Available online at www.sciencedirect.com
SCIENCE DIRECT®
Reliability Engineering and System Safety 86 (2004) 269–284
RELIABILITY ENGINEERING & SYSTEM SAFETY
www.elsevier.com/locate/ress

Dynamic modeling of the tradeoff between productivity and safety in critical engineering systems

Michelle M. Cowing[a,*], M. Elisabeth Paté-Cornell[b], Peter W. Glynn[b]

[a]*Business and Economics Department, Pacific University, 2043 College Way, Forest Grove, CA 97116, USA*
[b]*Department of Management Science and Engineering, Stanford University, Stanford, CA 94305-4026, USA*

Received 21 June 2003; accepted 4 February 2004

Abstract

Short-term tradeoffs between productivity and safety often exist in the operation of critical facilities such as nuclear power plants, offshore oil platforms, or simply individual cars. For example, interruption of operations for maintenance on demand can decrease short-term productivity but may be needed to ensure safety. Operations are interrupted for several reasons: scheduled maintenance, maintenance on demand, response to warnings, subsystem failure, or a catastrophic accident. The choice of operational procedures (e.g. timing and extent of scheduled maintenance) generally affects the probabilities of both production interruptions and catastrophic failures. In this paper, we present and illustrate a dynamic probabilistic model designed to describe the long-term evolution of such a system through the different phases of operation, shutdown, and possibly accident. The model's parameters represent explicitly the effects of different components' performance on the system's safety and reliability through an engineering probabilistic risk assessment (PRA). In addition to PRA, a Markov model is used to track the evolution of the system and its components through different performance phases. The model parameters are then linked to different operations strategies, to allow computation of the effects of each management strategy on the system's long-term productivity and safety. Decision analysis is then used to support the management of the short-term trade-offs between productivity and safety in order to maximize long-term performance. The value function is that of plant managers, within the constraints set by local utility commissions and national (e.g. energy) agencies. This model is illustrated by the case of outages (planned and unplanned) in nuclear power plants to show how it can be used to guide policy decisions regarding outage frequency and plant lifetime, and more specifically, the choice of a reactor tripping policy as a function of the state of the emergency core cooling subsystem.

Keywords: Risk management; Safety; Production; Trade-off; Maintenance; Critical systems; Nuclear reactors; Dynamic modeling; Probabilistic risk assessment

1. Managing the short-term productivity/safety trade-off

Maintaining high levels of productivity and safety is a primary objective for many industries, especially those that operate complex, hazardous production systems such as nuclear power plants, offshore oil platforms, chemical plants, and space transport systems. In the long term, accident avoidance and high productivity are tightly linked. In the short term, however, these systems may operate under tight resource constraints and there are often tradeoffs between immediate productivity and safety [1]. A number of accidents, for example, the loss of the space shuttle Challenger and of the Piper Alpha oil platform, have occurred because upgrades and/or maintenance operations were delayed in order to meet production goals or deadlines [2–4].

Safety management strategies for critical systems involve multiple dimensions including design philosophy, maintenance policies, and procedures of personnel hiring, training, and evaluation. At one end of the spectrum, the most conservative approaches rely on a robust system design, frequent preventive maintenance, and early response to warnings. At the other end, aggressive strategies are driven by demanding production schedules, single-string system designs, and minimal inspection and maintenance to obtain maximum production with minimum interruptions. The difference, of course, lies in the immediate costs and in the resulting level of system failure risk. Further discussion of comprehensive risk management strategies in critical

* Corresponding author. Tel.: +1-503-352-2184; fax: +1-503-352-3195.
E-mail address: cowing@pacificu.edu (M.M. Cowing).

0951-8320/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ress.2004.02.003

Nomenclature

M	decision variable cycle length or time between planned maintenance periods
L	decision variable number of cycles in the lifetime of the system
X	$(X_n : 0 \leq n \leq M)$, discrete time Markov chain with time index n
i	index of the cycle of operation
$\{S\}$	Markov chain state space, as defined in Section 3, $\{S\} = \{x_1, x_2, \ldots, x_{k_1+k_2}\}$
$\{S_1\}$	subset of overall state space, planned maintenance states, $\{S_1\} = \{x_1, x_2, \ldots, x_{k_1}\}$
$P_{11}(i)$	corresponding sub-matrix of transition probabilities for maintenance states during the ith cycle of operation
$\{S_2\}$	planned operating state space subset, $\{S_2\} = \{x_{k_1+1}, x_{k_1+2}, \ldots, x_{k_1+k_2}\}$
$P_{22}(i)$	corresponding sub-matrix of transition probabilities for planned operating states during the ith cycle of operation
$P_{12}(i)$	matrix defining transitions from S_1 states to S_2 states during the ith cycle of operation
$\Pi(i)$	one-step partitioned transition matrix, operating cycle i
$\Pi_{x,y}(i)$	probability of transition from state x to state y during cycle i, element (x, y) of $\Pi(i)$
$\mu_i(n)$	state distribution vector at time n within cycle i, $\mu_i(n) = [P(X_n = x_1), P(X_n = x_2), \ldots, P(X_n = x_{k_1+k_2})]$
f	state occupancy cost per unit time, column vector
$f(x)$	element of the f vector, rate at which cost accrues in state x per time unit
g	state transition cost vector, column vector
$g(x)$	element of the g vector, cost incurred for each visit to state x
$C_i(\mu_i(0), M)$	accumulated monetary value (costs or benefits) during cycle i of duration M with initial state distribution $\mu_i(0)$.

systems and their effect on system productivity and safety is presented elsewhere [5].

Production interruptions can be caused by scheduled inspection and maintenance, or by unplanned circumstances including warnings of transients, production failures, and accidents. Maintenance operations can include inspection, parts replacement, repair, and refueling or restocking. Many maintenance tasks can be performed while the system remains in operation; but major scheduled shutdowns are needed in critical systems and have a significant effect on system productivity, cost, and safety. Unplanned production interruptions may also be costly, but they are generally benign. Accidents, by contrast, can result in catastrophic and unrecoverable failures with large financial, human, and environmental costs. Therefore, the choice of an appropriate risk management strategy depends on the probabilities and consequences of production failures and of accidents of variable levels of severity under different options. It also depends on the preferences and risk attitude of the decision maker.

There are two basic approaches to the mathematical modeling of safety and productivity. For systems with sufficient historical data about production interruptions, one can use a global empirical method to obtain a measure of the system's failure risk based on a statistical analysis of production failures and on a set of hypotheses about the link between production interrupters and the likelihood of catastrophic failures [6,7]. When this information is not available, another approach is to rely on systems' analysis and probabilistic risk assessment (PRA) to compute the probabilities, the effects, and the costs of both production interruptions and catastrophic failures. For critical systems (such as nuclear reactors), the empirical approach may provide sufficient data to analyze production interruptions because these occur relatively frequently. However, empirical *safety* assessments based on production measures rely on the assumption that production failures are highly correlated with catastrophic failures. This is not necessarily true. The correlation depends on the overlap between events that may lead only to production interruptions, and those that may also lead to accidents. In reality, these two types of events are generally distinct. Therefore, we use a systems' analysis approach and probabilities to model productivity and safety over the life of a critical facility.

One problem of short-term operations policies and production decisions is that they are often myopic. For example an individual operator or manager might have different or limited concerns and objectives than those best overall for the organization (e.g. a short-term departmental production or cost focus rather than a long-term, organization-wide focus). Therefore, there may be a 'principal-agent problem' associated with the discrepancies of preferences (and incentives) between the different parts of the organization, and perhaps, the public interest in cases such as nuclear power plants. Thus, there is a need for an overarching organizational objective function by which alternative operating and maintenance policies are evaluated. In this paper, we do not attempt to analyze the effects of the behaviors of operators or managers who may find it convenient to bypass the procedures required by the organization. Instead we try to provide managers with a decision support system that accounts for these possible discrepancies.

We thus assume that the long-term preferences of the firm's top management prevail and are effectively implemented in operations decisions. A decision support system that reflects these preferences is then an asset because it can be used across the board to reflect consistently the preferences of the firm's highest level of

management. Therefore, each decision is made in a larger context, and not separately on the basis of variable objectives that may depend on the time horizon of individual operators and managers.

The question is then: whose value function should be used here? In a nuclear power plant, for example, many parties are involved in informing or making daily and long-term decisions, e.g. plant management, regulatory agencies, an independent safety committee, the world experience evaluation group, and various internal groups as identified by Vaurio [8]. Industries that manage critical systems plants and for which the safety of operations is paramount to operating decisions are generally regulated by federal, state, and local agencies (e.g. the US Nuclear Regulatory Commission or the Federal Aviation Administration). This is to ensure that the profitability requirements of utility or airline companies, for instance, do not jeopardize system safety. These firms must thus first satisfy these rules and regulations, then manage their operations in a prudent and cost-conscious way. Therefore, in what follows, we assume that the interests of the public are protected by these regulations and guidelines and that there is no attempt on the part of industry to take short cuts with respect to safety rules or a mandated price structure. The question here is thus to satisfy an objective function that reflects the long-term interests of the firm (which of course, are very dependent on safe and reliable operations) within the requirements of the different elected or appointed bodies.

The purpose of this paper is to illustrate the use of mathematical and engineering models to support consistent choices of general operations policies and of short-term management options, based on a long-term assessment of their effects on the system's productivity and safety, and a long-term vision of the consequences of immediate decisions. Our analytical framework combines PRA, stochastic systems, decisions and cash flows. We include these in a dynamic model of the evolution of a safety-critical production system through states characterized by different levels of safety, productivity, and cost. System performance attributes that are the output of our dynamic model (including production revenue, costs, human casualties, and environmental impacts) are then variables of an objective value function specified by the decision maker.

Related approaches to the dynamic modeling of component or system reliability, productivity, or safety have been proposed in the literature [9–13]. For example, Vesely [9] used a Markov model to quantify maintenance effects on component availability and risk. Similarly, we model the effects of maintenance costs and benefits on component performance, but with the purpose of measuring availability and risk at the system level in order to support a range of risk management decisions. Furthermore, Vesley's optimum is based on availability alone, whereas our optimal policy is based on a balance between measures of failure risk and productivity. Martorell et al. [10] describe an approach that considers the integrated effect of various surveillance and maintenance tasks on specific critical components and seek to explore impacts on both risk and cost. However, their approach is not dynamic, and they do not explicitly attempt to resolve the cost versus risk tradeoff.

We illustrate our model by the case of the management of outages in a nuclear power reactor system to link measures of short-term and long-term productivity and safety to policy decision variables (i.e. the choice of an interval between maintenance operations and of the plant's lifetime). We then illustrate the specific use of the model to determine a decision rule for reactor trips caused by failures of the emergency core cooling system (ECCS).

2. Failure types and initiating events

We define two types of system failures: accidents (A), and unplanned shutdowns or breakdowns (B). Whereas accidents affect measures of both productivity and safety, unplanned shutdowns only affect system productivity. In a nuclear reactor, accidents are the rare events that can result in core damage and the release of radioactive material inside or outside the containment building. Unplanned shutdowns (such as trips or SCRAMS), result in production failures. They may reflect the system's physical inability to operate, for example because of the failure of a turbine generator, or a technical decision to interrupt production in order to repair a deteriorating component or a failed standby safety feature. Such a decision may thus be the result of a prudent response to a system challenge. While potentially costly, unplanned shutdowns do not generally threaten health and safety. Yet, they may lead to an accident if the shutdown itself involves safety issues. For example, in the case of nuclear power plants, the rapid shutdown process and the restarting of a reactor after a SCRAM may cause transient events that can be accident initiators. Our probabilistic model thus involves explicitly accidents, breakdowns and shutdowns.

2.1. Link between accidents and breakdowns

In some cases, the link between accidents and breakdowns is clear. A single event may cause a production interruption and increase the risk of an accident. For example, in a nuclear power plant, a steam leak in the piping between the steam generator and the turbine renders the production system inoperable. The same event can also be an initiator of an accident sequence. Not all production interruptions, however, result in an increase of accident risk: well-performed maintenance on demand requires production interruptions and improves safety. Conservative operating policies in which operations are interrupted for minor abnormal events also result in additional unplanned shutdowns. They generally, but not always, lead to a higher level of safety.

2.2. Failure modeling

PRA often starts with the identification of a set of initiating events (IEs) of accident sequences [14]. They are structured into an exhaustive set of mutually exclusive events. The occurrence of an initiating event does not necessarily imply that an accident (A) will occur, and it may or may not cause an interruption in production (B). To model all relevant states of operation, we expand the traditional set of initiating events to include not only potential accident initiators, but also events that may lead to production failures only. Our set of initiating events thus consists of two subsets: accident initiators (IE_A) and production failure events (IE_B). For example, in a car, accident initiators include a blown tire or failure of the breaks, and shutdown initiators include engine failure and running out of gas. By definition, IE_A events may lead to accidents (A) or to unplanned system downtime (B), while IE_B events generally only lead to unplanned downtime (B).[1]

3. System states and model structure

The initiating events of accident sequences and breakdowns are included here in a dynamic model based on transitions among operating and failure states over the life of the system. The modeling is performed at the system level, but the states and state transitions are defined in terms of sub-system and component failure rates.

3.1. Model assumptions

The lifetime of the system is divided into distinct cycles of operation defined as the time elapsed between successive starts of planned maintenance periods. The model requires that each cycle length be specified according to the chosen maintenance policy. The duration of these cycles is assumed to be constant throughout the life of the plant. This assumption can be relaxed if needed to allow for planned cycles of different duration.

The division of a system's lifetime into cycles of operation allows representation of the correction, at the time of planned maintenance, of the short-term deterioration that occurred during the previous cycle. The system state at the end of a cycle thus depends on the choice of the cycle length. If a long time elapses between planned maintenance operations, we assume that more extensive system wear and deterioration occurs before adjustment. This is represented, in our model, by higher probabilities of transition to different failure states within a cycle of operation. The probabilities of transition to failure states are thus functions of management policies such as the choices of maintenance cycle length, and of the total number of cycles in the system's lifetime (i.e. the system's age at decommissioning time).

Each cycle of operation consists of a planned maintenance phase (PM) followed by a period of planned operation (Op). Within each of these two modes of operation, the system may evolve through a number of distinct states, such as unplanned outages (UPO) during the planned operating phase. Let L be the life of the system, and $M_i = M$ be the maintenance cycle length (assumed equal over time) for each cycle i. Over the total time horizon there are L/M distinct cycles of operation. While the length of each cycle is deterministic (M time units), the amount of time spent in the planned maintenance phase and in planned operation within each cycle is the result of a stochastic analysis of transitions in which the first k time units are spent in maintenance states, and the remaining time in planned operating states.[2] Each cycle of operation is modeled as a Markov process.[3] A discussion of the strengths and weaknesses of alternative modeling approaches for risk assessment in dynamic systems is presented by Siu [11].

Within each cycle, we assume that the transition probabilities per time unit are constant, but that they are a function of chosen management policies such as cycle length.[4] In addition, transition probabilities may vary from one cycle to the next to account for effects such as long-term system deterioration. These cycles are then linked to determine life-time performance measures for the system.

3.2. Model structure

During each cycle of operation, we model separately planned maintenance (typically for inspection and preventive maintenance activities) and planned operation. The overall system state space $\{S\}$ is thus divided into two subsets: $\{S_1\}$, planned maintenance states and $\{S_2\}$, planned operations states. While in a state of planned maintenance, the model tracks transitions through different states of the shutdown sequence, including compromised system configurations during the shutdown process, stable

[1] In rare circumstances an event classified as a breakdown initiator may result in an accident (e.g. engine failure on a busy freeway). Events that may result in either breakdowns or in rare instances accidents, should either be classified as an accident initiating event or the event rate can be split into two initiators, one for accidents and one for breakdowns.

[2] If the duration of the planned maintenance state within a cycle of operation is known with certainty or if the need to model transitions among distinct maintenance states is unnecessary, the approach can be simplified to skip the stochastic analysis of PM transitions, and simply add fixed times and associated costs between our linked cycles of operation to obtain an appropriate lifetime objective function.

[3] Markov modeling is a standard technique for the mathematical representation of dynamic systems. It easily handles the time dependent nature of the system and allows for the explicit specification of unique system configurations. We discuss some of its specific limitations in Conclusions.

[4] In the application to nuclear reactors, the maintenance cycle is specified in quarters (3-month intervals). State transitions are defined in terms of a finer time scale, per week in the example. For a car system, the time unit might be defined in terms of weeks or months.

shutdown configurations, and possible accidents. While in most systems, shutdown states preclude the possibility of an accident, a system such as nuclear power plant (or a car immobilized on a freeway) is at risk of an accident even during system shutdown. For generality, we thus assume that accidents can occur during planned shutdowns. During operations, the Markov model tracks transitions through periods of system operation under normal and compromised conditions and into unplanned shutdowns and accidents.

For each cycle of operation, the overall transition matrix can be represented as a partitioned matrix composed of the transition matrices for the state space subsets $\{S_1\}$ and $\{S_2\}$. The overall transition matrix for cycle i is then of the following form:

$$\Pi(i) = \begin{array}{c} \\ S_1 \\ S_2 \end{array} \begin{array}{c} \begin{array}{cc} S_1 & S_2 \end{array} \\ \begin{bmatrix} P_{11} & P_{12} \\ 0 & P_{22} \end{bmatrix} \end{array}. \quad (1)$$

P_{11} and P_{22} sub-matrices are the matrices of transition probabilities among states of planned maintenance and planned operation respectively. P_{11} and P_{22} are thus square matrices of dimension k_1 and k_2, respectively, where k_1 represents the number of distinct states in $\{S_1\}$ and k_2 in $\{S_2\}$. P_{12} is therefore an k_1 by k_2 matrix that accounts for transitions from planned maintenance to operation. Transitions in the lower left part of $\Pi(i)$ are not allowed during a single cycle of operation. They would represent transitions from operating states to planned shutdown states, which by construction, occur in subsequent cycles of operation. The duration M of each operating cycle is fixed and specified as part of the risk management strategy. Therefore, the transition matrix $\Pi(i)$ is valid for the M time units of cycle i. Each cycle begins with transitions among maintenance states as defined by the P_{11} sub-matrix. After some k (random) periods, transition occurs from a maintenance state to an operating state (a P_{12} transition). The remaining $M-1-k$ transitions occur among the planned operating states, defined by the P_{22} sub-matrix.

3.3. State space

Within each of the state space subsets $\{S_1\}$ (planned shutdown) and $\{S_2\}$ (planned operations), we distinguish unique system shutdown and operating states to a level of detail necessary to capture differences in productivity or safety measures. We group these states into the high-level sets listed below. While the actual state space is unique to each system, the types of states required for modeling purposes is general.

$\{S_1\}$: Planned Maintenance States
- Set of Planned Shutdown and Maintenance States
- Set of Accident States possible during Planned Maintenance

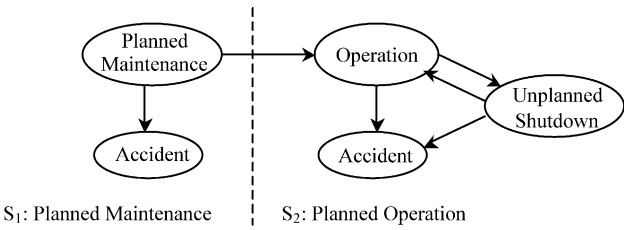

Fig. 1. Basic system states types and transitions for a cycle of operation.

$\{S_2\}$: Planned Operation States
- Set of Operating States
- Set of Unplanned Shutdown States
- Set of Accident States possible during Planned Operation

While the state space is representative of the system's evolution over its lifetime, the probabilities within the transition matrix may change from one cycle to the next because the terms of this matrix are a function of both management policies and system age. Fig. 1 indicates the types of high-level state and the transitions that may occur during a particular cycle of operation. Because we focus here on critical systems, we model accidents as absorbing states. Recoverable accident states are modeled as unplanned shutdowns with relatively low exiting transition rates and high costs.

In practice, the development of a PRA model for the system is thus an important part of our approach because it allows us to identify unique states of system operation, and to represent unique failure probabilities and rates of deterioration at the component or subsystem level where the data are often more readily available.

3.4. External events

External events, such as storms and earthquakes, affect system safety to the degree that they affect the probability of component or subsystem failures, or the ability of the system to respond to initiating events. These external events increase the probabilities of accident initiators and/or decrease the system's ability to respond to the challenge. Therefore, they influence the probabilities of failure scenarios and their consequences given the occurrence of an initiating event. Because of their potential effects on system performance, external events are part of our dynamic model. External events are treated explicitly in the PRA method. For example, see Mosleh [15] and Budnitz [16].

In our model, the probabilities of initiating events and of subsequent evolution paths that constitute the transition matrix are conditioned by external events, and the state space reflects the system's evolution with and without the occurrence of a particular external event. This is important because the choice of a risk management strategy may influence the effect of an external event on the system.

Managers can set policies that limit operations during (or in anticipation of) certain types of external events. For example, production on an offshore oil platform may be limited in anticipation of a severe storm, and a nuclear power plant can automatically shut down during an earthquake. Such policies may reduce the consequences of an initiating event, but they also reduce temporarily the system's productivity.

3.5. Output, costs, risks and value model

Production and risk patterns are tracked by our dynamic stochastic model as the system evolves through the state space over time. We model each of these factors through the succession of cycles of operation under a Markov assumption. We then link successive Markov cycles to compute the different performance measures (e.g. through expected values) over the life of the system. The model computes the expected time spent in each state and the expected number of transitions to each state. Output measures from the model are then used as inputs to a decision model to provide support for the design of a risk management strategy. The decision model includes an objective function that aggregates the values attributed by the decision maker to productivity and safety measures.

Our approach is to encode and to use the value function of an actual and specific decision maker. We assume that the organization can identify this actual decision maker (e.g. the plant manager of power generation at a nuclear power station), and that the decision maker and the organization will benefit from consistency of goals and objectives. Therefore, this approach requires that the organization (via its highest levels of management, informed by various internal and external parties) explicitly consider the relative values of short-term and long-term costs, risks, production levels, etc.

Each defined system state is characterized by two value components: one based on time spent in the state, and the second based on transitions to the state. These values are an aggregation of both financial measures (i.e. costs and revenues) and of non-financial attributes such as safety effects and aversion to certain high-risk states. Values may be a function of the cycle number (or other attributes) to account for such factors as long-term system wear. Productivity and safety measures that are a function of either the time in-or transitions to-a particular state are modeled in the objective function. For example, to determine the productivity of a nuclear reactor, we keep track of the length of planned shutdown periods as well as the amount of time spent in non-producing states of planned operation.

We characterize these values by 'costs' that represent actual or hypothetical monetary values (costs or benefits) of the following performance measures: (1) revenues from operations, (2) fixed and variable operations and maintenance costs, and (3) penalties for downtime, near-misses, and transitions to high-risk states. These penalties represent a 'willingness-to-pay' to avoid these system states, and therefore, the utility, preferences and risk attitude of top management. We also assume that they satisfy first the existing laws, regulations and decrees of the relevant jurisdiction. The optimum strategy is the solution that corresponds to a net maximization of benefits (or to a cost minimization). Value tradeoffs among monetary and non-monetary attributes are handled through the process of assigning these monetary values. An alternative would have been to use explicitly a multi-attribute utility model [17]. We chose to use cost values instead because they are easier to encode and to understand intuitively. This approach, however, assumes that the different cost components are additive and independent of each other, which may not always be true (e.g. willingness to pay to avoid problems may depend on the revenues so far). In either case, an aggregation of the different attributes must ultimately be performed, a difficult and unavoidable process which, in our approach, requires explicit consideration by managers of the classic tradeoffs between human safety and monetary measures [18] as opposed to implicit tradeoffs through intuitive choices.[5]

4. Dynamic modeling of one cycle of operation

The computations require first an estimation of the system's performance over time. The resulting objective function during a particular cycle of operation involves the transitions among states, the time spent in the different states, and the values assigned to corresponding levels of productivity and safety in the objective function.

4.1. Objective function

The objective is to maximize the expected equivalent monetary value associated with operations over the life of the system. Because we use a discrete Markov model, the expected monetary value is the product of the cost components and the probability of being in each state at each time period. The overall objective function is the sum (or discounted sum) of expected monetary values for each of the distinct operating cycles:

$$C_{\text{total}} = \sum_{i=1}^{\#\text{ of cycles}} C_i(\mu_i(0), M). \tag{2}$$

[5] It is assumed in this paper that the system operates within a range of safety that justifies cost-benefit analysis [38]. This implies that the probability of severe accident is low enough for the individual risk to be acceptable in the first place both on-site and off-site. Otherwise, the system should simply not operate until it has reached first that minimum safety level on the basis of the threshold and the time horizon (e.g. annual probability) specified by the regulatory authorities.

The value associated with operation and shutdown over M time periods in the ith cycle of operation consists of two monetary components.[6] C_{if} is the cost of *spending time* in the various states, and C_{ig} is the expected cost of *transitions* to the different states. The total cycle cost is the sum of these cost elements:

$$C_{\text{cycle}} = C_i(\mu_i(0), M) = C_{\text{if}} + C_{\text{ig}}. \tag{3}$$

4.2. Cost computation for a single cycle of operation

We compute first the cost C_{if} of occupying specified states. It is computed as the conditional expectation:

$$C_{\text{if}} = E_{\mu_i(0)}\left[\sum_{n=0}^{M-1} f(X_n)\right]. \tag{4}$$

The expected cost of spending time in the various states during each time period is given by the vector product of the state distribution and the f-cost vector. The cumulative expected cost associated with time spent in different states during the M time periods of cycle i can be computed recursively, updating the cumulative expected cost function as we update the state distribution at time n, μ_n. Because the Markov assumption holds within each cycle of operation, the state distribution at time n is:

$$\mu_i(n) = \mu_i(n-1)\cdot\Pi(i) \tag{5}$$

The cumulative expected state occupancy cost at time zero, with initial cycle state distribution $\mu_i(0)$, denoted as $C_{\text{if}}(\mu_i(0), 0)$, is zero, because no costs have been incurred. With the passing of each successive time period, one incurs the expected cost associated with the state distribution during each time period. The cumulative cost function is therefore computed recursively as:

$$C_{\text{if}}(\mu_i(0), n) = C_{\text{if}}(\mu_i(0), n-1) + \mu_i(n-1)\cdot f \tag{6}$$

Next we compute the expected cost of state transitions, C_{ig}, also computed as a conditional expectation of the Markov Chain, where expected costs are incurred at the time of transitions. $I(X_n \neq X_{n-1})$ is an indicator function that is equal to one if an actual transition to a different state occurs (i.e. if $X_n \neq X_{n-1}$) and zero otherwise. Thus:

$$C_{\text{ig}} = E_{\mu_i(0)}\left[\sum_{n=1}^{M} g(X_n)\cdot I(X_n \neq X_{n-1})\right]. \tag{7}$$

The expected cost of transitions at the end of each time period is given by the vector product of the expected number of transitions to each state and the g-cost vector. Let the expected number of transitions to a particular state x during a single time period n be $\alpha_n[x]$. The expected number of transitions to each state during a particular time period can be represented by a modified state distribution vector, in which a transition is recorded only if the new state is different from the previous one:

$$\alpha_n[x] = \mu_i(n)[x]\cdot I(X_n \neq X_{n-1}) \tag{8}$$

To represent the indicator function I, we define a new matrix $H(i)$ with the same dimension as the one-step transition matrix $\Pi(i)$ and containing the same elements except for the diagonal elements which are zero so as to account only for actual transitions. The elements of the H matrix are thus:

$$h_{x,y}(i) = \begin{cases} \pi_{x,y}(i) & x \neq y \\ 0 & x = y \end{cases} \tag{9}$$

The vector of expected transitions for one time period α_n in cycle i is then computed as:

$$\alpha_{n(i)} = \mu_i(n-1)\cdot H(i) \tag{10}$$

The cumulative transition cost at time zero, with initial cycle state distribution $\mu_i(0)$, is denoted as $C_{\text{ig}}(\mu_i(0), 0)$. It is equal to zero because no transition costs have been incurred. Thereafter, transition costs are incurred during each new time period. The cumulative transition cost function is computed recursively:

$$C_{\text{ig}}(\mu_i(0), n) = C_{\text{ig}}(\mu_i(0), n-1) + \mu_i(n-1)\cdot H(i)\cdot g \tag{11}$$

The combined conditional expectation for cost in a single cycle i is therefore:

$$C_i(\mu_i(0), M) = E_{\mu_i(0)}\left[\sum_{n=0}^{M-1} f(X_n) + \sum_{n=1}^{M} g(X_n)\cdot I(X_n \neq X_{n-1})\right]. \tag{12}$$

5. Successive cycles of operation

The model described in Section 4 is used to compute the net expected value of operation and maintenance over a single cycle of operation. In this section, we discuss the differences among cycles of operation and we describe the computations necessary for linking successive cycles to determine the total expected cost over the life of the system.

5.1. Linking operating cycles

Successive cycles of operation are linked because the final state distribution of one cycle (after M time units) affects the initial distribution of the subsequent cycle, and because deterioration that occurs in one cycle is present in the next one. The initial distribution for a new cycle i is determined as a function of the ending distribution of the previous cycle using a matrix T which represents the probabilities of all possible transitions from cycle-end states to next-cycle planned maintenance states. T is a square matrix of order $k_1 + k_2$ corresponding to $\{S_1\}$ and $\{S_2\}$.

[6] Again, in all that follows the term 'costs' is used to represent monetary values that may be positive or negative to reflect revenue, actual costs, and penalties.

Because, by construction, each cycle begins in a planned maintenance state, the transformation matrix T is sparse. Transitions may only be to a subset of the planned shutdown states of $\{S_1\}$. The initial distribution for each new cycle i is therefore determined probabilistically as:

$$\mu_i(0) = \mu_{i-1}(M) \cdot T \tag{13}$$

5.2. Modeling deterioration effects

Within each cycle of operation, we assume that the transition probabilities are constant and represent mostly system deterioration. During one cycle, the extent of the eventual deterioration is determined, in part, by management policies such as the cycle length: less frequent maintenance leads to more severe component deterioration before the subsystem is fixed. For example, planned maintenance on a 2-year instead of a 1-year schedule yields higher cycle failure probabilities and therefore higher transition probabilities. We also assume that equipment failure probabilities generally increase with successive cycles of operation as the system ages (i.e. for higher cycle numbers). Therefore, the deterioration at the end of the system's life also depends on the policy decision of how long the system is or will be kept in operation. Defining base transition probabilities for a given cycle as a function of the *cycle length* and the *cycle number* allows us to capture both the short-term and the long-term effects of these two management decisions on the system's performance over time during its total lifetime.

We model equipment failure probabilities as an increasing function in the cycle number (i) in the life of the plant (i.e. system age) and in maintenance cycle length (M). A transition probability can be any function of i and M that is consistent with previous operating experience or expert opinion. In general, transition probabilities may be monotonic or not, and increasing or decreasing in time. Individual failure probabilities can change over time according to system and component characteristics. For example, we model our transition probabilities to unplanned outage states as an increasing function in the number of previous unplanned outages to account for the cumulative deterioration effects of such abrupt and stressful shutdowns.

This formulation assumes that the deterioration depends only on the system's age from the start ($i = 0$) and is independent of specific events in previous cycles. If appropriate, one can capture relevant mid-life events such as the complete replacement of a deteriorated subsystem by resetting the age index. This simply requires the definition of specific age indices for the subsystems that are likely to sustain periodic but infrequent replacements. These indices can be reset upon entry into unique states of planned maintenance that occur in rare instances (e.g. every 10 or 20 years) to account for major overhauls and replacements.

Another potential effect of system deterioration is an increased probability of both planned and unplanned shutdown complications resulting in longer shutdown durations. Assuming that such deterioration effects increase with both i and M, the transition probabilities out of planned and unplanned outage states may decrease with i and M. The result is an increase in expected planned and unplanned shutdown durations with system age and with longer intervals between planned maintenance. In our application of the model, we use exponential functions of i and M to reflect both effects. Additional issues in the choice of deterioration functions are discussed by Bier [19]. Martorell et al. [20] describe their approach for modeling deterioration in NPP safety components while considering the effects of both maintenance activities and working conditions.

5.3. Discounting and non-linear costs

The objective function is generally non-linear with respect to time and number of transitions. In the model, we take into account two important non-linear effects: the time value of money (i.e. the opportunity cost of capital), and increase in fixed costs associated with successive transitions to the different states, and in particular, transitions to unplanned shutdown states. We discount the costs on a per-cycle basis, assuming end-of-cycle lump sums.[7] The social rate of discount that we use is the result of an equilibrium between the marginal rate of transformation (that reflects the ability of the economy to generate capital over time) and the marginal rate of substitution that reflects the desire of individuals that constitute society to consume now rather than later. Over the long run, and in constant monetary units, such a rate can be set in the order of 3–3.5% [21]. We also make the classic assumption of a common discount rate for dollars and human safety. The core of the argument is based on equity, with the objective to provide all generations with the same amount of safety measures over time in a situation where wealth accumulates over years. Further discussion of this classic choice is provided in Paté-Cornell [21]. While discounting at the same rate all attributes of the cost function (including the number of lives saved by various safety measures) seems to imply that 'future lives are cheaper than present ones', the opposite is true: if all generations are to be equally protected, the same amount of life saving technology should be made available to all individuals at all times. By virtue of the geometric growth of the money accumulated by society at the social rate of discount, this implies that the ratio of number of lives saved to safety investments should remain the same (in constant

[7] This is a reasonable approximation to the extent that the cycle length is in the order of 1 year or less.

monetary units). It should be noted that this result is based on the conservative hypothesis that life saving technologies will remain at the same cost level in the future (in constant monetary units). In reality, one might expect technological progress to reduce these costs, and the rate of discount of the safety attribute could be increased to reflect this rate of progress. Therefore, in what follows, we assume a single social rate of discount of 3%, free of risk and free of inflation, applied to all aspects of the failure costs integrated into a single figure. Again, the risk is included in the probability of the failure scenarios and all computations are in constant monetary units.

Let t be the cycle number in the overall system life and i_M the chosen per-cycle discount rate. The present value (CycleCost_0) of the cost incurred during the tth cycle of operation is:

$$\text{CycleCost}_0 = \frac{\text{CycleCost}_t}{(1+i_M)^t}. \tag{14}$$

For costs that increase non-linearly with the number of transitions to a particular state, we modify the corresponding elements in the cost vectors. For example, because unplanned shutdowns often cause severe stresses in the system that may have a cumulative deterioration effect, we model the cost of successive transitions to such states as an increasing function of the transition number.

For example, let $g(x_i)$ be a cost element of the g (transition) cost vector associated with a transition to state x_i, an unplanned shutdown state. If the cost of the first transition to state x_i, $g_1(x_i)$, is C_{base}, successive transition charges can be assess as a function of the number of previous transitions to an unplanned shutdown state. Let k be a cost rate modifier factor. One can use, for example, a simple linear model to represent the increasing cost of a transition to state x_i as a function of the transition number j :

$$g_j(x_i) = -[C_{\text{base}} + k(j-1)]. \tag{15}$$

6. Computational improvement

Both cost terms C_{if} and C_{ig} can be computed more efficiently by taking advantage of the partitioning of the state space into the two sets $\{S_1\}$ (planned maintenance) and $\{S_2\}$ (planned operation). We use the same recursive relations as in Section 4.2. Cost quantities, however, are computed separately for each mode of operation. We represent the state distribution vector at any time n during a cycle in terms of its block components: $\mu_{i1}(n)$, the distribution of planned outage or maintenance states at time n, and $\mu_{i2}(n)$, the distribution of planned operating states at time n :

$$\mu_i(n) = [\mu_{i1}(n), \mu_{i2}(n)]. \tag{16}$$

Because, by definition, each cycle begins in a planned maintenance state, the initial state distribution for each cycle i can only be non-zero in $\{S_1\}$. Therefore, we can represent the initial state distribution in block-partitioned format as:

$$\mu_i(0) = [\mu_{i1}(0), 0]. \tag{17}$$

In Section 4.2, C_{if} and C_{ig} are computed recursively as a function of the successive calculations of the state distribution, $\mu_i(n)$. In this section, we compute $\mu_i(n)$ in terms of its block-partitioned components as $\mu_i(n) = [\mu_{i1}(n), \mu_{i2}(n)]$.

Because of the structure of the $\Pi(i)$ transition matrix, the distribution of planned maintenance states in $\{S_1\}$ at time n during cycle i, $\mu_{i1}(n)$, is a function of the previous distribution of planned maintenance states and of the $P_{11}(i)$ sub-matrix. Within a cycle of operation, new transitions to an S_1 state can only occur from within the $\{S_1\}$ state space according to the $P_{11}(i)$ matrix. Therefore, we compute $\mu_{i1}(n)$ according to the following relation:

$$\mu_{i1}(n) = \mu_{i1}(n-1)\cdot P_{11}(i). \tag{18}$$

The distribution of $\{S_2\}$ (planned operating states) at time n during cycle i is a function of both the prior distribution of S_1 states and S_2 states. Changes in the S_2 state distribution may occur because a transition occurs from an S_1 state to an S_2 state during the time period as described by the elements of the $P_{12}(i)$ sub-matrix. Transitions may also occur if the system is already in an S_2 state. In that case, the system evolves according to the $P_{22}(i)$ sub-matrix that defines transitions among the S_2 states. The distribution of S_2 states is therefore computed as:

$$\mu_{i2}(n) = \mu_{i1}(n-1)P_{12}(i) + \mu_{i2}(n-1)P_{22}(i) \tag{19}$$

Objective function computations are performed as described in Section 4, using corresponding block components of the f and g cost vectors.

7. Illustration: management of outages in nuclear reactors

As an illustration of this model, consider its application to the management of outages in a nuclear power plant. In the nuclear power industry worldwide, one can observe a spectrum of safety management strategies ranging from the general conservatism of Western designs and operations to the approach sometimes used, for example in Eastern Europe, of running a plant without interruption until serious problems occur [22].

7.1. Nuclear power plant outages

Some of the inspection and maintenance activities required for the reliable operation of a nuclear power

system are performed while the reactor is producing electricity. Much of this work, however, is accomplished during outages (planned and unplanned)[8] and their management has implications for plant safety and productivity, both in the long term and in the short term [23].

The frequency and extent of planned outages are generally determined by the reactor type and regulatory requirements, and often coincide with refueling activities. Planned outage policies and execution affect long-term safety through the frequency and quality of maintenance and have an important effect on system capabilities in emergency situations. The scheduling of outage tasks also determines the configuration of available subsystems during the outage itself that, in turn, affects short-term reactor safety. Unplanned outages, in contrast, are fairly rare and can represent major disruptions of productivity. They have direct implications for safety because they may indicate degraded conditions, inadequate maintenance, poor operator practice, or legitimate challenges to the system [24].

PRA models have been used extensively in the US nuclear power industry since the Reactor Safety Study [25] to estimate reactor risks during operation and more recently, to quantify and study short-term reactor risks sustained during planned outages [26,27]. These studies, however, do not include an analysis of the system's evolution, of the long-term effects of specific outage management policies, and of the dynamics of deterioration. PRAs are typically used in the design phase, in a static mode, for licensing decisions and other regulatory purposes. As currently employed in some industries, PRA models can give an inaccurate (often on the optimistic side) description of the safety of the system because they are seldom updated over time as the system wears and as new operating and outage policies are adopted. However, in the nuclear power industry, the use of PRAs is much more extensive. In nearly all nuclear power plants, PRAs are routinely updated to reflect changing component and system characteristics. Several plants even have on-line risk monitors that are updated as systems and data change. Such 'living' PRAs help to evaluate and plan maintenance activities and to make short-term decisions about safe operations [8].

Our modeling approach thus uses PRA as one of several analytical tools and allows for the explicit evaluation of dynamic effects of particular types of management policies over time. Our approach is to evaluate alternative operations and maintenance policies proactively, *before* they have been in effect to produce the system and data changes picked up by such on-line risk monitors as described above. Rather than relying on a sequence of static PRA risk snap-shots to inform reactive maintenance policies, we model the long-term dynamics of system risk and reliability as a function of alternative policies to identify the optimum at the beginning of the system's lifecycle.

[8] In reality, reactors may also operate at a reduced power level, i.e. a partial outage, due to equipment problems or restrictions imposed by the regulators. Such conditions and their corresponding contribution to productivity can be modeled in our approach.

7.2. Outage decision making

Until relatively recently, most of the attention on the efficient management of planned outages has been focused on the minimization of downtime. The concepts of Reliability Centered Maintenance (RCM) and Reliability-Focused Maintenance (RFM) have more recently been developed to manage maintenance activities based on the 'criticality' of components. The objectives here are to improve reliability and to reduce costs [28,10]. A limited number of utility companies have sponsored their own studies on effective outage management programs to address system vulnerabilities during planned outages [27]. The US Nuclear Regulatory Commission (USNRC) provides Technical Specifications (TS) requirements for operating and maintaining safety-related systems and components. For example, TSs define Limiting Conditions for Operation (LCOs) to ensure safety during operation, Allowed Outage Times (AOTs) to limit safety-system repair time, and Surveillance Test Intervals (STIs) to specify regular test intervals. However, the rationale behind certain outage-related policies, such as AOTs and STIs, has not always been clear or consistent and some requirements may be much more conservative than others. Recent efforts based on risk- and reliability-based methods have been proposed to improve TS requirements in general [29], and STIs or AOTs in particular [10,30,31]. In terms of objectives, Yang et al. [30] and Harunuzzaman et al. [31] seek to minimize cost while holding safety at a fixed level. Martorell et al. [10] compute policy impacts on both cost and risk (i.e. unavailability), but as noted earlier, do not propose to resolve the tradeoff.

Our model provides additional guidance to improve LCOs and AOTs, as well as establishing risk-informed planned maintenance schedules, reactor trip criteria, and plant life decisions. For example, the framework can be used to evaluate the long-term effects of alternative AOT policies on system productivity and safety. One can define, for each subsystem, unique AOT policies by defining the corresponding allowable Markov state space and transition probabilities. One can then use PRA to link the performance of individual components and subsystems to the performance of the overall system. Like Čepin et al. [32], our analysis depends on the identification of several unique plant configurations with differing levels of safety (i.e. conditional probability of core damage as computed using PRA). The differences in core damage frequency (and associated accident costs) that correspond to different plant configurations are then balanced through our objective function against corresponding differences in production capabilities. While many of the nuclear studies related to

the evaluation of alternative AOTs or LCOs focus either on schedule optimization or on safety alone, our model provides an integrated approach for balancing the production and safety effects of policy options.

7.3. Nuclear system model

Our illustration focuses on the three primary subsystems of a hypothetical nuclear production system: the reactor core, the reactor cooling system, and the electricity generation system. The reactor core contains the fuel rod assembly that, through the nuclear fission process, generates heat to produce steam that drives the electricity generation system. The reactor cooling system consists of a primary cooling loop and a number of redundant and independent ECCS. The electricity generation system consists primarily of a turbine and a generator that convert heat energy from the core into electricity. Failures within any one of these subsystems may affect both system production and safety. In addition, subsystem operations rely on common support systems such as a source of electricity.

We assume that failure of the cooling system may lead to excessive core heating and to accidents, while failure of the electricity generation system creates only an interruption in planned production. (In reality, failures in this system can also create safety problems). We also assume that core damage results only from failures in the core cooling system. Furthermore, loss of off-site electrical power renders subsystems such as core cooling inoperable until power is restored or supplied by backup generators. We consider the following sets of initiating events:

Accident initiating events

IE_1: pipe breaks-leading to a potential loss of coolant accident (LOCA)
IE_2: losses of off-site power (LOOP)

Shutdown initiating events

IE_3: turbine-generator trips

We consider here only one type of accident: severe core damage. Accident Initiating Events may result in severe core damage (with major safety and productivity costs), unplanned production interruptions, or no significant consequences for productivity or safety. Shutdown Initiating Events may result only in production interruption and/or no significant consequences.

In our illustrative application, we define 41 distinct states of operation and shutdown, namely 17 planned shutdown states $\{S_1\}$ and 24 planned operating states $\{S_2\}$. The states are based on unique system conditions and configurations that affect reactor productivity or system safety. We use the failure modes identified by a reactor PRA study and the results of a reactor outage study to identify relevant system configurations, conditions, and events. For example, in $\{S_1\}$, x_1 is a specific unstable state through which the system must pass during the normal planned shutdown sequence, x_9 represent cold shutdown mode, and x_{14} characterizes a state of loss of off-site power during cold shutdown and compromised system configuration (loss of ECCS redundancy). In $\{S_2\}$, x_{18} is normal, full power operation, and x_{23} represents operation at full power with the loss of an ECCS subsystem.

The P_{11} sub-matrix is a 17×17 matrix that defines transitions among the 17 planned shutdown states defined in $\{S_1\}$. The P_{22} sub-matrix is a 24 by 24 matrix defining transitions among the planned operating states of $\{S_2\}$. The time unit is one week. State transitions can be triggered by accident sequence initiators identified in a system PRA (e.g. coolant pipe rupture), production interrupters (e.g. turbine trip), or subsystem breakdowns leading to compromised modes of operation (e.g. loss of standby emergency core cooling system). They can also be the result of external events (e.g. an earthquake). Transition probabilities are a function of system operating capabilities, and are obtained using failure rates from the plant's PRAs, modified as appropriate to reflect the impact of time and of alternative operations and maintenance policies as described in Section 5.2.

The model is run using a set of base costs (or benefits) corresponding to the time spent in each state (f cost vector) and to state transitions (g cost vector). These costs (or benefits) are monetary values that represent the preferences and risk attitude of the decision maker for both monetary and non-monetary attributes of the consequences. Cost values must be estimated for each system state. They reflect real operations costs as well as any non-direct costs and penalties that the decision maker chooses to represent his willingness to pay to avoid some system states as described in Section 3.5. For example, the values that correspond to a planned maintenance state include inspection and maintenance costs, loss of revenue or cost of replacement energy, failure costs, and penalty costs. As an illustration, in our base model, the cost of normal (cold) downtime, $f(x_9)$, is \$500,000 per week.[9] A *transition* to an accident state, $g(x_{17})$ or $g(x_{41})$, involves a large cost, \$1 billion. We also specify costs for each state of planned operation. For example, normal system operation, $f(x_{18})$, provides a net benefit of \$5 million per week. The cost of transition to an unplanned outage state is a function of the expected number of previous unplanned outages: the first one costs \$200,000 and increases thereafter with each successive transition to account for effects of system

[9] Downtime is expensive because of the loss in revenue, and because of the cost of maintenance personnel and materials and the relatively high cost of replacement energy that must be purchased when the system itself is unable to generate electricity. Our cost figure for shutdown states includes downtime costs in excess of lost revenue, and is based on the incremental increase in cost of maintenance operations and the cost of supplying replacement energy relative to normal operating costs.

deterioration. The total cost of unplanned outage *time* is a loss of net revenue plus penalty costs for occupying certain high-risk states (e.g. $200,000 per week).

7.4. Choice of maintenance cycle length and system life

Table 1 presents a summary of the effects of different combinations of maintenance cycle lengths and system lifetimes on the overall objective function for the base-case model. For a given system life and for increasing maintenance cycle lengths, there is an increasing followed by a decreasing trend in the objective function. At one extreme, an overly conservative (i.e. frequent) maintenance policy is suboptimal because it severely constrains revenues from operations. At the other extreme, infrequent maintenance leads to greater system deterioration and therefore to a larger number of transitions to unplanned outages (reflected in higher transition probabilities in each cycle). This limits production and reduces safety in the long-term. For each cycle length and for increasing system lifetimes, there is also an increasing followed by a decreasing trend in the value function as very long lives eventually result in many production interruptions and safety problems. Furthermore, as a system ages, maintenance costs typically increase, particularly when essential subsystems (e.g. steam generators) must be replaced. These non-routine costs can be incorporated in the model, for example, by defining an exceptional state of costly planned maintenance, which is entered in rare instances (e.g. every 20 years on average) and for which the time index is reset to zero upon exit.

Fig. 2 illustrates the effect of maintenance cycle and system life (a) on the average percentage of operating time spent in a shutdown state (including both planned and unplanned outages), and (b) on the probability of severe core damage. The objective function includes both productivity and safety and depends on the relative costs assigned by the decision maker to the various states and transitions. The optimal maintenance cycle and system life are the ones that best balance the different costs (and probabilities) associated with system operation, planned and unplanned shutdown, and possible accidents. The effect of maintenance cycle and system life on the overall objective function for our illustrative nuclear production system is shown in Fig. 2(c).

7.4.1. Optimal policy

In our illustrative case, and for the considered decision maker, the preferred policy is a 100-year life with planned maintenance every 2.75 years. If the expected life of the plant is set at 40 years, as is the case for most US reactors, the optimal maintenance interval is 3 years. In our illustration, the optimum maintenance cycle is relatively insensitive to a range of system life-times. This is a consequence of our assumptions regarding deterioration. It is not a general result.

The optimal policy is generally not the policy that yields the minimum system risk or the policy that minimizes system downtime. It depends on the costs (or benefits) assigned to the different states and transitions. This effect is examined next.

7.4.2. Comparison of optimal policies under different cost assumptions

We evaluated the effects of alternative maintenance and decommissioning policies under two additional cost assumptions: high risk-cost (and average outage-cost) assumptions, and high outage-cost (and average risk-cost) assumptions. These two sets of cost assumptions reflect the preferences of two hypothetical decision makers, one driven primarily by safety and the other by productivity. The high risk-cost assumption represents a decision maker who is relatively risk averse in terms of negative safety impacts and therefore places a higher cost value on transitions to and time in accident or other unsafe states. The high outage-cost assumption reflects a decision maker who is relatively risk averse in terms of negative production effects and therefore places relatively high costs on non-producing states.

Table 1
Objective function under alternative maintenance and system life policies

Maintenance cycle length	System life (decommissioning horizon)					
	10	25	40	50	100	150
4	1071.8	1783.6	2011.2	2061.0	2085.9	2085.5
8	1350.5	2307.8	2642.4	2720.0	2763.0	2761.0
10	1389.1	2293.8	2707.5	2783.8	2824.2	2822.8
11	1393.4	2382.8	2716.6	2791.8*	2830.2**	2828.9*
12	1398.5	2385.6*	2717.0*	2789.8	2825.9	2824.8
13	1405.0*	2384.0	2710.3	2780.5	2814.2	2813.2
14	1402.4	2379.6	2698.1	2766.0	2797.1	2796.2
16	1389.3	2356.4	2664.4	2725.0	2751.6	2750.9
20	1377.8	2293.8	2567.8	2618.5	2636.1	2635.7

For each proposed system life (decommissioning horizon), a single asterisk indicates the preferred maintenance cycle length. The double asterisk marks the overall best combination of system life and maintenance cycle for our example.

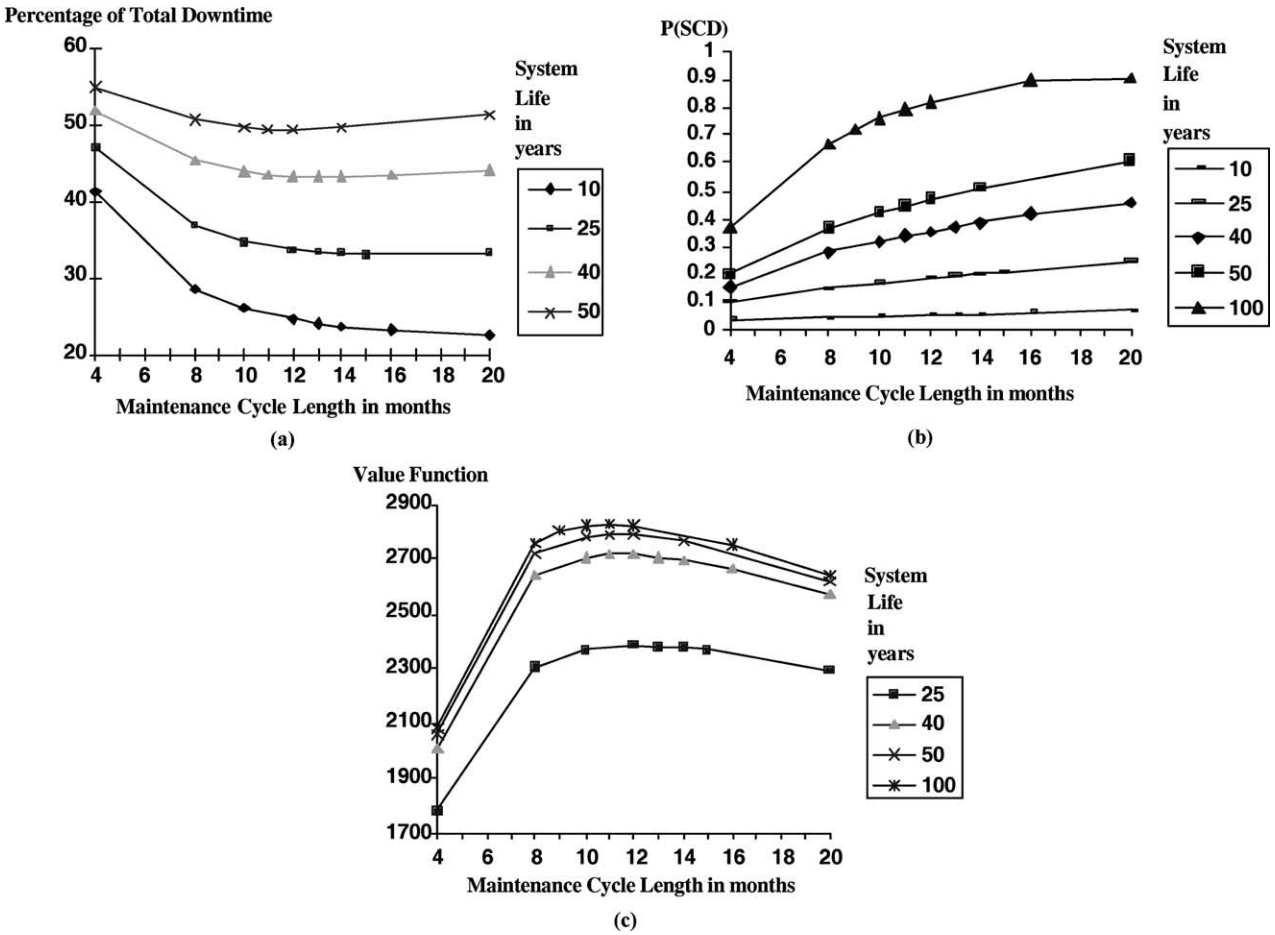

Fig. 2. Illustrative effects of alternative maintenance cycle lengths and system lives on (a) the percentage of time spent in planned or unplanned outage states, (b) the probability of severe core damage, and (c) the overall objective function value.

These alternative perspectives are modeled by adjusting elements of the f and g cost vectors.

Table 2 presents a summary of comparison of preferred planned maintenance and system life policies under the three cost assumptions. High risk-penalty and high outage-penalty each reduce the optimal system life as discussed earlier. The optimal maintenance cycle is reduced under high risk-costs to avoid accidents, and is increased under high outage-costs to avoid costly planned downtime. The optimal policy is clearly dependent on cost assumptions and, in particular, the relative costs of planned versus unplanned outage time.

7.5. The decision to interrupt production

Next, the model is used to determine a decision rule for 'tripping' a reactor because of failures in ECCS. It is assumed that the ECCS consists of three redundant subsystems, at least one of which must be operational for the backup cooling to function when needed. When one or more of the three ECCS subsystems fails, managers (or regulators) face the primary decision of whether to continue operations while the failures are repaired ('hot repair'), or shut the plant down ('cold repair') [33]. The tradeoffs between productivity and safety in this case are similar to those discussed in the warning systems literature [34,35] and in Vaurio's discussion of maintenance decisions and temporary configurations [8].

In our base model, the state space was designed as if the reactor tripped when 2 or more ECCS subsystems failed. Therefore, we assumed continued operation under

Table 2
Comparison of optimal maintenance and life policies under different cost assumptions

Optimal policy	Cost assumption		
	High risk-cost	Base cost assumption	High outage-cost
System life (years)	40	100	50
Maintenance cycle (quarters)	7	11	13

Table 3
Comparison of reactor trip policies under different cost assumptions

	Cost assumption					
	High risk-cost		Base-cost trip policy		High outage-cost	
	New	Base	New	Base	New	Base
Optimal maintenance cycle (quarters)	9	7	13	12	14	13
Expected cost function at optimal point	1959*	1524	2957*	2717	2361*	2089

In the illustrative case, the new more conservative reactor trip policy is preferable to the base trip policy for each of the three cost assumptions. This demonstrates a robust policy improvement under a range of decision-maker value preferences and risk attitudes. A 40-year reactor life is assumed here.

a compromised condition in which one of the backup system had failed (state x_{35}). To consider the effects of a more conservative policy, we modify the state space to prohibit continued operation if any of the ECCS subsystems fails. Under this new policy, the state space is reduced because the probability of transition into a number of states is now zero (namely, states $x_5-x_8, x_{13}-x_{16}, x_{24}-x_{29}$, and $x_{36}-x_{40}$). The probability of transition to an unplanned outage state under the more conservative policy is computed as the probability of one or more ECCS subsystem failures. It is higher than under the original policy that resulted in shutdown only after 2 or 3 subsystem ECCS failures. The probability of a transition out of the unplanned outage resulting from the ECCS policy is also increased because the expected repair time is reduced as only one subsystem, instead of 2 or 3 in our base model, needs repair.

7.5.1. Optimal reactor trip policy

Again, we evaluate the policy alternatives as a function of maintenance cycle length under each of the three cost assumptions: base-cost, high risk-cost and high outage-cost. We assume here a 40-year plant life. Under each cost assumption, the optimum maintenance cycle length has been extended under the more conservative reactor trip policy relative to the base policy. This implies that a more prudent unplanned shutdown policy allows for extended periods of operation between planned maintenance periods. The disadvantages of increased deterioration effects for extended cycles can thus be offset by more conservative shutdown policies.

Table 3 summarizes the effects of the new policy on the optimal maintenance cycle length for each of the three different cost scenarios, and shows the optimum policy for each cost assumption. In this example, the new policy is preferred to the base policy in all cases, which, again, is a result of the chosen cost functions rather than a general finding.

8. Conclusions

Balancing productivity and safety is an essential objective for the designers and managers of safety-critical production systems. Without explicit recognition of the potential tradeoffs and a structured framework for evaluating the long-term effects of operations policies (beyond existing safety regulations), the decision makers can knowingly or inadvertently increase overall failure risks for increases in short-term productivity.

The model presented in this paper provides a framework for the evaluation of alternative risk management strategies based on the predicted operating performance of a critical system in terms of these two objectives. This model requires an explicit statement of the value function of the decision maker assumed here to be top management of the organization. It can be applied to systems such as nuclear power plants, offshore oil platforms, or simply a car. Long-term system performance measures are computed and used as inputs to the decision maker's objective function to determine the preferred risk management strategy. By combining a risk analysis for the physical system, the characteristics of strategic options, and the risk attitude of the decision maker, this model can provide decision support for the design of risk management strategies and, for example, as shown in this paper, to set planned and unplanned outage policies for a nuclear power plant.

In an effort to address with clarity a relatively broad and complex management problem, we have applied a relatively simple modeling approach, relying on Markov assumptions. Along with the clarity afforded by using such an approach, come potential limitations for modeling real critical systems. We have in some cases avoided the potential problems associated with the Markov assumption of memoryless transitions by defining as unique states of operation those from which future transitions depend on the history. For example, the cause of an unplanned outage is likely to determine the rate of return to an operating state. Therefore, we define unique unplanned outage states for each of the possible causes. Still, alternative or more complex models, for example, genetic algorithms as used by Yang et al. [30], might be a fruitful future extension.

A specific problem with the Markov approach is that it cannot accurately handle standby systems that are periodically tested. The failure rate of such systems are not well reflected in the exponential assumption of the Markov model; they are in fact essentially uniformly distributed

over the test interval due to latent failures in such systems that are not detected between tests. The use of Apostolakis et al.'s [36] analytic unavailability equations for our standby systems or Vaurio's [37] related approach that evaluates the impact of alternative maintenance policies on standby system availability, might improve the result. Because this situation is quite relevant to our illustrative nuclear plant example, further study should be carried out to explore the significance of this limitation.

A final challenge in applying our approach is that it currently requires the subjective linking of alternative operations and maintenance policies to our system. Consequently, the results of the optimization model are only as accurate as are the decision maker's assumptions about the impact of specific policy options on the parameters of the models. In an industry that possesses an extensive data base on operating experience that is clearly linked to specific elements of different risk management strategies, it might be possible over time to develop more objective connections between policy choices and system performance characteristics.

Acknowledgements

This work was funded in part by the National Science Foundation under grant SBR 94-22946. This support is gratefully acknowledged.

References

[1] Starr C, Whipple C. Coping with nuclear power risks: the electric utility incentives. In: Buchanan JR, editor. Nuclear safety, vol. 1; 1982. p. 1–7.
[2] Paté-Cornell ME. Risk assessment and risk management for offshore platforms: lessons from the piper alpha accident. J Offshore Mech Arctic Engng 1993;115:179–90.
[3] Cullin. The Hon Lord. The public inquiry into the piper alpha disaster, vol. 1. Report to Parliament by the Secretary of State for Energy; 1990.
[4] Presidential Commission. Report on the Space Shuttle Challenger Accident, vol. 1, Washington, DC, 1986.
[5] Baron MM, Paté-Cornell ME. Designing risk management strategies for critical engineering systems. IEEE Trans Engng Mgmt 1999; 46(1):87–100.
[6] Rothwell G, Rust J. A dynamic programming model of US nuclear power plant operations. Working paper; 1995.
[7] David PA, Maude-Griffin R, Rothwell GS. Learning by accident? Reductions in the risk of unplanned outages in US nuclear power plants after three mile island. J Risk Uncertainty 1996;13:175–98.
[8] Vaurio JK. Safety-related decision making at a nuclear power plant. Nucl Engng Des 1998;185:335–45.
[9] Vesely WE. Quantifying maintenance effects on unavailability and risk using Markov modeling. Reliab Engng Syst Safety 1993;41: 177–87.
[10] Martorell S, Munoz A, Serradell V. An approach to integrating surveillance and maintenance tasks to prevent the dominant failure causes of critical components. Reliab Engng Syst Safety 1995;50: 179–87.
[11] Siu N. Risk assessment for dynamic systems: an overview. Reliab Engng Syst Safety 1994;43:43–73.
[12] Aldemir T, Siu NO, Mosleh A, Cacciabue PC, Göktepe BG, editors. Reliability and safety assessment of dynamic process systems. NATO ASI series, Berlin: Springer; 1994.
[13] Vaurio JK. Optimization of test and maintenance intervals based on risk and cost. Reliab Engng Syst Safety 1995;49:23–36.
[14] Garrick BJ. Recent case studies and advancements in probabilistic risk assessment. Risk Anal 1984;267–79.
[15] Mosleh A. Lessons learned from nine plant-specific external flooding analysis. Presentation at the Severe Accident Policy Implementation External Events Workshop Annapolis, Maryland; August 4–5 1987.
[16] Budnitz RJ. Current status of methodologies for seismic probabilistic safety assessment. Reliab Engng Syst Safety 1998;62: 71–88.
[17] Keeney RL, Raiffa H. Decisions with multiple objectives: preferences and value tradeoffs. New York: Wiley; 1976.
[18] Howard RA. On making life and death decisions. In: Howard RA, Matheson JE, editors. Reprinted in: readings on the principles and applications of decision analysis, vol. 2. Matheson. Palo Alto: Strategic Decisions Group; 1984.
[19] Bier VM. Issues in the estimation of ageing in event frequencies. In: Proceedings of an International Symposium on the Use of Probabilistic Safety Assessment for Operational Safety, International Atomic Energy Agency, American Nuclear Society, European Nuclear Society, and the Nuclear Energy Agency of the OECD. Vienna; June 1991.
[20] Martorell S, Sanchez A, Serradell V. Age-dependent reliability model considering effects of maintenance and working conditions. Reliab Engng Syst Safety 1999;64:19–31.
[21] Paté-Cornell ME. In: Grigoriu M, editor. Discounting in risk analysis: capital vs. human safety. Risk, structural engineering and human error. University of Waterloo; 1984.
[22] EQE International, Soviet nuclear plant safety: Bulgaria's challenge. EQE Rev 1992;Fall.
[23] US Nuclear Regulatory Commission. Shutdown and low-power operation at commercial nuclear power plants in the United States. NUREG-1449; 1993.
[24] Denton HR. Reactor scrams in the US: a regulators point of view. In: Proceedings of an NEA Symposium on Reducing the Frequency of Nuclear Reactor Scrams, Tokyo: OECD; 1987. p. 66–74.
[25] US Nuclear Regulatory Commission. Reactor safety study: an assessment of accident risks in US. Commercial nuclear power plants. NUREG-75/014 (WASH-1400); 1975.
[26] Budnitz RJ, Davis PR. A scoping evaluation of severe accidents at the Surry and Grand Gulf nuclear power plants resulting from earthquakes during shutdown conditions. Report; August 1991.
[27] EPRI. Safety Assessment of Diablo Canyon risks during shutdown operations. EPRI Outage Risk Assessment and Management Program, Nuclear Safety Department, EPRI; 1992.
[28] US Nuclear Regulatory Commission. A process for risk-focused maintenance. NUREG/CR-5695; 1991.
[29] US Nuclear Regulatory Commission. Handbook of methods for risk-based analyses of technical specifications. NUREG/CR-6141; March 1995.
[30] Yang J-E, Sung T-Y, Jin Y. Optimization of the surveillance test interval of the safety systems at the plant level. Nucl Technol 2000; 132:352–65.
[31] Harunuzzaman M, Aldemir T. Optimization of standby safety system maintenance schedules in nuclear power plants. Nucl Technol 1996; 113:354–67.
[32] Čepin M, Martorell S. Evaluation of allowed outage time considering a set of plant configurations. Reliab Engng Syst Safety 2002;78: 259–66.
[33] Mankamo T. Operational decision alternatives in failure situations of standby safety systems—development of probabilistic approach

and PC program TeReLCO. Reliab Engng Syst Safety 1992;36: 29–34.
[34] Paté-Cornell ME. Warning systems in risk management. Risk Anal 1986;5(2):223–34.
[35] Paté-Cornell ME, Lee HL, Tagaras G. Warnings of malfunction: the decision to inspect and maintain production processes on schedule or on demand. Mgmt Sci 1987;33(10):1277–90.
[36] Apostolakis G, Chu TL. The unavailability of systems under periodic test and maintenance. Nucl Technol 1980;50:5–15.
[37] Vaurio JK. On time-dependent availability and maintenance optimization of standby units under various maintenance policies. Reliab Engng Syst Safety 1997;56:79–89.
[38] Paté-Cornell ME. Quantitative safety goals for risk management of industrial facilities. Struct Safety 1994;13:145–57.