vesely [21]. Jonathan Young, an expert in fault trees from Boeing, was one of the leaders in the project with six or seven reactor experts developing prototype fault trees for both the BWR (Peach Bottom Atomic Power Station, Unit 2) and PWR (Surry Power Station Unit I) designs. Although fault trees were developed for almost all of the major safety-related systems, the team realized that integrating the overall fault tree analysis for the entire nuclear power plant was too complex of an undertaking for the RSS, given constraints in time and resources [22]. This led to the development of the event tree concept to model the approximate time-line of the possible accidents scenarios. Originally borrowed from the decision analysis field, Professor Rasmussen proposed the event tree method that was later matured by Mat Taylor, one of the AEC team members of the RSS. Event tree methodology remedied the constraints in time and resources associated with relying on fault tree analysis alone. Later the event tree approach became a predominant force in PRAs.

The event trees looked at two separate areas. The first covered failures in major systems, such as the engineered safety systems. The second investigated the ability of a nuclear power plant's containment system to prevent the spread of radiation in the case of an accident. Event trees start with an initiating event that causes a reactor enter a transient from its steady state operating condition. Initiating events, usually a breach in the coolant system integrity or a reactor transient, covered several possibilities that the AEC did not consider at that time. These potential problems included the possibility of reactor vessel failure and steam generator failures-which had been treated by the AEC as events with negligibly small likelihoods due to the stringent quality requirements for the components. The use of event trees was a pivotal decision that made PRA a practical reality.

[1] The LOFT experiments were later shown to have scaling issues that showed the concerns with the ECCS did not translate into full-scale power reactor safety systems.

Six specific LOCAs were analyzed in detail as initiating events by the RSS:

- Large pipe breaks (6″ to 3″ in diameter)
- Small/intermediate pipe breaks (2″ to 6″ in diameter)
- Small pipe breaks (less than 2″ in diameter)
- Large disruptive reactor vessel ruptures
- Gross steam generator ruptures
- Ruptures in systems that interface with the reactor coolant system

After a LOCA occurs, the engineered safety features (ESFs) of a power plant are used to reduce or minimize the amount of radioactive material that reaches the environment. One key system of the ESF is the ECCS (discussed above). Examples of ESFs include an airtight containment building and large tanks of water and pumps to ensure water flows into a reactor vessel in case of a leak. The goal of the event tree was to decompose any possible process, which could occur following an initiating event that results in the release of radiation, into a set of discrete failure events, such that the probabilities of such events can be estimated. Fault trees were used to model the probability of the events included in the event tree. Therefore, each event tree traced the initiating event all the way through to the eventual failure of containment, and determined the probability of the event.

Besides LOCAs, the RSS team investigated several types of reactor transients as possible initiating events for reactor system failure. In the study, a transient was defined as any significant deviation from the normal operating value of any of the key reactor operating parameters—including all non-LOCA situations that could lead to fuel heat imbalances. Transients could occur from a variety of means, such as equipment failure or human error, and often led to reactor shutdown to limit potential damage to the fuel. The three principal interest areas for transients were instances where the reactor power increased, the coolant flow decreased, or the coolant pressure increased. Each of the three areas could lead to core melting or breach of the reactor coolant system.

Transients were broken into two broad categories—anticipated transients, such as loss of off-site power and loss of feedwater transients, and unanticipated transients (reactor vessel rupture, turbine missiles, sabotage, etc.). After some preliminary assessment of the frequency and consequences of these transients, the team decided that the unanticipated transients' potential contribution to overall risk was small compared to the anticipated transients that produced the same consequences. The relatively low frequency of many of the anticipated transients also led to their removal from overall risk calculations. Preliminary results indicated that the most important transients involved the loss of offsite power and the loss of plant heat removal systems.

By narrowing down the possible initiating events that could cause a radiation release, the event trees allowed the RSS team to reduce the particular fault trees that needed to be investigated. One problem with the method proved to be lack of data to estimate the probability of failure for many of the components involved—a fact noted both in the RSS itself, and in the various criticisms of it. RSS used generic data derived primarily from similar basic components (pumps, valves, etc.) used by related industries or the military. Additional analyses into both common-cause failures, addressing interdependencies within the system, and human error, addressing external factors, had to be performed.

Using fault trees, very complex systems could be broken down into constituent parts and failure probabilities could be assigned to each segment. The failure probabilities took into account the human and common cause failures mentioned above. For the RSS, fault trees were developed for essentially all of the major individual systems included in the event trees.

The accuracy of the study was undoubtedly the best when real-world data based on the same type of equipment used in the reactors was available. Often, other industrial data had to be used—with component failure rates and uncertainties increased, in some places substantially, due to the unique operating environment of a nuclear reactor, especially during an accident, where exposure to high-temperature steam and radiation could cause component failure. In addition judgment from the experts in the field was used. To account for uncertainties due to limited data, failure probabilities of the events modeled were represented by lognormal distributions instead of point estimates. Because the typical probability of failure for the complete system was so small, even a factor of 100 alteration in the failure probability of most components would not produce much over-all change in the system failure probability or, consequently in the reactor safety calculations. Some system components carried more weight through the process, and their associated uncertainties would produce a greater uncertainty for the entire system, but typically these systems had a much smaller uncertainty associated with their failure probability.

Most of the failure rates and probabilities used in the study had uncertainties on the order a factor of 10–100; in some cases—for very low failure rates—the error factor was as much as 1000. The study used a Monte Carlo method to calculate the overall uncertainty associated with estimated risks, using lognormal distribution assigned to the probability of failure of components and events.

Within both the fault trees and event trees, there was no guarantee that all of the various modes leading to significant reactor failure were captured in the study, but it must be noted that the team had significant modeling experience as well as reactor operation and safety experience. All of the systems associated with the primary and secondary core cooling loops and the safety systems were analyzed. The level of detail used in the fault tree could also be questioned, but the study used sensitivity tests to determine whether

the fault trees had gone into more detail than was needed, and determined that the level of detail was adequate.

Following the attempts to model what happen in the reactor during an accident required corresponding calculations of the potential radiation release from the reactor into the containment and ultimately into the environment. After the amount of radiation release was known, the consequences (expected human, economics, and environmental losses) could be estimated. The consequence section of the study was in some ways the most surprising. Before the RSS was released, the general feeling in the nuclear industry was that the consequences of a severe reactor accident would automatically be massive, but the RSS showed that most accidents that led to radiation release would only have small consequences.

The most important element of the consequence analysis was estimation of human exposures and subsequent fatalities and health effects due to any radiation released to the environment. Using the known meteorological and demographic data for each of the existing or planned 68 sites for nuclear reactors in 1974, the RSS was able to calculate the expected radiation pathways and the effects on the near-by residents. The goal was to estimate the most 'realistic' radiation effects, by relying on best estimate values and avoiding, as much as possible, any use of conservative assumptions. Doses from five potential exposure modes were used:

1. the external dose from the passing cloud (plume),
2. the dose from internally deposited radionuclides which are inhaled from the passing cloud,
3. the external dose from the radioactive material which is deposited on the ground
4. the dose from internally deposited radionuclides which are inhaled after resuspension
5. the dose from internally deposited radionuclides that are ingested after ground contamination

Three kinds of effects of radiation were calculated from the total human population dose—early fatalities (within 1 year of exposure), early illnesses (people needing medical treatment), and long-term health effects (additional cancers occurring after a few years). A member of the RSS team, Dr Ian Wall then of AEC, made major contributions in the development of the consequence analysis. He worked closely with Professor Rasmussen and a team of medical and health physics advisors to investigate health consequences. One assumption of the RSS was that medical care would be available for the exposed population. In addition to health effects on the population, the RSS attempted to predict property damage associated with a reactor accident.

The RSS team investigated more than a thousand core melt sequences for the PWRs. Rather than using valuable computer processing time to calculate the amount of radiation release to the environment due to each of the sequences, the team sorted them into 38 general sequences. After running a computer code especially developed to calculate the amount of various radioisotopes released (the CORRAL code) for each of the 38 sequences, the results were sorted into one of nine broad release categories. Similarly, for the BWR, the team created five broad release categories. A brief description of each broad category can be found in Table 1. Of significance was Category 7 for PWRs, which was largely an unknown scenario in the nuclear industry.

Aside from the LOCA and transient initiating events, the study attempted to estimate the general magnitude of risk associated with earthquakes, floods, tornadoes accidental aircraft impact (note—the effect of a deliberate crash on a nuclear power plant was not addressed), and turbine missiles. The study calculated that the probability of reactor failure associated with each of the external events was small compared to the overall calculated risk. During the preparation of the report, the Brown's Ferry fire of 1975 occurred, and was commented upon as requiring further

Table 1
Radioactive release categories of RSS

Category	Brief release description
PWR Cat 1	Steam explosion in the reactor vessel, containment failure
PWR Cat 2	Core melt—with failure of radioactive removal systems
PWR Cat 3	Core melt—partial success of radioactive removal systems
PWR Cat 4	Core melt—containment not fully isolated, radioactive removal system fails
PWR Cat 5	Core melt—similar to cat 4 with partial radioactive removal system success
PWR Cat 6	Core melt through reactor vessel—radioactivity removal system operates
PWR Cat 7	Core melt through reactor vessel—radioactivity removal system failure
PWR Cat 8	Fuel failure—containment fails to isolate properly
PWR Cat 9	Fuel failure—containment operates correctly
BWR Cat 1	Steam explosion in the reactor vessel, containment failure
BWR Cat 2	Core melt after containment rupture—no or little internal deposition
BWR Cat 3	Overpressure rupture of containment, with significant deposition
BWR Cat 4	Containment isolation failure, no rupture
BWR Cat 5	Fuel failure—release through stack

study. The RSS also noted that straightforward measures to improve fire prevention and fire fighting capabilities could significantly reduce the risk of reactor failure from a fire [23].

Possibly the most controversial part of the RSS was the comparison of risks from nuclear power plant failures to other more common or extreme rely remote risks encountered by the general public. Among the risks used to illustrate the issue were automobile accidents, hurricanes, tornadoes, earthquakes, meteorites, airplane crashes, explosions, dam failures, fires, and industrial accidents leading to hazardous chemical releases. For chemical releases, the RSS used a generic chlorine release on a major rail line in Ohio. The closest in risk to 100 operating nuclear power plants of any of the examples illustrated in the RSS was the risk of a large meteorite impact-estimated to be about 10^{-4} (or 1-in-10,000 chance) that 10 people being killed and about 10^{-7} that 10,000 people being killed.

The RSS calculated the frequency of core melt for a PWR to be $\sim 6 \times 10^{-5}$ per reactor year and for a BWR to be $\sim 3 \times 10^{-5}$ per reactor year. The major change in the failure frequency for PWRs from early studies (usually about 10^{-6}) is that prior estimates tended to ignore or downplay the small LOCAs' contribution to core melt, whereas the RSS determined that small LOCAs had the highest contribution to the overall risk. The RSS determined that the two largest contributors to BWR frequency of core melt were the failure to rapidly shut down the reactor when needed and the failure of the decay heat removal system after transient-caused shutdowns. Previous calculations and estimates had also concentrated on worst-case scenarios when determining the consequence of reactor accidents. The RSS showed that the majority of core melt accidents would produce modest consequences; with only a very small portion of the core melt scenarios causing catastrophic off-site damage as envisioned in WASH-740. The RSS found that the conditional probability of containment failing, given occurrence of an accident sequence that releases radiation into the containment atmosphere, was higher than originally believed, although often with much of the radioactive material being deposited inside the containment building before the containment failure.

The RSS modeling effort succeeded in producing an accurate and far more realistic result, compared to previous efforts, by using event and fault trees, looking at the interaction of a molten core with the containment system, investigating common cause and human failures, understanding the safety significance of support and other 'non-safety' systems and structures, and determining possible problems in operating, testing and maintenance procedures. In many cases, the RSS pioneered the investigation of nuclear safety issues from a risk perspective.

It took the team working on the RSS about 3 years to complete the study. While much of the time was spent preparing fault or event trees or running calculations, there was a 2-month hiatus while Professor Rasmussen had frank discussions with one or more AEC commissioners regarding the RSS's treatment of the reliability of the ECCS. Other members of the regulatory staff intensively reviewed the RSS; many staff members were skeptical that the study could determine the risk associated with nuclear power [24]. Outside of the team actively working on the project and a few supporters elsewhere in the AEC (and later the NRC), there was little support for the effort.

The AEC finally published an initial draft of the RSS in August 1974, and attempted to gather peer review comments. Following its release, the study received significant public and media attention, leading to influential and effective interviews, representations, defense and debate over the results and methods of the RSS by Professor Rasmussen. One panel of scientists, organized by the American Physical Society, criticized much of the report, especially the fatality estimates that considered only fatalities from radiation absorbed in the first 24 h after an accident. The APS considered radioactive cesium and strontium (both with half-lives near 30 years) to be major contributors to any radioactive exposure of the population. Other reviewers within the APS group also criticized the treatment of ECCS in the RSS. Other groups contributed extensive comments to the RSS team as well, although the technical sophistication of some of the groups was low [25].

When the NRC was finally established in 1975 and took the ownership of Draft RSS report, it published the report in October of 1975 in final form. The one section of the report that was most commonly read was the Executive Summary, which had two sections—a section summarizing the results and another comparing the risk associated with nuclear reactor failures with risks from other man-made events and natural occurrences along with a frequently asked questions. Using a few diagrams, the RSS effectively communicated that the risk associated with the operation of 100 nuclear power plants is much lower than the risk associated with automobiles, airplane crashes, or hurricanes. Unfortunately for the report, the Executive Summary became a controversial issue later.

Following the publication of the RSS, many members of the NRC attempted to disown the study. Members of the NRC staff, who had worked on the study, had to answer internal questions for a period of 4 or 5 months [26]. In June 1976, the Committee on Interior and Insular Affairs of the US House of Representatives held hearings on the findings of the RSS. Chaired by Representative Morris Udall, the hearings found that RSS seemed to be misleading in the certainty and comprehensiveness of its conclusions. The committee also focused on the worst possible postulated accidents without taking into account the probabilities associated with them compared to less damaging accidents. Rep. Udall suggested that a new executive summary could be written that would solve these problems, and that an outside review panel be formed to take a closer look at how the study arrived at its conclusions. Marcus Rowden, the chairman of the NRC, disagreed that a re-written executive

summary would be of much use, but did agree that an outside panel looking at the study would be of value. Commissioner Rowden asked Dr Harold Lewis of the University of California-Santa Barbara to chair the Risk Assessment Review Group. Dr Lewis had also chaired the American Physical Society review of the RSS in 1975.

To its authors, it was evident that the RSS methodology was a unique and powerful means by which to improve regulations and licensing of nuclear power plants [27]. However, it was unclear how the methodology could be used in regulatory decision-making. Unfortunately, because the study became the centerpiece of a fierce political controversy over the safety and acceptability of nuclear power in the US, the NRC did not immediately embrace the power of the new methodology, and inhibited the overall acceptance of RSS's methodology for safety design and regulation. If it were not for its steadfast defense by Professor Rasmussen following the release of the study, the PRA methodology and RSS results might have been doomed altogether. He also recognized the importance of educating a new generation of nuclear engineers with in-depth knowledge of the field. As a first step he developed a graduate level course at MIT in reliability and risk analysis which ultimately drew several students into the field. Subsequently, in 1974–1984, he supervised several doctoral and master theses in this area, including the PhD Thesis of the second author of this paper.

3. The Lewis committee report on WASH-1400

The first meeting of the Risk Assessment Review Group (also known as the Lewis Committee) was in August 1977. The seven-member group, chosen by Dr Lewis, included three university professors, two scientists at national laboratories, one scientist from the EPA, and an engineer from the Electric Power Research Institute. Over the next 13 months, the review group met monthly and in September 1978 finished its report, now almost universally known as the Lewis Committee Report. The Lewis Committee expressly thanked the NRC for its assistance in the preparation of the report and for the complete autonomy given to the group [28].

During its review of the RSS, the Lewis Committee found several good qualities, along with many issues the RSS handled poorly. The good qualities included the use of fault-tree/event-tree methodologies, the improvement in making the study of reactor safety more rational, the identification of important accident pathways previously under-investigated by the NRC (specifically the containment bypass, anticipated transients without scrams, small LOCAs and consideration of human errors), the first systematic calculation of accident consequences beyond early fatalities, and establishment of the procedures so that quantitative estimates of risk could be further explored. The report went on to say, 'Despite its shortcomings, WASH-1400 provides at this time the most complete single picture of accident probabilities associated with nuclear reactors. The fault-tree/event-tree approach coupled with an adequate data base is the best available tool with which to quantify these probabilities'.

Among the shortcomings that the Lewis Committee identified in the RSS were the lack of scrutability of the calculation/analysis process, the lack of accurate data on which to base component reliability estimates, the finding that some external accidents (earthquakes, fires, human accident initiation) contribute negligibly to the overall risk, the peer review process used by the NRC, the difficulty in finding within the report some of the health impacts caused by a radiation release, and the poorly conveyed information in the Executive Summary. Probably the most important criticism of the report was not about the report itself, but rather how the report was being used: 'There have been instances in which WASH-1400 has been misused as a vehicle to judge the acceptability of reactor risks. In other cases it may have been used prematurely as an estimate of the absolute risk of reactor accidents without full realization of the wide band of uncertainties involved. Such use should be discouraged'.

Although the use of lognormal distributions to model the probability of failure uncertainties was criticized by the Lewis Committee, especially when the failure data did not appear to correspond to a lognormal distribution, the Committee also acknowledged that the use of lognormal distributions when the true nature of uncertainty was unknown meant that the calculated value would probably remain within a factor of two or three of the 'true value' of the system. The Committee essentially indicated that the use of lognormal distributions was a problem, but the committee did not have a better solution to how to account for the 10–1000 fold uncertainties in failure probabilities due limited data, nor did they think the error due to the use of lognormal distribution model as opposed to another distribution model caused the risk estimates to be changed by more than a factor of two or three—far smaller than the overall uncertainties and errors associated with failure data and other models used by the RSS. Another criticism involved statistical issue of the 'Square-Root Bounding Model', which was an attempt by the RSS team to estimate common cause failure probabilities. The problems with the bounding model were the arbitrariness of the method and the lack of reality in the estimates of the lower (and in some cases upper) bounds. The Committee recommended staying away from point estimates, but instead recommended using bounds to account for uncertainties based as much as possible on actual operating data, which was admittedly lacking during the RSS calculations. The Committee also recommended that common sense be applied to bounding estimates to inject realism into areas where data did not exist. An additional problem identified was the use of medians instead of means whenever the RSS gave

a lognormal distribution, merely from a scrutability standpoint.

Possibly the biggest problem that the Lewis Committee had with the RSS was the way that the RSS identified and carried uncertainties through the calculations. The committee identified several areas where the RSS did not address uncertainties well, including the use of models, variations between reactors, propagation of errors-where assumptions were used rather than experimental data or models, assignment of uncertainties to the assumptions, and how the overall uncertainties were calculated.

The radiation effects calculations was another area where the Lewis Committee had a series of contentions with the RSS, although many of the problems listed in the Lewis report (effect of very-low doses on populations, acute lethal dose, cancers and other long-term affects, in vitro radiation, rate of radiation exposure) were not simply criticisms of the RSS itself, but rather criticisms with the health physics community for not having consensus answers for the questions asked in the study. Because there were few data available on the effects of radiation exposure in humans, with the survivors of Nagasaki and Hiroshima by far the largest population to study,[2] the questions raised in this area by the Committee Report still remain valid today.

The Lewis Committee indicated that the area of human errors was one in which the NRC should conduct further research, because the committee was unable to determine if the contribution to the estimated risk associated with human interventions was accurately portrayed in the RSS.

One final area of concern was the presentation of risk and the attempt by the RSS to manage the risk perception of the public at large without regard to potential benefits associated with other man—caused accidents (such as the convenience and time saving that the air transportation offers). By displaying the risk information the way it was in the RSS, the committee indicated that the RSS was prejudging an acceptable level of risk for nuclear energy, especially when the whole fuel cycle (mining, milling, waste disposal, transportation, processing, and plutonium production) are not taken into account.

In summary, the Lewis Committee commended the RSS for its ambitious undertaking and the methods used. The report indicated that the assessments in the report appeared to be good where the data were firm, but in many places the data were limited, and the potential errors associated with poor data and subjective reasoning were neither properly labeled nor explained. Finally, the Committee blasted the Executive Summary, calling it not a summary of the report but a separate, advocacy document.

[2] The Chernobyl accident may also provide a significant population of radiation exposure victims.

4. Post RSS review and the three mile island accident

Following the Lewis Committee Report of September 1978, the NRC withdrew its support of the RSS results and disavowed the Executive Summary, but the Commission tried to get the NRC staff to use PRA techniques in general. However, as the staff only understood deterministic analysis, nor cared to learn probabilistic analysis, the Commission's call for uses of probabilistic techniques fell into deaf ears. In fact, many in the NRC took the Lewis Committee Report to be damning criticism (the NRC stated, 'In particular, in light of the Review Group conclusions on accident probabilities, the Commission does not regard as reliable the Reactor Safety Study's numerical estimate of the overall risk of a reactor accident [29]'). A team of NRC employees spent 2 months rewriting NRC papers and documents by removing any reference to the RSS. About the time this task was completed, the accident at the Three Mile Island Unit 2 occurred in March 1979 where about half of the reactor core melted. Fortunately, the crisis ended without a major release of radiation or a need to order a general evacuation, but the event indicated that major accidents not necessarily addressed in the formal reactor licensing process are possible, and that new approaches to nuclear regulation were essential.

While the RSS had considered a similar sequence of events for a reactor other than TMI, and showed that this sequence was not among the risk-significant contributors for that reactor design, the TMI accident confirmed a major RSS insight that small LOCAs are more risk-significant than large LOCAs that the NRC used as a design basis accident for worst-case LOCAs in licensing reactors.[3] Also, the RSS pointed out the potential role of human error, which showed itself to be a highly significant factor in the TMI accident when operators turned off the ECCS (despite the fact that this particular error was not considered by the RSS). As a result, the NRC had a change of heart and it spent another 2 months restoring all of the modified references to the RSS in its documents [30]. The NRC made the decision in part because one of the accident sequences studied in the RSS—where the pressurizer relief valve failed to close—was very similar to what actually occurred at TMI accident, when adjusted for the differences in the reactors. That particular sequence had not been identified as a potential problem before the RSS. Subsequently, the NRC placed much greater emphasis on operator training and 'human factors' in plant performance, investigating severe accidents that could

[3] This should not be confused with the fact that Large LOCAs still cause the highest peak clad temperature and the most physically limiting operating conditions. But since small LOCAs are more frequent, despite of their less severe physical conditions, their total contribution to the total plant risk, Core Damage Frequency (CDF), and Large Early Release Frequency (LERF) is higher than large LOCAs. As such, the worst-case design basis accident approach used by the NRC overlooked the potential consequences of small LOCAs as significant contributor to risk.

occur as a result of small equipment failures (as occurred at Three Mile Island), emergency planning, plant operating histories, and other similar matters.

The report of Presidential Commission (Kemeny Commission) appointed by President Carter to evaluate the accident at TMI contained many criticisms of the NRC as well. The RSS, while not being the complete answer to these criticisms, at least showed one possible solution for the NRC to consider in terms of safety regulations. Another key finding in the report was, 'With its present organization, staff, and attitudes, the NRC is unable to fulfil its responsibility for providing an acceptable level of safety for nuclear power plants'. Although not specifically discussing the disdain much of the NRC felt toward the RSS, the report recommended that the NRC establish a requirement to analyze safety-cost trade-offs; something much more easily performed using RSS type analyses. Also, the Commission noted that the NRC and the industry were too focused on the 'design-basis (large) accidents' that have great consequence, but low probability of occurrence [31].

During the hearings associated with the TMI investigation, both Mr Levine and Professor Rasmussen gave depositions indicating the areas where the RSS could have been used, and possibly just as important, the fact that the NRC, had endorsed a plan to apply the same techniques to other reactors. One additional area where the experience gained from the RSS had an input into the TMI report was in the area of future research. At the time, reactor safety research had concentrated on the large LOCA, usually the extremely rare event of a double-ended guillotine pipe break. The TMI report suggested that future reactor safety research should be consistent with priorities determined by their relative risk contributions, and should look not only at LOCAs, but also at transients. Only the RSS provided much information about relative risk contributions of accident scenarios. The TMI report also commented on the need to be cognizant of the fact that operators were running the nuclear power plant, and that a 'mindset' down-playing the role of humans in the safety process existed; as a result, the NRC increased its human-factors studies as noted above.

5. Post-TMI accident and revival of the uses of PRA

Following both the publication of the Lewis Committee report and the accident at Three Mile Island, the NRC began to devote some additional resources towards the expansion of the uses of PRA in the industry. During the late 1970s, Mr Levine, in his role as the Director of the Office of Nuclear Regulatory Research, along with a few other NRC staff initiated a number of PRA studies. For example, the first risk-based prioritization of generic safety issues, and completed a study on the reliability of auxiliary feedwater systems [32]. The NRC also initiated research on steam explosions under various conditions, possible interactions between a molten core and the underlying concrete, and study of the basic processes involved with release of radioactive materials from molten fuel-three areas identified in the RSS as needed further study.

During 1979–1982 the NRC undertook two sets of follow-up PRA studies. The Reactor Safety Study Methodology Application Program (RSSMAP) to apply the RSS methodology to additional reactor designs [33]; and the Interim Reliability Evaluation Program (IREP) was a planned multiplant reliability evaluation program to develop and standardize the reliability methodology involved in performing reliability and safety studies. The IREP was conceived in NRC as reported in 'Action Plan Developed as a Result of the TMI-2 Accident' NUREG-0660 [34] was a pilot study with a scaled-up evaluation of an additional six plants. Recognizing lack of any formal guiding documents to perform PRAs, an effort to develop such documents was proposed by the American Nuclear Society and the Institute for Electrical and Electronics Engineers to the NRC and initiated in 1983. The resulting documents published methods for performing PRAs for nuclear power plants (NUREG/CR-2300, 'PRA Procedures Guide: A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants') [35].

In the early 1980s the NRC relied on PRA techniques in addressing a few of its unresolved safety issues involving beyond design basis accidents. The most notable were the Anticipated Transient Without Scram (ATWS) and Station Blackout rules. The so-called backfit rule which attempted to remedy some of the safety concerns that surfaced following the TMI accident was also addressed with the help of the PRA techniques. Finally, the risk significance of incidents reported to the NRC by the plant owners (in the so-called Licensee Event Reports) were analyzed using the PRA methods. In these studies the conditional frequency of core melt due to the occurrence of these incidents were estimated, by viewing at these events as 'precursors' to severe accidents. The early precursor studies received a lot of publicities because they predicted rather higher conditional frequency of core melt than was anticipated. The controversy diminished as the methods for accounting for these precursor events improved in the subsequent precursor studies.

Parallel to the NRC's efforts related to PRAs, during the early- and mid-1980s, several plant owners completed PRAs of their own in order to facilitate technical upgrades, or characterize risk to local populations. When the PRAs for Zion and Indian Point -2 and -3 were published in 1981 and 1982, they showed the risk associated with earthquakes and fires was not negligible, as was concluded in the RSS, but a significant one that required further study. However, all the industry-supported studies confirmed the general insights of the RSS. Furthermore, the studies focused on more advanced methodologies to determine the uncertainties more systematically than the RSS's approach. In the meantime, the nuclear industry sponsored additional efforts

to improve PRA techniques (for example in modeling common cause failures), and undertook additional PRAs in support of their licensing efforts (e.g. the Seabrook Station PRA effort in 1982–1983). Most of these efforts were done through the consulting firm Pickard, Lowe and Garrick (PLG) that employed an impressive PRA team of engineers and consultants, led by a pioneer in this field, Dr B. John Garrick.

Subsequent to the Kemeny Commission's report in 1986, and based on the idea of the long-time ACRS member Dr David Okrent, ACRS started an extensive debate, public workshops and meetings, which led to the release of the NRC's final policy statement establishing qualitative safety goals, and associated quantitative health objectives to be used for measuring the attainment of these goals. The policy statement was intended to come to grips with the integration of the quantitative assessment of risk into the regulatory system. During the deliberations by the ACRS, RSS methodology was the clearest approach to measure the quantitative safety goals proposed. The primary issue for the NRC in developing safety goals was how to use the PRA techniques to help articulate a level of acceptable risk-in other words, to define 'how safe is safe enough'. [36].

Two safety goals introduced by the NRC were stated in terms of public health risk—one addressing individual risk and the other addressing societal risk. The risk to an individual is based on the potential for death resulting directly from a reactor accident—that is, a prompt fatality. The societal risk is stated in terms of nuclear power plant operations, as opposed to accidents alone, and addresses the long-term impact on those living near the plant. The goals were expressed in qualitative terms, perhaps so the philosophy could be understood. The NRC also expressed the qualitative goals for the safety of nuclear power plants in terms of individual and societal 'quantitative health objectives'. The quantitative goals did indirectly impacted the NRC's regulations, as the goals provided indices as to the level of 'public protection which nuclear plant designers and operators should strive to achieve'. They were also meant to provide additional guidance to the NRC staff as part of their regulatory decision-making process. While the safety goals provided a metric to address the question of 'how safe is safe enough', practical implementation of the NRC's guidance proved to be difficult because of the large uncertainties involved in calculation of risk [36].

In 1986, the NRC started work on what would become NUREG-1150, 'Severe Accident Risks: An Assessment for Five US Nuclear Power Plants', which was essentially an update of the RSS, with 10 more years' additional operating experience, PRA knowledge, and methods gained following the RSS release. NUREG-1150 was published in final form in December 1990, following a long and extensive review process, both internally within the NRC, and also by the American Nuclear Society and the International Atomic Energy Agency. This study was the most important step forward for the NRC following the release of the RSS; several areas of safety, such as mechanisms of failure and potential large loads, were investigated. NUREG-1150 showed that the risks of severe accidents were lower than those calculated in the RSS, primarily through the use of a larger database and more sophisticated models, but due to large uncertainty bands in RSS, the risks outlined in the RSS were captured within NUREG-1150's results.

To make the risk technology and methods available to the industry, in November 1988, the NRC issued Generic Letter 88-20, 'Individual Plant Examination for Severe Accident Vulnerabilities'. This letter acknowledged that each nuclear power plant is unique and may have plant specific vulnerabilities. The NRC required each plant owner:

1. to develop an appreciation of severe accident behavior;
2. to understand the most likely severe accident sequences that could occur at its plant;
3. to gain a better quantitative understanding of the overall probabilities of core damage and fission product releases; and
4. if necessary, to reduce the overall frequency of core damage and fission product releases by modifying, where appropriate, hardware and procedures that would help prevent or mitigate severe accidents.

The individual plant examination (IPE) laid out the process for each plant owner to gain experience with PRA[4] by using its own staff as much as possible to perform the examination. Furthermore, the Generic Letter gave several additional benefits for performing PRAs—support for licensing actions, license renewals, risk management, and integrated safety assessment.

In Generic Letter 88-20, the NRC discussed what a PRA was and how the industry use it in the future. As a result of the Generic Letter, 74 PRAs with varying degrees of detail, representing 106 US nuclear power plants were completed by 1992. Each of the PRAs looked at the reactor core damage frequency (CDF) and the Large Early Release Frequency (LERF), giving the utilities appreciation for PRA methods and a method of tracking the improvements made on the reactor in terms of risk abatement and cost-effectiveness.

In 1990, the NRC provided additional guidance to the staff regarding the Safety Goals, endorsing surrogate objectives concerning the frequency of core damage accidents and large releases of radioactivity [37]. The numerical value of one-in-ten-thousand for core damage frequency (CDF) was cited as a 'very useful subsidiary benchmark…'. In addition, a conditional containment failure probability of one-tenth was approved for application to evolutionary light water reactor designs. This resulted in a large release frequency of one in

[4] The NRC did not require the plant owners to use the PRA, but encouraged them to do so.

one-hundred-thousand, since containment failure is necessary for a large release to occur. These values later evolved into the 'benchmark' values of 10^{-4} for CDF and 10^{-5} for LERF.

As a direct outcome of the nuclear industry's knowledge of PRA methods, results and uses gained through their IPE studies, some industry leaders began to lobby the NRC commissioners and staff in 1992–1993 to base some of their regulatory and enforcement efforts on PRA results, tools and techniques. By 1995, the use of PRAs had been well established in the nuclear industry. As a result, the NRC issued its PRA policy statement directing that the NRC staff use PRA for all regulatory matters to the extent supported by the state of the art in the field. However, the NRC also made clear that the defense-in-depth policy would remain as an important element of licensing and regulatory decision making. This policy statement effectively introduced a new regulatory paradigm called Risk-Informed Regulation (RIR) whereby PRA results in concert with traditional deterministic analyses is to be used for regulatory decision making. Also in 1995, the Electric Power Research Institute published, 'PSA Applications Guide' to help the industry formalize decision-making processes using probabilistic safety assessments (PSAs) [38]. The NRC published a series of Regulatory Guides (RGs) in 1998 to further define how PRA results should be evaluated in its newly adopted RIR approach. These guides included:

- RG 1.174 for changes to plant licenses
- RG 1.175 for in-service testing
- RG 1.176 for graded specifications
- RG 1.177 for technical specifications
- RG 1.178 for in-service inspection of piping

The move toward RIR was a significant transition at the NRC fueled by the premise that a reduction in unneeded expenditure of resources on matters that are not safety significant is required to make nuclear power safer. The industry welcomed RIR because it also observed that nuclear plants can be run more effectively and economically. The challenge, however, has been to accomplish this transition while maintaining the basic objectives of adequate protection of the health and safety of the public [39]. The safety principles articulated in the Regulatory Guide 1.174 address additional considerations relevant to adequate protection that are not directly or fully captured in PRA (e.g. maintaining sufficient safety margins and monitoring performance).

Possibly the greatest benefit that the NRC has found in the use of PRAs by utilities is that it has required the utilities to write down all of the assumptions involved in reactor operation and safety systems. Prior to PRAs, many of the assumptions would never be explicitly stated in technical specifications or other design documents. By stating the assumptions explicitly, the utility can gain a better understanding of the function of the reactor and its safety systems.

In 1998, the NRC introduced its new Reactor Oversight Process (ROP), with the concept of seven cornerstones as a basis for defining the safety scope in its new safety oversight model to be consistent with its mission of protecting the public health and safety with respect to civilian nuclear power plant operation. This mission is then broken down into three strategic safety performance areas of reactor safety, radiation protection, and safeguards. The cornerstones of safety that were associated with each of these strategic performance areas are basically the safety functions or objectives that are needed to meet each of the strategic areas and assure that the overall safety mission objective is met. The NRC announced its safety philosophy by defining objective thresholds for these seven cornerstones. The first four cornerstones are primarily derived from the PRA approach to plant safety: Initiating Events, Mitigating Systems, Barrier Integrity, and Emergency Preparedness. The others include Occupational Radiation Safety, Public Radiation Safety, and Physical Protection.

In addition to the cornerstones, the reactor oversight program features three 'cross-cutting' elements, so named because they affect and are therefore part of and influence each of the seven cornerstones:

1. Human performance,
2. Management attention to safety and workers' ability to raise safety issues (The so-called 'safety-conscious work environment'), and
3. Finding and fixing problems [40].

From 2000 to the present, to facilitate the move toward RIR, the NRC has initiated many activities in all areas of ROP, developing guidance on risk-informed licensing basis changes, and developing risk analysis methods, tools, and data. Examples of these are:

- risk-Informing Part 50
  Hydrogen Control Requirements (10 CFR 50.44)
  Emergency Core Cooling System (ECCS) Acceptance Criteria (10 CFR 50.46)
- risk-informed technical specification initiatives
- performance-based risk-informed fire protection standards and analysis
- draft Regulatory Guide, DG-1122 including work with ASME and ANS for developing PRA standards
- developing technical bases for revision of the Pressurized Thermal Shock (PTS) screening criteria in the PTS Rule (10 CFR 50.61)
- revising NUREG/CR-6595, 'An Approach for Estimating Frequencies of Various Containment Failure Modes and Bypass Events'
- assessing risk of dry storage of spent nuclear fuel and amending 10 CFR Part 72
- developing risk guidelines and a risk-informed decision-making process for nuclear materials and medical applications

- developing human reliability analysis methods and tools for reactor, materials and waste applications
- investigating feasibility of risk-informed technology—neutral regulations.

Areas of further study remain. Among these are further research efforts investigating:

- human reliability, especially acts of commission; the effects of aging infrastructure and facilities;
- treatment of uncertainties;
- events at low power or during changing power levels;
- containment modeling under a variety of accident conditions;
- terrorist-caused events, such as sabotage, truck bombs, or other physical insults to a reactor facility.

6. Current Issues with the Use of PRAs

Clearly the NRC is still grappling with how to use the mix of the traditional regulatory approaches such as the defense-in-depth, safety margin, adequate protection and conservative deterministic calculations along with PRA approaches. Despite of the attempts in the Regulatory Guide 1.174 to clarify and articulating the role of each, a consistent approach to reactor safety philosophy has yet to come. It is important and entirely appropriate to establish a clear, consistent, and well-understood statement of safety philosophy and the meaning of adequate protection. In current NRC practice there is a legal presumption that substantial compliance with the deterministically based regulations provides adequate protection of public health and safety [41].

The move to RIR and the change in the NRC culture will also be based on the ability of the NRC to establish policies based on defining adequate protection and on creating a clear, consistent, and well-understood statement of safety philosophy [41]. The concept of 'adequate protection' and the equivalent phrase 'no undue risk' are not explicitly and concisely defined in the Atomic Energy Act of 1954. Quantitative (absolute) risk estimates serve as an important measure of plant safety, but do not embody the full range of considerations that enter into the judgment regarding adequate protection. The judgment regarding adequate protection derives from a more diverse set of considerations, such as acceptable design, construction, operation, maintenance, modification, and quality assurance measures, together with compliance with NRC requirements including, license conditions, orders, and regulations. Recently ACRS has explored regulatory frameworks with and without the concept of the defense-in-depth [42]. If one had complete confidence in the accuracy of PRAs, one might conclude that defense-in-depth could be ignored if the risk were sufficiently low [43].

Dr Richard Meserve, former Chairman of the NRC, put the issues into clear perspective: "…we are grappling with the possibility that we may have to develop a new regulatory system that, unlike the focus of the current rules on light water reactors, will be independent of technology. The foundation of any such system must inevitably include compliance with the safety goals—or their subsidiary objectives-as demonstrated by PRAs" [43].

7. Conclusions

The use of PRAs may not eliminate the need for properly designed engineering safety features, the use of engineering design safety criteria, safety margins, or defense-in-depth. PRAs do allow the operators of a nuclear power plant to determine where the weaknesses are in safety systems, and to properly allocate resources to correct important potential safety problems. Although using risk analysis to help with decision making has a number of advantages, it took over twelve years from the publication of the Reactor Safety Study in 1975 until the NRC produced Generic Letter 88-20 in 1988, formally enabling the use of PRAs in the industry.

There are several reasons for this delay; foremost was the lack of understanding of just what a risk assessment was, and how it would be used. Second, most engineers tend to stick with the methods that they learned, and through the 1960s and 1970s, risk analysis education was not widespread and the NRC was dominated by staff comfortable and familiar with a deterministic/structuralist school. Finally, the administration of the NRC was not comfortable with the concept, partly because of the initial reception of the Reactor Safety Study and partly due to the idea behind the quote from Max Planck in the introduction, "A new scientific truth triumphs not because its opponents become convinced and finally see the light, [but] rather, because they eventually die and a new generation is born which is familiar with the new concepts". This statement might be too black and white, but for the most part adequately describes how the use of PRA became acknowledged as useful and later as fundamental by the NRC—engineers and scientists familiar with the process had to move into positions of power and policy-making to facilitate the use of PRAs. Ultimately, NRC appreciated the power of PRA and became the pioneer, the leading governmental agency, and the source of knowledge in the use of risk information in safety regulation.

The RSS was a pioneering event in the use of risk assessment in the nuclear industry, and looking at broader regulatory issues, it influenced not just the NRC, but also federal agencies such as the EPA and NASA. Professor Rasmussen's hard work and dedication to undertake such a monumental task, especially considering the state of risk analysis in the early 1970s, the complexity of the task, and the lack of support by parts of the regulatory agency was extraordinary. Professor Lewis sums up the Professor Rasmussen's work the best, "I always had enormous respect for him—he was an honest man, doing a tough job". [44] (Box 1).

Norman C. Rasmussen was born in Harrisburg, Pennsylvania on 12 November 1927. Following naval service in World War II, undergraduate education at Gettysburg College and graduate education at the Massachusetts Institute of Technology (MIT), he became a physics instructor at MIT in 1956. Dr Rasmussen received a professorship at MIT in 1958, and served in the nuclear engineering department until 1994. His initial research concentrated on investigating radiation and gamma rays. He was the head of the nuclear engineering department from 1975 to 1981. Among his numerous honors was his election to both the National Academy of Engineering (1977) and the National Academy of Sciences (1979), as well as serving a 6-year term on the National Science Board during the Reagan Administration. Professor Rasmussen won the Enrico Fermi Award for excellence in the field of nuclear energy in 1985 for his 'pioneering contributions to nuclear energy in the development of probabilistic risk assessment techniques that have provided new insights and led to new developments in nuclear power plant safety'. Perhaps his most remembered moment was his televised debate with activist Ralph Nader over the safety of nuclear power. Professor Rasmussen passed away on 18 July 2003.

References

[1] Planck M. Autobiographia scientifica, Turin; 1956.
[2] US NRC. WASH-1400: reactor safety study (NUREG-75/014); 1975.
[3] US NRC. A short history of nuclear regulation, 1946–1999 (NUREG/BR-0175, Rev. 1).
[4] Rhodes R. The making of the atomic bomb. New York: Simon and Schuster; 1986.
[5] Carlisle R. Probabilistic risk assessment in nuclear reactors: engineering success, public relations failure. Technol Culture 1997; 38:920–41.
[6] Frankel E. Systems reliability and risk analysis, 2nd ed. Boston: Kluwer Academic Publishers; 2002.
[7] Green A, Bourne A. Reliability technology. London: Wiley; 1972.
[8] US AEC. WASH-740, Theoretical possibilities and consequences of major accidents in large nuclear power plants (AKA The Brookhaven Report); 1957.
[9] Wood W. Nuclear safety, risks and regulation. American Enterprise Institute—Public Policy Research; 1983.
[10] Farmer F. Reactor safety and siting: a proposed risk criterion. Nuclear Safety 1967;539–48.
[11] Starr C. Social benefit versus technological risk. Science 1969;19: 1232–8.
[12] Carlisle R. Probabilistic risk assessment in nuclear reactors: engineering success, public relations failure. Technol Culture 1997;38:920–41.
[13] US AEC., Minutes of the AEC general advisory committee; 12–14 July 1966.
[14] Ford D. A history of federal nuclear safety assessments: from WASH 740 through the reactor safety study. Washington: Union of Concerned Scientists; 1977.
[15] Ericson C. Fault tree analysis—a history. 17th International System Safety Conference, Orlando; 1999.
[16] Federal Aviation Administration. Advisory Circular 25.1309-1A; 1988.
[17] Colglazier E, Weatherwas R. Failure estimates for the space shuttle. Abstracts of the society for risk analysis annual meeting, Boston; 1986. p. 80.
[18] Kouts H. History of safety research programs and some lessons to be drawn from it. 26th water reactor safety information meeting, Bethesda; 1998.
[19] Interview with Joseph A. Murphy; Dec. 22, 2003.
[20] Kouts H. History of safety research programs and some lessons to be drawn from it. 26th water reactor safety information meeting, Bethesda; 1998.
[21] Interview with Dr William Vesely; Dec. 12, 2003.
[22] Interviews with Joseph A. Murphy Dec 22, 2003 and Dr. William Vesely, Dec. 12, 2003.
[23] US NRC. WASH-1400: reactor safety study (NUREG-75/014); 1975.
[24] Interview with Joseph A. Murphy; Dec. 22, 2003.
[25] Lewis H, et al. American physical society reactor study review group report on WASH-1400; 1975.
[26] Interview with Dr William Vesely; Dec. 12, 2003.
[27] Wall I, Haugh J, Worledge D. Recent applications for PSA for managing nuclear power plants. Progr Nucl Energy 2001;39:367–425.
[28] Lewis H, et al. Risk assessment review group report to the US Nuclear Regulatory Commission. NUREG/CR-0400; 1978.
[29] US NRC. Statement on risk assessment and the reactor safety study report; 1978.
[30] Interview with Dr William Vesely; Dec. 12, 2003.
[31] Kemeny et al. Report of the President's Commission on the Accident at Three Mile Island; 1979.
[32] Murphy J, Budnitz R. Newsletter—nuclear installations safety division of the American Nuclear Society, Spring; 2001.
[33] US NRC. NUREG/CR-1659. Reactor safety study methodology applications program. US Nuclear Regulatory Commission, (Vol. 1) April 1981, (Vol. 2) May 1981, (Vol. 3) June 1982, (Vol. 4) November 1981.
[34] US NRC. NUREG-0660. NRC action plan developed as a result of the TMI-2 accident. US Nuclear Regulatory Commission, May, (Rev. 1); August 1980.
[35] Institute for Electrical and Electronics Engineers. NUREG/CR-2300: PRA procedures guide: a guide to the performance of probabilistic risk assessments for nuclear power plants; 1983.
[36] Meserve R. The evolution of safety goals and their connection to safety culture, speech delivered at: ANS topical meeting on safety goals and safety culture, Milwaukee, WI; June, 2001.
[37] US NRC. SECY-89-102. Implementation of the safety goals; 1990.
[38] Electric Power Research Institute. PSA applications guide, EPRI TR-05396; August 1995.
[39] Ahearne J, et al. The regulatory process for the nuclear power reactors: a review. A report of the CSIS nuclear regulatory process review steering committee; 2001.
[40] US NRC Staff Briefing On Reactor Inspection. Enforcement, and assessment; January 1999.
[41] US NRC. Regulatory Guide 1.174: An approach for using probabilistic risk assessment in risk-informed decisions on plant-specific changes to the licensing basis. (Revision 1 issued November 2002); 1998.
[42] Sorensen J, Apostolakis G, Kress T, Powers D. On the role of defense-in-depth in risk-informed regulation. Proceedings of PSA'99, Washington, DC. la Grange Park, IL: American Nuclear Society; 1999.
[43] Meserve R. The evolution of safety goals and their connection to safety culture. Speech delivered at the American Nuclear Society topical Meeting on safety goals and safety culture, Milwaukee, WI; June, 2001.
[44] Personal communication with Dr Harold Lewis; April 2004.