Working Paper

A Decision Model for the Risk Management of Hazardous Processes

Jan Holmberg
YSSP Participant 1995

WP-95-95
September 1995

IIASA International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 807 □ Fax: +43 2236 71313 □ E-Mail: info@iiasa.ac.at

A Decision Model for the Risk Management of Hazardous Processes

Jan Holmberg
YSSP Participant 1995

WP-95-95
September 1995

International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 807 □ Fax: +43 2236 71313 □ E-Mail: info@iiasa.ac.at

Foreword and Acknowledgments

The research work was carried out at the International Institute for Applied Systems Analysis (IIASA) during the Young Scientists' Summer Program 1995 in the project Risk, Policy and Complexity. The study is part of the Reliability and Risk Analysis project in the Finnish nuclear energy research program Reactor Safety (RETU), 1995–1998. The study was financed by the Ministry of Trade and Industry in Finland and VTT Automation. Financial Support for visiting IIASA was also given by the Finnish Committee for IIASA. The author would like to thank Yuri Ermoliev at IIASA and Urho Pulkkinen at VTT Automation for their helpful discussions and comments.

Abstract

We formulate a decision model for the risk management of hazardous processes as an optimization problem of a point process. The essential features of the model are: long-term (process lifetime) objective function which is a risk-averse utility function, a dynamic risk model (marked point process model) representing the stochastic process of events observable or unobservable to the decision-maker and a long-term control variable guiding the selection of optimal solutions for short-term problems.

The model is demonstrated by a case study of a hazardous process with reparable safety systems, such as a nuclear power plant. The short-term decision problem of the case study is whether it is sometimes beneficial to temporarily shut the process down in order to cut off the high risk periods. The long-term decision problem is to optimize a long-term control variable that determines which decision alternative is preferred in a case of increased risk in the process: (1) to shut the process down during the repair time or (2) to continue the operation. Several long-term strategies are analysed and compared. As a solution approach for the optimization problem, we use the stochastic quasi-gradient procedure.

Contents

A Decision Model for the Risk Management of Hazardous Processes

Jan Holmberg
YSSP Participant 1995 *

1 Introduction

We consider processes which may lead to catastrophic consequences with a low probability. Examples of such processes are nuclear power plants, chemical plants, transportation of hazardous materials, air traffic, seafaring. We can also mention similarity to such processes as environmental degradation although our discussion is in the context of technological processes.

To control hazardous processes is a problem of risk management. Problems in risk management mostly arise from the complexity of the process, and several approaches are usually applied to confront complexity (Wahlström 1992), such as inherently safety design of the systems, safety regulations, quality control, safety analyses and operating experience feedback. We would like to build a model which incorporates the *short-term* operational risk management with the *long-term* safety objectives. Therefore we consider risk management to be a process where several interrelated problems are solved driven by various events which brings new information about the process.

By safety related operational problems we mean, for instance, questions like what to do when failures degrade the safety level of the process, and how to schedule the surveillance of the safety systems. Risk management must daily solve this kind of problems even if the process is in a good condition, because, particularly then, safety is one decision criterion compared with the economical consequences of the decision. In a way, safety management searches all the time for a balance between safety and economy. For instance, in the nuclear safety context the so called ALARA principle — as low as reasonably achieved — is applied when decisions are made about how far the risks should be minimized. What is then a "reasonably" low risk level is a decision problem.

Of course, if some disturbances or incidents occur in the process, safety becomes the primary concern for the management since without rapid actions an accident may happen or the consequences of the accident may become catastrophic. This area of risk management, called emergency or accident management, is usually controlled by procedures. However, a process is typically most of the time in normal conditions and this is also our application area.

The problems have different time spans. We can divide problems into two categories: (1) long-term problems, and (2) short-term problems. In long-term problems, the decision-maker (DM) wants to improve the system by making permanent changes in the design, procedures or other practices. In short-term problems, temporary safety related problems are solved. A typical example is to decide whether to shut down the operation of the plant in a case of a failed condition in safety related systems.

Traditionally, decision analysis has been applied in individual problems, particularly in long-term problems. Short-term problems are a newer application area. The need arises from the

*Member of IIASA's Young Scientists' Summer Program 1995. Home Institute: VTT Automation, Industrial Automation, P.O.Box 1301, FIN-02044 VTT, Finland

fact that complex systems can be threatened by unexpected events whose uniqueness requires taking actions that are different from the procedures (Peroggi and Wallace 1994). If decision analysis is applied, the structure of the decision model must be prepared in advance. However, the basic issue of accepting large losses with small probability remains in all cases.

The solving of individual problems does not necessarily lead to a satisfactory result since the problems are interrelated in a complex way. Decision trees and influence diagrams are too simple because, for instance, the order in which events can take place is stochastic. Further, the risk aversion towards accidental outcomes should be accounted for. This can be expressed by a utility function. However, then the outcome function is not anymore linear. It means that, from the long-term perspective, the operating history not only affects the probability of the future events, but it also affects the value of future prospects. In order to maximize the long-term objective function, we need a global (long-term) control over the short-term decision making.

The aim of this paper is to model the risk management process as sequential decision making in a stochastic environment. In our approach, we interpret the controls as decision rules for the management. Mathematically, the problem is to find an optimal control for a point process. Optimization of such processes can be performed analytically only in very simple cases. Appropriate search procedures for this type of problem have been promoted at the International Institute for Applied Systems Analysis (IIASA), see e.g. (Ermoliev and Wets 1988). The research is presently towards dealing with "surprises", abrupt transition jumps (Ermoliev et.al. 1995, Oortmarssen and Ermoliev 1994), which occur in our applications, too. Therefore we take advantage of the results of this research, and we apply the stochastic quasi-gradient method as one solution approach.

Another aim of this paper is to introduce a utility function that represents the risk aversion of the decision-maker. By manipulating the shape of the utility function, we can study how different probabilistic decision criteria in the long-term level are reflected in short-term decision making. To our knowledge, long-term probabilistic criteria and short-term risk based operational rules have not been linked this way before. Perhaps, reasons for this have been the lack of dynamical decision models for risk management of hazardous processes and the lack of appropriate search procedures. Now, facilitated with a point process model of the system and the stochastic quasi-gradient algorithm, we can develop a decision analytic approach to integrate the probabilistic safety assessment (PSA) into risk management. This is the novelty of our paper.

The stochastic optimization methods have already been applied in the risk management context to evaluate optimal test intervals and inspection strategies (see e.g. Pulkkinen and Uryas'ev 1990). In those decision models, the problem is to monitor the ageing of hazardous processes. We will study reparable safety systems which have a cyclic reliability dynamic. In other words, ageing deteriorates the components of the system but once in a while the components are maintained, repaired or replaced, which improves the reliability of the systems. In our case study, the main problem is whether it is sometimes beneficial to temporarily shut down the process in order to cut off the high risk periods.

The paper is organized as follows. In Section 2, we describe the general structure of the model. In Section 3, we demonstrate the approach by a case study. In Section 4, we discuss the applicability of the model and extensions of the case study.

2 General description of the model

In this section, we formulate a decision model for risk management. First we give an overview of the model, and then, in the following subsections, we explain the details of the model. Table 1 summarizes our notations.

Figure 1 illustrates the time axis of the process. The sequence $T_0, T_1, \ldots$ denotes randomly occurring events or predetermined time epochs, when operational decisions can be made. The sequence $Z_0, Z_1, \ldots$ denotes marks corresponding to the nature of the events. In other words, we obtain various kinds of information about the process at discrete time epochs. The information

Table 1: Notation index.

Ω	sample space
ω	outcome in a set Ω
$\mathcal{F}$	σ-algebra of Ω
Z_n	mark of the nth event
T_n	time of the nth event
T^L	licensing time of the process
(T_n, Z_n)	a marked point
$(\bar{T}, \bar{Z})$	the terminal point of the process
E	set of marks
E^0	set of initiating event categories
$\mathcal{E}$	σ-algebra of E
$\hat{E}$	set of observable marks
$\bar{E}$	set of termination marks
$N_t(E_1)$	counting process of marks $z \in E_1$
$H_t / \hat{H}_t$	full/observed process history
h	a sample path of the full process history
$\mathbf{H}$	space of process histories
$\lambda_t(y)$	accident hazard rate of category y accident at t
$\hat{\lambda}_t(y)$	monitored accident hazard rate
$\lambda^n(y)$	nominal accident hazard rate
$\lambda^b(y)$	baseline accident hazard rate
$\lambda^0(y)$	inherent accident hazard rate
$\lambda_t(z)$	initiating event intensity of category $z \in E^0$
$c(t, z; y)$	conditional probability of consequence y when z takes place at t
$c^n(z; y)$	nominal safety system failure probability
x	control
X	set of control variables
a	decision option
$A(Z_n)$	space of decision options depending on the mark Z_n
J	process lifetime outcomes
J_t	outcomes up to t
M	costs of an accident
β	cost rate function
κ	discrete cost function
ζ	process availability
S_t	cumulative operation time of the process up to t
t^r	repair time of a failed component
η	degradation degree of the system
$u(\cdot)$	utility function
$V(\cdot)$	short-term decision function

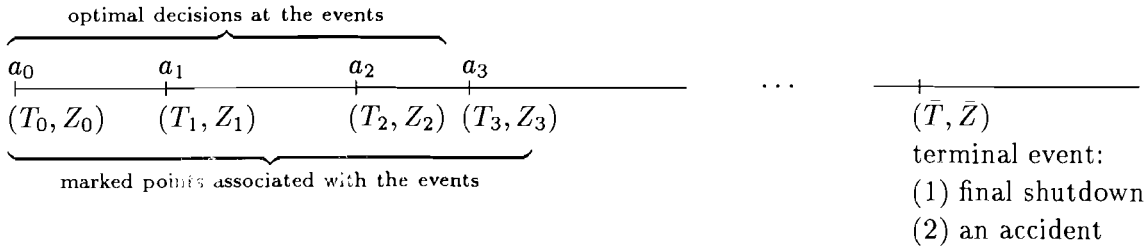

Figure 1: Marked point process of events and decisions.

can be a failure or degradation in the process or it can be a shock to the process like loss of the external power of a technical system or an earthquake. The terminal event of the process $(\bar{T}, \bar{Z})$ can be either the final shutdown of the process or an accident. Not all events require problem solving, and there can be latent events which are not observed.

For short-term risk management, an optimal solution a_n is chosen by maximizing a decision function

$$a_n = \arg \max_{a \in A(Z_n)} V(a, \hat{H}_{T_n}, x),$$

where $A(Z_n)$ denotes the set of decision options depending on the problem Z_n, $V(\cdot)$ is the decision function, $\hat{H}_{T_n} = \{(T_i, Z_i); T_i \leq T_n, T_i \in \hat{E}\}$ is the observed history up to T_n and $x \in X$ is the vector of long-term control variables such as parameters of the decision model and indicator variables when to make decisions. The definition of the control variables is the essential modelling problem of this formulation.

The lifetime profits depend on the history of events and past decisions. The profit function can be divided into time intervals according to the decision making epochs as follows,

$$J(x, \omega) = \sum_{n=0}^{\bar{n}-1} \left(\int_{T_n}^{T_{n+1}} \beta(t, a_n, \omega) dt + \kappa(T_n, a_n, \omega) \right) + \kappa(\bar{T}, \bar{Z}),$$

where $\beta(\cdot)$ is the profit rate function (profits or costs per time unit), $\kappa(\cdot)$ represents profits or costs associated with discrete time points and $\bar{n}$ is the index of the terminal event. The long-term decision problem is to adjust the control variables so that the expected life time profits are maximized, i.e.,

$$\begin{aligned} &\text{maximize } F(x) = \mathbf{E}_\omega[J(x, \omega)] \\ &\text{subject to } x \in X. \end{aligned}$$

The above formula assumes that the decision-maker (DM) is risk neutral. However, it would be more reasonable to assume that the DM is risk averse, i.e., great losses are avoided more than the expected value formula would suggest. Therefore we introduce a utility function to account the DM's preferences over uncertain outcomes. Thus the objective function is

$$\begin{aligned} &\text{maximize } F(x) = \mathbf{E}_\omega[u(J(x, \omega))] \\ &\text{subject to } x \in X. \end{aligned} \tag{1}$$

In the above formulation, we have incorporated the long-term decision problem related to the control variable $x \in X$ with the short-term decision problem of choosing between decision options $a \in A(Z_n)$, $n \geq 1$.

The optimization of (1) is a complicated task. The objective function may be discontinuous with respect to the argument x, and generally the expected value $F(x)$ cannot be evaluated analytically. Particularly in our applications the computational difficulties are due to feedback mechanisms and due to the non-linear utility function. However, by simulation of sample paths of the process history ω, for some x a simulated life time utility $u(J(x, \omega))$ is obtained. Applying the stochastic quasi-gradient algorithm, we can approach the solution of the optimization problem.

2.1 A marked point process model

A marked point process $\{(T_n, Z_n); n = 1, 2, \ldots\}$ is an ordered sequence of time points T_n and marks $Z_n \in E$ associated with the time points. The marked point process framework allows us to model processes where relevant information consists of various type of discrete events. A counting process $N_t(E_1)$ counts the number of marked points (T_n, Z_n) with marks in a set $E_1 \in \mathcal{E}$ up to time t, i.e.,

$$N_t(E_1) = \sum_n 1_{\{T_n \leq t, Z_n \in E_1\}}, \; t \geq 0, \; N_0(E_1) = 0.$$

$N_t(E_1)$ is thus a step function taking a jump of size 1 when a mark belonging to E_1 occurs. In the particular case where E_1 is a singleton, say $E_1 = \{z\}$, $z \in E$, we denote the counting process by $N_t(z)$.

The history process H_t is formed by marked points up to time t

$$H_t = \{(T_n, Z_n); T_n \leq t\},$$

and H_{t-} is defined in the same way except that the inequality is strict: $T_n < t$. H_t takes values in the space $\mathbf{H}$ which is a subset of $[0, \infty) \times E$.

The z-specific hazard rate or intensity at t given the history H_{t-} can be written as

$$\lambda_t(z) = P(dN_t(z) = 1 \mid H_{t-})/dt = \lambda(t, z \mid H_{t-}),$$

where it has been assumed that the corresponding measure is absolutely continuous with respect to the Lebesgue measure. More generally, if absolute continuity cannot be assured, we can use a hazard measure $d\Lambda_t(z)$ with the interpretation

$$d\Lambda_t(z) = P(dN_t(z) = 1 \mid H_{t-}).$$

2.1.1 Decomposition of process histories

The marked point process model represents only the most important part of the actual process history, forming a "landmark process". From an observer's point of view, the landmark process may contain marked points that remain *latent*, unobserved, at least for a while. In our case study, there are no latent events. However, we take this possibility into account for future extensions of the problem, and in order to apply similar denotations when sampling partial process histories in the stochastic quasi-gradient algorithm.

Let $\hat{H}_t$ denote the *observed* pre-t process history data,

$$\hat{H}_t = \{(T_n, Z_n); T_n \leq t, Z_n \in \hat{E}\},$$

where $\hat{E}$ is the set of observable marks. Each observed history $\hat{H}_t$ is fully determined by the underlying full marked point process history H_t. Consequently, the observed hazard rate can be expressed as an expected hazard rate as follows

$$\begin{aligned} \hat{\lambda}_t(x) &= P(dN_t(z) = 1 \mid \hat{H}_{t-})/dt \\ &= \int_{h \in \mathbf{H}} P(H_{t-} \in dh \mid \hat{H}_{t-}) \, \lambda(t, z \mid h) \\ &= \mathbf{E}[\lambda(t, z \mid H_{t-}) \mid \hat{H}_{t-}], \end{aligned} \tag{2}$$

where $P(H_{t-} \in dh \mid \hat{H}_{t-})$ is the conditional probability that the full process history H_{t-} is in the elemental volume dh of $\mathbf{H}$, given the observed, strict pre-t process history. A sample path of the full process is denoted by $h = \{(t_n, z_n); \; n \geq 1\}$, and the corresponding pre-t histories by $h_t = \{(t_n, z_n); \; t_n \leq t\}$ and $h_{t-} = \{(t_n, z_n); t_n < t\}$.

2.1.2 Terminal event

We denote the mark corresponding to the terminal event by $(\bar{T}, \bar{Z})$, i.e.,

$$(\bar{T}, \bar{Z}) = \{(T, Z) \mid T = \min_{T_n}(T_n, Z_n), Z_n \in \bar{E}, Z = Z_n\}, \tag{3}$$

where $\bar{E}$ is the set of termination marks. The terminal event can be an accident or the final shutdown of the process. We assume that the operation time of the process is limited by a licensing time T^L. If no accident happens, then $\bar{T} = T^L$.

2.2 Accident hazard

The accident hazard is estimated by the model of the process, in this paper called the risk model. Our concept for the risk model is based on the event tree-fault model used in a probabilistic safety assessment (PSA) for the risk analysis of the operation of a nuclear power plant. The model could as well represent other processes where the course of an accident is caused by randomly occurring system disturbances and subsequent failures of the safety barriers. The accident process is thus a compound process of an initiating event process and a safety system process. In the nuclear power plant context, the accident is a core damage.

Traditionally, PSA models have been static, expressing average conditions at the systems. A static risk model may be applicable for long-term problems, but it is insufficient for our purposes. Therefore, we apply a dynamic risk model based on the marked point process framework (Arjas and Holmberg 1995). We also introduce the basic risk measures of the static risk model since they have a practical meaning as reference risk levels, when the long-term safety objectives are considered.

2.2.1 Dynamic risk model

A dynamic risk model expresses the momentary risk as a consequence of the actual conditions of safety related equipment in the plant. The *instantaneous* accident hazard rate is the basic risk measure. We denote it by

$$\lambda_t(y) = \sum_{z \in E^0} \lambda_t(z)\, c_t(z, y), \tag{4}$$

where $z \in E^0$ indexes the initiating event categories, $\lambda_t(z)$ is the initiating event intensity of category z, and $c_t(z, y)$ the conditional probability that consequence y results in when z takes place at time t. The instantaneous accident hazard rate is obtained by using initiating event intensities as well as component unavailabilities based on the up-to-date operating experience of the system.

In the risk monitoring, the accident hazard rate is evaluated dynamically based on the observed history $\hat{H}_{t-}$

$$\hat{\lambda}_t(y) = \sum_{z \in E^0} \left(\int_{h \in \mathbf{H}} P(H_{t-} \in dh \mid \hat{H}_{t-})\, \lambda(t, z \mid h)\, c(t, z; y \mid h) \right), \tag{5}$$

where $P(H_{t-} \in dh \mid \hat{H}_{t-})$ is the conditional probability that the full process history H_{t-} is in the differential volume dh of $\mathbf{H}$, given the observed, strict pre-t process history. $\hat{\lambda}_t(y)$ is called the *monitored* accident hazard rate.

2.2.2 Static risk measures

In the PSA context, we can define three static risk measures: *nominal*, *baseline* and *inherent* accident hazard rate (Holmberg et al. 1993). The nominal accident hazard rate represents the average accident hazard rate of the system. It is obtained by the use of nominal or time-average

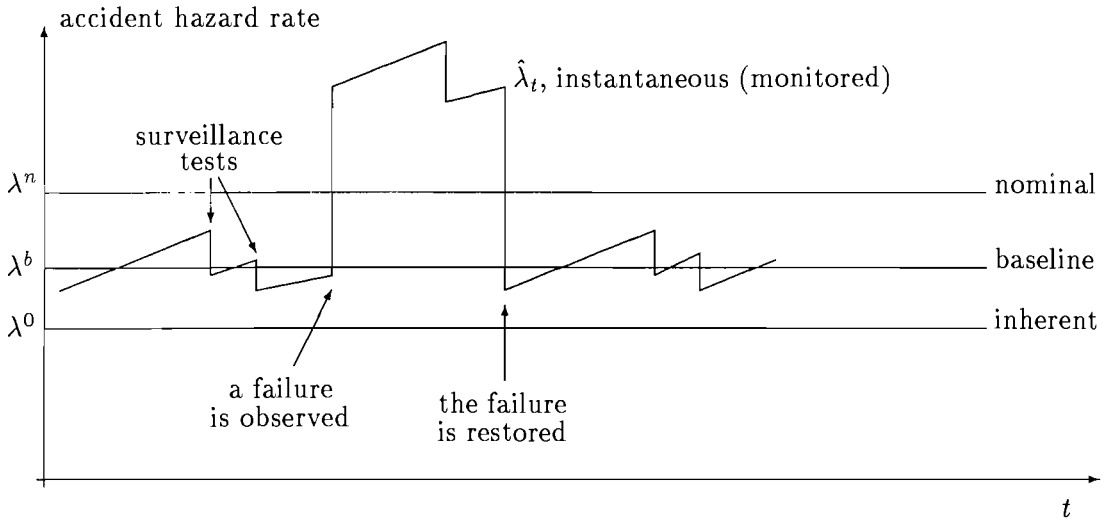

Figure 2: Different accident hazard rates.

unavailabilities for the components and by the use of nominal initiating event intensities. We denote the nominal accident hazard rate by

$$\lambda^n(y) = \sum_{z \in E^0} \lambda^n(z)\, c^n(z, y), \tag{6}$$

where $\lambda^n(z)$ is the nominal initiating event intensity of category z, and $c^n(z, y)$ is the nominal conditional probability that consequence (accident category) y results when z takes place.

The momentary variations in the instantaneous accident hazard rate are mainly caused by two kinds of events. Firstly, there are evident failures or other evident events like maintenance of the systems which temporarily increases the monitored accident hazard rate, $\hat{\lambda}_t(y)$. If the evident unavailabilities are excluded from the nominal accident hazard rate, a baseline accident hazard rate, $\lambda^b(y)$, is obtained. Normally (not nominally), the risk level of the process should be close to the baseline hazard rate, particularly, if the failures in the safety systems are not very frequent. Therefore the baseline accident hazard rate is an applicable reference risk level, for instance, for the evaluation of the unavailabilities allowed by the Technical Specifications of a nuclear power plant (IAEA 1993).

The second category of events causing variation in the instantaneous accident hazard rate is latent failures. Some of the latent failures can be detected by surveillance tests. The inherent accident hazard rate, $\lambda^0(y)$, corresponds to conditions of the safety systems where no component is unavailable due to maintenance or repair (as in the baseline accident hazard rate), and standby components have recently been tested without any failure indications. It represents the "lowest theoretically achievable" accident hazard rate with the current design of the systems.

Different hazard rates are illustrated in Figure 2. The saw-teeth shape of the instantaneous accident hazard rate is due to the contribution of latent failures which increases the hazard rate between test epochs. Most of the time, the instantaneous accident hazard rate varies around the baseline, it can never go below the inherent level, and when significant failures are detected, it rises above the nominal level.

2.3 Controls

Three types of control of point processes can be distinguished: optimal stopping, intensity control and impulsive control (Bremaud 1981). Optimal stopping means the possibility to determine the terminal point of the system. After that the system will neither give profits nor cause costs.

Optimal stopping time may be restricted by the licensing time of the system. We assume that there is a licensing time T^L which cannot be exceeded.

In the intensity control, the intensity of some events like failures can be modulated at some expense. The intensities can be affected e.g. by changing the components of the system. In our formulation, the intensity control is a long-term problem.

In the impulsive control, the decision-maker can add or erase points of the process. The decision-maker can add points by determining beforehand time points when some actions are to be taken, such as surveillance tests. Then by choosing the decision option at a time epoch, the decision-maker attaches a mark to the process history. We will consider two kinds of impulsive controls: temporary shutdowns of the process and surveillance tests. An impulsive control can be sometimes interpreted as an intensity control, too.

2.4 Lifetime profit function

The profit function has two parts: (1) profit rate function as long as the system is operated and (2) various kinds of discrete costs depending on the events and decisions made. We assume that all profits and costs can be represented in a monetary scale, as if other types of losses or benefits can be exchanged into monetary units.

In this paper, we assume a constant rate of incomes, $\beta > 0$, if the system is operated and one accident category with the costs $-M < 0$. The lifetime profits depend on how long time the system has been operated and whether an accident happened or not. Let $\bar{T}$ be the terminal time of the operation (formula (3)). Then the lifetime profits are

$$J(x,\omega) = \begin{cases} \beta S_{\bar{T}}(\omega) & \text{if no accident happens} \\ \beta S_{\bar{T}}(\omega) - M & \text{if an accident happens,} \end{cases} \tag{7}$$

where $S_t(\omega) \leq t$ is the cumulative operation time of the process. Note that $J(x,\omega)$ is positive if no accident happens, and it can be negative only if an accident happens.

Depending on the required realism of the model, the profit function can be made more accurate. One question is, whether costs or incomes should be discounted. We do not discount the rate of incomes or the costs of an accident since, for instance, in the nuclear power plant context and from the power company point of view, the rate of incomes depends on the price of electricity and the costs of an accident is at least the price of a new nuclear power plant (that is the smallest accident category we are considering).

2.5 Utility function

The utility function represents the DM's preferences over uncertain outcomes. The maximum costs associated with the accident correspond to the utility 0, i.e., $u(-M) = 0$. The best possible outcome depends on the terminal time.

The form of the utility function is crucial for the rest of the decision analysis. We provide here one approach to formulate it. It is based on the acceptance of the present safety level of the process as if the safety authority and the responsible company had implicitly agreed on the utility function in the licensing phase of the process. On the other hand, as is well known, accepted risk levels vary between different hazards in the society. Therefore, the utility function used in the operation of a nuclear power plant is probably not applicable in another context. The key assumption is that the risk management applies the same utility function for all decision making concerning the process they are responsible for.

We consider a differential time unit and compare the shutdown option to the operation of the process. The utility function should be such that in normal conditions the operation is preferred to shutdown. Assuming that the lifetime profits so far are J_t, the expected utility of the shutdown option a^s is

$$\mathbf{E}[u \mid a^s] = u(J_t).$$

Given the accident hazard rate λ, the operation alternative corresponds to the lottery a^c : $\langle e^{-\lambda dt}, J_t + \beta dt; 1 - e^{-\lambda dt}, J_t - M \rangle$. The expected utility of the operation over a differential time unit dt is

$$\mathbf{E}[u \mid a^c] = (1 - e^{-\lambda dt})u(J_t - M) + e^{-\lambda dt}u(J_t + \beta dt).$$

Requiring $a^s \prec a^c$, i.e. $\mathbf{E}[u \mid a^s] < \mathbf{E}[u \mid a^c]$, we obtain an inequality

$$\frac{u(J_t + \beta dt) - u(J_t)}{1 - e^{-\lambda dt}} > u(J_t + \beta dt) - u(J_t - M). \tag{8}$$

When $dt \to 0$,

$$\frac{\beta}{\lambda}u'(J_t) > u(J_t) - u(J_t - M). \tag{9}$$

The condition (9) can also be obtained in another way by considering the optimal terminal time. Let x be the decided terminal time and ω (unknown) time of the accident. The outcome of the operation is

$$J(x, \omega) = \begin{cases} \beta x & \text{if } x < \omega, \\ \beta\omega - M & \text{if } x \geq \omega. \end{cases}$$

If λ is the constant accident hazard rate, then the expected utility is

$$\mathbf{E}[u(J(x,\omega))] = \int_0^x \lambda e^{-\lambda t}u(\beta x - M)\, dt + e^{-\lambda x}u(\beta x).$$

The derivative of this expression is

$$\frac{d}{dx}\mathbf{E}[u(J(x,\omega))] = \lambda e^{-\lambda x}u(\beta x - M) - \lambda e^{-\lambda x}u(\beta x) + e^{-\lambda x}\beta u'(\beta x).$$

Requiring that the derivative is positive for all $x > 0$, i.e. the optimal solution is $x^* = \infty$, we obtain the condition (9).

Below we study two classes of utility functions which can be used for describing risk aversion: power function and exponential function. In the case study, we will use the exponential function.

2.5.1 Power function

One common suggestion (see e.g. IAEA (1989)) for a utility function in the risk management context is a power function

$$u(\gamma) = 1 - \left(\frac{\beta T^L - \gamma}{\beta T^L + M}\right)^\alpha, \; -M \leq \gamma \leq \beta T^L, \; \alpha \geq 1, \tag{10}$$

where γ is the outcome and α is a risk aversion factor. Note that we have scaled here the function to fit to our outcome space. If $\alpha = 1$ we have a linear function, and the DM is risk neutral. If $\alpha = 2$, then the utility function is consistent with the recommendations by the Dutch authorities for the management of major hazards (Anon. 1989). Following (9), we obtain a condition

$$\frac{\beta}{\lambda}\alpha(\beta T^L - J_t)^{\alpha-1} > (\beta T^L - J_t + M)^\alpha - (\beta T^L - J_t)^\alpha. \tag{11}$$

If $\alpha = 1$ (risk neutral case), then the condition for the preference of the operation is

$$\frac{\beta}{\lambda} > M, \tag{12}$$

independently of J_t. If $\alpha = 2$ (Dutch authority case), then the condition for the preference of the operation is

$$\frac{\beta}{\lambda} > M\left(1 + \frac{M}{2(\beta T^L - J_t)}\right). \tag{13}$$

2.5.2 Exponential utility function

The exponential utility function has the form

$$u(\gamma) = 1 - e^{-\alpha(\gamma+M)},\ \gamma \geq -M,\ \alpha > 0, \tag{14}$$

where α is a shape parameter. The greater α is, the more risk averse is the DM. For the exponential utility function, we obtain the condition

$$e^{-\alpha M} > \frac{\lambda}{\lambda + \alpha\beta}, \tag{15}$$

which does not depend on J_t. Given an upper limit for the accident hazard rate λ^*, (15) provides an upper limit for α. On the other hand, we can note that, if $\lambda^* \geq \beta/M$ is accepted (c.f. (12)), then the operation is accepted for all $\alpha > 0$. The interpretation of λ^* is that it is a certain maximum allowed accident hazard rate. Whenever the risk increases above a certain level, the process should be shut down. However, the shutdown of the process itself includes a discrete risk which should be accounted. This is taken into account in the case study.

We think that an exponential function might be an appropriate choice of the utility function even for practical applications. The question is how α or λ^* should be defined. Perhaps λ^* is given by the safety authority and the responsible management then chooses an acceptable utility function. After the selection of the utility function, the management tries to operate the system within the allowed "safety margin" in an optimal way.

3 Analysis of a reparable safety system

The case study is a test problem by which we can compare the stochastic quasi-gradient method with analytical results. Once the results are confirmed, the case can easily be extended to a more realistic one. Even though the system is simple, we think that it is rather illustrative concerning dynamical safety evaluation of the operation of a nuclear power plant. We point out that *the results and conclusions depend on the chosen model parameters and the chosen forms of the cost and utility functions.* The methodological part is more invariant.

We consider a process with one initiating event category and a safety system with two redundant components. An accident takes place if an initiating event occurs and both the components fail. We have only one accident category, so we omit the accident category variable y in our denotations.

Randomly occurring shocks degrade the other component of the safety system causing an increase of the accident hazard rate. Effectively, it means that the probability of the safety system failure, term $c(t, z \mid h)$ in (5), increases. When the degraded condition is detected, the reparation of the system is started in order to restore the normal condition of the safety systems.

We consider the problem of a temporary shutdown of the process for the reparation period, i.e., for the period of increased risk. The approach can be extended to other short-term decision making problems as will be discussed. Figure 3 shows an example realization of the process from the point of view of process availability, safety system failure probability and monitored accident hazard rate. When the process is operated, the production is at a 100% level, and when it is shut down it is at a 0% level. The decisions concerning temporary shutdowns are made when a mark '1' occurs. If the shutdown is chosen, then a mark '4' is attached to the same time epoch. There are two failure epochs in this realization. The first is the continued operation. The second is that the process is shut down. The variables t_1^τ and t_2^τ denote the repair times and η_1 and η_2 are some measures of the degradation degrees of the safety system. An initiating event marked by '0' takes place between the two failures, but it does not result in an accident. Table 2 summarizes the possible marks in the process history.

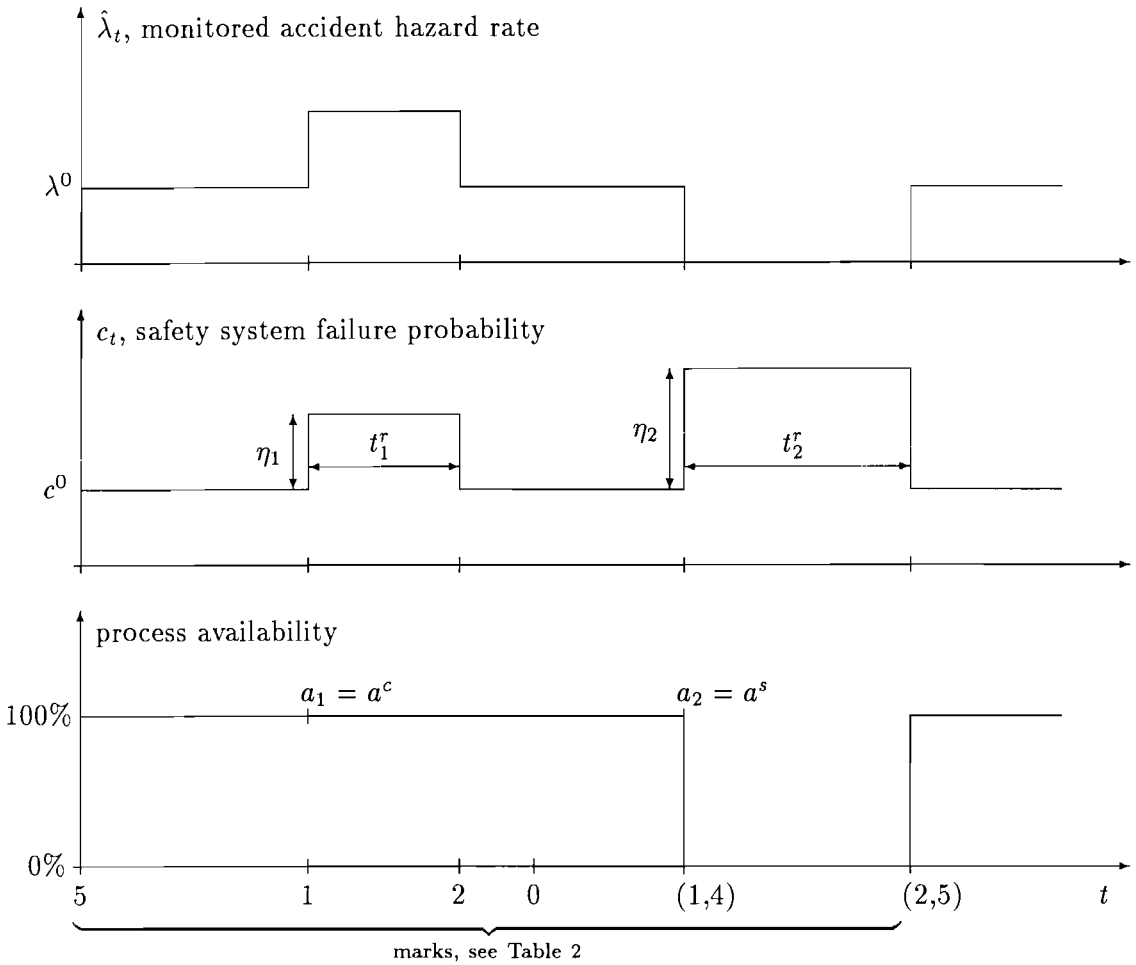

Figure 3: An example realization of the process availability and the safety system failure probability.

Table 2: Marks of the example process history.

Mark	Explanation
0	initiating event
1	degrading failure of the component
2	repair of the degraded component ends
4	planned (temporary) shutdown of the process
5	startup of the process from the shutdown state

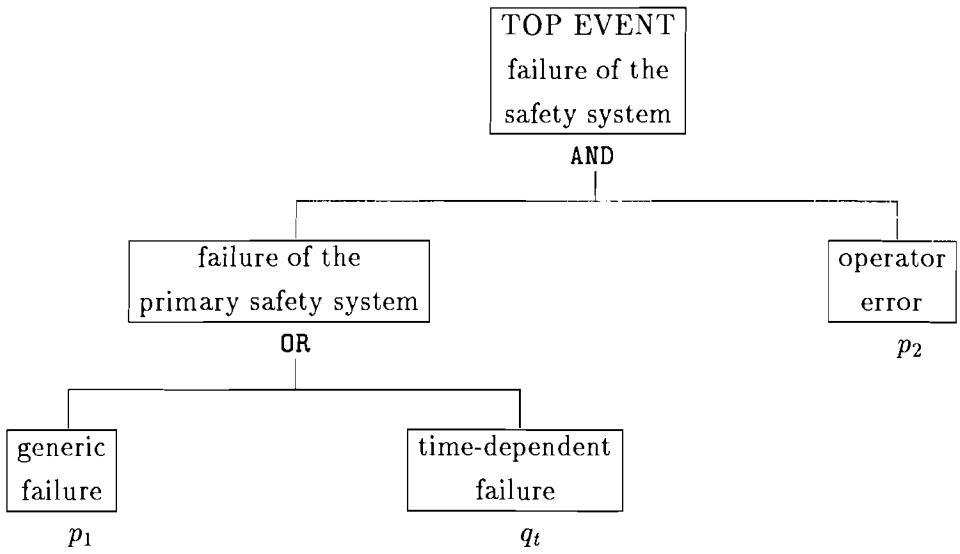

Figure 4: Fault tree of the safety system.

3.1 Process description

3.1.1 Safety system failure probability

We introduce a safety system whose probability of failure varies between $0 < c_t < 1$. In normal conditions, the system failure probability is very low, and the operation of the process is then acceptable. Occasionally, a part of the safety system is degraded, which increases the failure probability to a rather high level. In order to have this kind of cyclic reliability performance, we consider the following system.

The safety system consists of a primary safety system and a back-up operator action. The operator action has a constant failure probability denoted by p_2. Figure 4 shows the system fault tree.

The primary safety system can be unavailable for two reasons. Firstly, there are time-independent, inherent causes which can make the system inoperable. The probability of a failure by this kind of causes is p_1. Secondly, shocks occur with the intensity $\lambda(1)$. A shock degrades the reliability of the system by making some of its subsystems unavailable. The degradation degree is random denoted by η_n, $0 \leq \eta_n \leq 1$, where n is the index of the failure. The safety system failure probability is then

$$c_t = \begin{cases} p_1 p_2 & \text{if the system is at the inherent state,} \\ (p_1 + (1 - p_1)\eta_{N_t(1)})p_2 & \text{if the system is at a degraded state,} \end{cases} \tag{16}$$

We assume that the degradation degrees are independent, identically distributed uniform random variables, i.e., $\eta_n \sim U(0, 1)$.

The repair times of the degrading failures are identical, independently distributed exponential random variables with the parameter $\lambda(2)$. They are independent of degradation degrees (η-variables). The repair time and degradation degree are assumed to be known when the failure is detected.

3.1.2 Accident hazard rate

The inherent accident hazard rate is

$$\lambda^0 = \lambda(0)p_1 p_2. \tag{17}$$

Since there are no latent failures, this is also the baseline accident hazard rate. The nominal accident hazard rate depends on the chosen control strategy.

Table 3: The parameters of the model in case 1.

symbol		unit	explanation
M	$15 \cdot 10^9$	FIM [1]	costs of an accident
	1000	MWe	net electrical effect of the plant
	100	FIM/MWh	price of electricity
β	100 000	FIM/h	rate of incomes ($\approx 8 \cdot 10^8$ FIM/a) [2]
$\lambda(0)$	0.05	1/a	initiating event intensity
$\lambda(1)$	1.0	1/a	failure rate
$\lambda(2)$	100	1/a	repair rate (≈ 0.0011 1/h)
p_1	0.001		probability of a generic failure
p_2	0.1		probability of an operator error
$c(4)$	$1 \cdot 10^{-6}$		probability of core damage given shutdown
T^L	50	a	licensing time
α	$2 \cdot 10^{-10}$	1/FIM	parameter of the exponential utility function

[1] USD 1≈FIM 4.5

[2] 1 operating year=8000 hours

When a failure with a degradation degree η takes place, the instantaneous accident hazard rate increases to

$$\lambda(\eta) = (1 + \frac{1 - p_1}{p_1}\eta)\lambda^0. \tag{18}$$

The ratio

$$\frac{\lambda(\eta)}{\lambda^0} = 1 + \frac{1 - p_1}{p_1}\eta, \tag{19}$$

is called the *risk increase factor*. Originally, the risk increase factor and other risk importance measures have been defined for a static risk model (Vesely et al. 1983). The meaning of risk importance measures is to present in a relative scale how much the importance of one component is to the reliability of the system.

3.2 Problem formulation

The short-term problem is to decide whether to shut the process down in a case of failure of the component (mark '1'), i.e., $A(1) = \{a^s, a^c\}$, where a^s denotes the shutdown option and a^c the continued operation option. Because there are only two decision options at each decision epoch, the short-term optimization problem is to compare the values of the decision function $V(a^s, \hat{H}_{t_i-} \cup \{(t_i, 1)\}, x)$ and $V(a^c, \hat{H}_{t_i-} \cup \{(t_i, 1)\}, x)$, where the arguments of the decision function $V(\cdot)$ are the decision option, the (observed) operating history up to the failure epoch and a control variable x. The chosen decision can be indicated by an indicator function

$$1_{\{a_i = a^s\}} = 1_{\{V(a^s, \hat{H}_{t_i-} \cup \{(t_i,1)\}, x) > V(a^c, \hat{H}_{t_i-} \cup \{(t_i,1)\}, x)\}},$$

which receives value 1 if shutdown is considered better than continued operation. $1_{\{a_i = a^c\}}$ is defined respectively corresponding to the superiority of the continued operation option. The long-term decision problem is to choose x so that the expected lifetime utility is maximized.

The model parameters, shown in Table 3, have been chosen to correspond to a 1000 MWe nuclear power plant. $\alpha = 2 \cdot 10^{-10}$ 1/FIM means according to (12) that $\lambda^* > 8 \cdot 10^{-3}$ 1/a which is about 1600 times higher than the inherent accident hazard rate. It is also more than the maximum instantaneous accident hazard rate process can ever have.

Even this simplified case poses a complex decision problem, and it is difficult to define a long-term optimal operation strategy. Therefore we approach the problem from several perspectives.

First we will study which is in general a better strategy: to always shut the plant down in case of a failure or to always continue the operation. Secondly, we make a pure short-term decision analysis, i.e., at each failure epoch the decision is made without a long-term control. Then, we study the use of two global control variables — one limiting the instantaneous accident hazard rate and the other limiting the time of having increased risk.

3.3 Solution approaches

3.3.1 Approximation of the expected utility

If the control variable is constant, then we can approximate the expected utility by approximating the nominal accident hazard rate and average process availability given the value of the control variable. Let $\lambda^n(x)$ be the nominal accident hazard rate (6) and $\zeta(x)$ the average process availability given the long-term control variable x. Descriptively, we can define the average process availability as

$$\zeta(x) \approx \lim_{t\to\infty} \frac{S_t(x,\omega)}{t},$$

where S_t is the cumulative operation time of the process.

The outcome of the operation is then approximately

$$J(x,\omega) \approx \begin{cases} \beta\zeta(x)T^L & \text{if no accident happens,} \\ \beta\zeta(x)\bar{T} - M & \text{if an accident happens.} \end{cases}$$

The expected utility is

$$\begin{aligned} \mathbf{E}[u(J(x,\omega))] &\approx \int_0^{T^L} \lambda^n(x)e^{-\lambda^n(x)t}(1 - e^{-\alpha\beta\zeta(x)t})\,dt + e^{-\lambda^n(x)T^L}(1 - e^{-\alpha(\beta\zeta(x)+M)}) \\ &= 1 - e^{-(\lambda^n(x)+\alpha\beta)\zeta(x)T^L}(e^{-\alpha M} - \frac{\lambda^n(x)}{\lambda^n(x)+\alpha\beta}) - \frac{\lambda^n(x)}{\lambda^n(x)+\alpha\beta}. \end{aligned} \tag{20}$$

The optimal x is found by maximizing this equation.

3.3.2 The stochastic quasi-gradient algorithm

An analytic or even approximative expression of the expected life time utility can be evaluated only in very simple cases. The stochastic quasigradient algorithm is a general method to approach the optimal solution by sampling process histories and choosing the next solution based on the calculated sample gradient. The sample gradient can be evaluated in many ways, and the selection of the appropriate approach is case-dependent affecting the speed of the convergence. In this paper, we have not compared various approaches, but we have only chosen one applicable way in order to demonstrate the approach.

In the stochastic quasi-gradient algorithm, the next solution in the sequence of trial solutions, $x^0, x^1, x^2, \ldots, x^s \in X$, is obtained by

$$x^{s+1} = \Pi_X(x^s - \rho^s\xi^s), \tag{21}$$

where $\Pi_X(\cdot)$ is the orthoprojection operation on the convex set X, ρ^s is a step size and ξ^s is a stochastic quasi-gradient satisfying the following property

$$\mathbf{E}[\xi^s \mid x^0, \xi^0, x^1, \xi^1, \ldots, x^s] = \nabla_x F(x^s) + \text{ possible bias},$$

i.e., the conditional expectation of the vector ξ^s is "equal" to the gradient of the performance function $F(x)$ at the point x^s. Since the stochastic quasi-gradient method is based on the sampling of process histories, it is useful to speed up the sampling by taking some expectations. In some cases, we can even smooth the sample performance function so that we can analytically calculate the gradient for each sample. However, in our problem context, we can seldom rely

on this possibility, and the sample performance function remains discontinuous. Then, the interchange of expectation and difference operators may not be valid, and the gradient must be approximated in other ways (see e.g. Ermoliev et al. 1995). Below we outline the formulas used in the optimization algorithm of this paper.

The discontinuities of (7) are caused by the cumulative operation time and the accident time. A conditional performance function is achieved, by the introduction of a σ-algebra $\mathcal{F}^*$ belonging to the σ-algebra $\mathcal{F}$ of the probability space $(P, \mathcal{F}, \Omega)$ where all random variables are specified. We choose an accumulating σ-algebra generated by the monitored history of the safety system $\hat{H}^*_{t-}$, i.e.,

$$\mathcal{F}^*_t = \sigma(\hat{H}^*_t),$$

where

$$\hat{H}^*_t = \{(T_n, Z_n) \mid T_n \leq t, Z_n = 1, 2, 4, 5\}.$$

$\hat{H}^*_t$ includes neither the initiating event marks nor the marks indicating accidents.

The sampled process history is divided into intervals according to the failure epochs

$$t_i = \{t \mid dN_t(1) = 1, N_t(1) = i\},\ t_0 = 0,\ t_{N_{T^L}(1)} = T^L.$$

A conditional expected utility is evaluated for each failure interval $[t_i, t_{i+1}]$, $i = 0, 1, \ldots$ accounting the probability that an accident happens during the interval. The conditioning is made with respect to the safety system history up to failure epoch t_i *including* the knowledge of the degradation degree and the repair time.

The sample performance function can be expressed as a sum

$$\hat{f}(x) = \sum_{i=0}^{N_{T^L}(1)} \hat{f}_i(x \mid \hat{H}^*_{t_i-}), \tag{22}$$

where

$$\hat{f}_i(x \mid \hat{H}^*_{t_i-}) = \prod_{j=0}^{i-1} P(\bar{T} \geq t_{j+1} \mid \hat{H}^*_{t_j-} \cup \{(t_j, (1, a_j(x)))\})$$
$$\int_{s=t_i}^{t_{i+1}} \mathbf{E}[u \mid \hat{H}^*_{t_i-} \cup \{(t_i, (1, a_i(x))), (\bar{T}, \bar{Z})\}] dP(\bar{T} = s \mid \hat{H}^*_{t_i-} \cup \{(t_i, (1, a_i(x)))\}), \tag{23}$$
$$i = 1, \ldots, N_{T^L}(1) - 1,$$

$$\hat{f}_0(x) = \int_{s=0}^{t_1} \mathbf{E}[u \mid \hat{H}^*_0 \cup \{(\bar{T}, \bar{Z})\}] dP(\bar{T} = s \mid \hat{H}^*_0) \tag{24}$$

and

$$\hat{f}_{N_{T^L}(1)}(x \mid \hat{H}^*_{T^L-}) = \prod_{j=0}^{N_{T^L}(1)-1} P(\bar{T} \geq t_{j+1} \mid \hat{H}^*_{t_j-} \cup \{(t_j, (1, a_j(x)))\}) \mathbf{E}[u \mid \hat{H}^*_{T^L-}]. \tag{25}$$

This formulation allows to also consider other than discrete decision spaces, like the selection of the next test epoch. The division of the safety system history into time intervals just has to be done accordingly.

Implementing the exponential utility function (14) and the knowledge about the process described in Section 3.1., we can write the exact equations of the $\hat{f}_i$ terms, (23)–(25). Let

$$\begin{aligned} g(\mu, \tau_1, \tau_2) &= \int_0^{\tau_1} \mu e^{-\mu s} (1 - e^{-\alpha\beta(\tau_2 + s)})\, ds \\ &= 1 - e^{-\mu\tau_1} - \frac{\mu}{\mu + \alpha\beta} e^{-\alpha\beta\tau_2} (1 - e^{-(\mu + \alpha\beta)\tau_1}). \end{aligned} \tag{26}$$

Then $\hat{f}_0$ is

$$\hat{f}_0 = g(\lambda^0, t_1, 0), \tag{27}$$

and $\hat{f}_i(x)$ is

$$\hat{f}_i(x) = e^{-\lambda^0 t_i} h^i(S_{t_i}, x) \prod_{j=0}^{i-1} e^{\lambda^0 t_j^r} \left(1_{\{a_j = a^s\}} (1 - c(4)) + 1_{\{a_j = a^c\}} e^{-\lambda(\eta_j) t_j^r} \right), \quad i \geq 1 \tag{28}$$

where $c(4)$ is the probability of the accident given a shutdown, t_j^r the repair time of the failure j, η_j the degradation degree, and

$$\begin{aligned} h^i(\tau, x) = \; & 1_{\{a_i = a^s\}} \left[c(4)(1 - e^{-\alpha\beta\tau}) + (1 - c(4)) g(\lambda^0, t_{i+1} - (t_i + t^r), \tau) \right] + \\ & 1_{\{a_i = a^c\}} \left[g(\lambda(\eta_i), t^r, \tau) + e^{-\lambda(\eta_i) t^r} g(\lambda^0, t_{i+1} - (t_i + t^r), \tau + t^r) \right], \\ & i = 1, \ldots, N_{TL}(1) - 1, \end{aligned} \tag{29}$$

and

$$h^{N_{TL}(1)}(\tau, x) = e^{\alpha(-\beta\tau + M)}, \tag{30}$$

is the expected utility given the failure during the interval $[t_i, t_{i+1}]$. We assume that the probability of an accident given the shutdown, $c(4)$, is independent of the condition of the safety systems, which generally is not true. However, we made this assumption only in order to simplify analytical evaluations.

For a chosen x^s and a simulated process history ω^s, we calculate the sample performance $\hat{f}^s(x^s, \omega^s)$. Since the function is discontinuous with respect to x, we use a finite-difference gradient approximation. We can, for instance, choose a new value $\hat{x}^s$ randomly in the neighborhood of x^s, and calculate a new sample performance for $\hat{x}^s$. On the other hand, we can use the results from the previous iterations, $x^{s-1}, \ldots, x^{s-k}$.

In this study, we utilize the knowledge that x controls the number of shutdowns. We find the closest $x_-^s < x$ and $x_+^s > x$ that changes the number of shutdowns by one. By calculating the variated sample performances $\hat{f}^s(x_-^s, \omega^s)$ and $\hat{f}^s(x_+^s, \omega^s)$, we can approximate the gradient by

$$\xi_s \approx \frac{\hat{f}^s(x_+^s, \omega^s) - \hat{f}^s(x_-^s, \omega^s)}{x_+^s - x_-^s}. \tag{31}$$

If x is a vector, then the direction of the variation can be chosen randomly.

In this study, we use a decreasing step size

$$\rho_i = \frac{1}{i} \rho_0, \; i = 1, 2, \ldots \tag{32}$$

A proper ρ_0 is found by experimenting. Too large ρ_0 causes fluctuation and too small ρ_0 makes the convergence slow.

The optimal expected utility can be estimated cumulatively by

$$\hat{F}^s = \frac{1}{s} \sum_{i=1}^{s} \hat{f}^i(x^i, \omega^i). \tag{33}$$

3.4 Comparison of the two extreme strategies

The two extreme strategies are: (1) to always operate the process regardless of the condition of the safety system or (2) to always shut down when a failure occurs regardless of the severity of the failure or the repair time. By defining the global control variable $x \in \{a^s, a^c\}$, and the short-term decision function as

$$V(a, x) = \begin{cases} 1 & \text{if } a = x \\ 0 & \text{otherwise,} \end{cases}$$

the long-term decision problem has been formulated.

3.4.1 Always shut the process down during the repair time

The control variable $x = a^s$ (always a shutdown) yields an average process availability

$$\zeta(a^s) = \frac{\lambda(2)}{\lambda(1) + \lambda(1)} \approx 0.9901.$$

The nominal accident hazard rate is

$$\lambda^n(a^s) = \lambda_0 + \lambda(1)c(4) = 6.0 \cdot 10^{-6} \text{ 1/a}.$$

By (20), the expected lifetime utility is $\mathbf{E}[u(J_{TL}) \mid x = a^s] \approx 0.999944$.

3.4.2 Always continue the operation during the repair time

The control variable $x = a^c$ (always operation) yields an average process availability $\zeta(a^c) = 100\%$.

The safety system has two possible states: (0) inherent state and (1) degraded state. Since the system behave like a two-state Markov chain, the steady-state probabilities are straightforwardly $P_\infty(0) = \lambda(2)/(\lambda(1) + \lambda(2))$ and $P_\infty(1) = 1 - P_\infty(0)$. The nominal failure probability of the safety system c^n can be derived as follows

$$c^n = P_\infty(0)p_1p_2 + P_\infty(1)(p_1 + (1 - p_2)\frac{1}{2})p_2 \approx 0.000595.$$

The nominal accident hazard rate is then

$$\lambda^n(a^c) = \lambda(0)c^n \approx 2.97 \cdot 10^{-5} \text{ 1/a}.$$

The expected lifetime utility is $\mathbf{E}[u(J_{TL}) \mid x = a^c] \approx 0.999798$ which is less than $\mathbf{E}[u(J_{TL}) \mid x = a^s]$.

3.5 Pure short-term decision analysis

Next we make decisions purely on a short-term basis without a long-term control. At the time epoch t a component failure occurs, i.e., $N_t(1) = N_{t-}(1) + 1$. Let t^r denote the estimated repair time of the failure and η the degradation degree. We apply the decision function

$$V(a, J_t, t^r, \eta) = \mathbf{E}[u(J_{t+t^r}) \mid a], \; a \in A(1).$$

where J_t are the cumulative costs.

The shutdown decision option corresponds to the lottery $a^s : \langle c(4), J_t - M; 1 - c(4), J_t \rangle$. The expected utility of a^s will be

$$\mathbf{E}[u \mid a^s] = c(4)(1 - e^{-\alpha J_t}) + (1 - c(4))(1 - e^{-\alpha(J_t + M)}).$$

Given the instantaneous accident hazard rate $\lambda(\eta)$, the continued operation alternative corresponds to the lottery $a^c : \langle 1 - e^{-\lambda(\eta)t^r}, J_t - M; e^{-\lambda(\eta)t^r}, J_t + \beta t^r \rangle$, where t^r is the repair time. The expected utility will be

$$\mathbf{E}[u \mid a^c] = (1 - e^{-\lambda(\eta)t^r})(1 - e^{-\alpha J_t}) + e^{-\lambda(\eta)t^r}(1 - e^{-\alpha(J_t + \beta t^r + M)}).$$

The difference between the expected utilities is

$$\mathbf{E}[u \mid a^s] - \mathbf{E}[u \mid a^c] = e^{-\alpha J_t}\left[1 - e^{-\lambda(\eta)t^r} - c(4) - e^{-\alpha M}(1 - c(4) - e^{-(\lambda(\eta) + \alpha\beta)t^r})\right]. \quad (34)$$

The sign does not depend on J_t.

We can notice that according to the parameters of Table 3, it is most unlikely that the shutdown would be a preferrable option. Therefore we can assume that the plant is never shut down in case of a component failure as in the strategy of Section 3.4.2. The expected lifetime utility $\mathbf{E}[u(J_{TL})] \approx 0.99979$. This is *not* an optimal strategy in the long run.

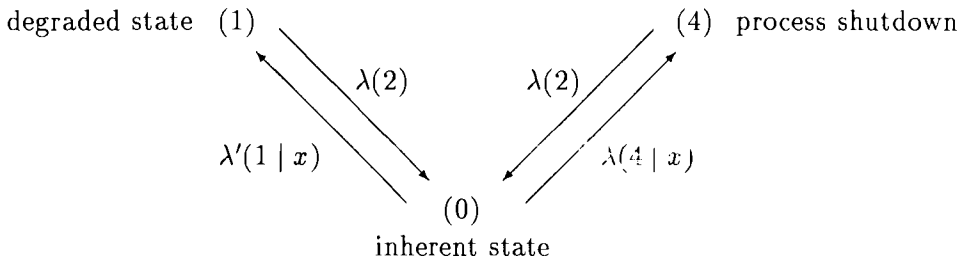

Figure 5: State diagram of the safety system.

3.6 Limited instantaneous accident hazard rate

If we limit the instantaneous accident hazard rate, we can apply as a decision function an indicator function

$$V(a, \lambda_t, x) = 1_{\{\lambda_t \leq x\lambda_0\}} 1_{\{a=a^c\}} + 1_{\{\lambda_t > x\lambda_0\}} 1_{\{a=a^s\}}, \ a \in A(1),$$

which receives value 1 for one of the decision options and 0 for the other one. The control variable x expresses here the maximum allowed risk increase factor with respect to the inherent accident hazard rate λ^0 (c.f. (19)).

3.6.1 Approximative analytical solution

A dynamic representation of the safety system can be given by a state diagram which has three states shown in Figure 5. State 0 corresponds to the inherent conditions between the end of last repair and the next failure. The probability of the failure of the safety system is then at the inherent level $c_t = c^0 = p_1 p_2$.

When a failure occurs, there are two possible transitions depending on whether the process is shut down or not. State 4 represents the shutdown option and state 1 the continued operation option. The transition rates from states 1 and 4 back to 0 are the same, i.e., the repair rate $\lambda(2)$.

Since $\eta_n \sim U(0,1)$, the probability that the failure (mark 1) increases the accident hazard rate above the level $x\lambda^0$ is

$$P(p_1 + (1-p_1)\eta_n > xp_1) = \frac{1-xp_1}{1-p_1}, \ x \in [1, 1000].$$

Therefore the hazard rate of the shutdown marks (4) is

$$\lambda(4 \mid x) = \lambda(1)\frac{1-xp_1}{1-p_1},$$

and the hazard rate of entering into degraded conditions without a shutdown is

$$\lambda'(1 \mid x) = \lambda(1)\frac{p_1(1-x)}{1-p_1}.$$

We denote the steady-state probabilities by $P_\infty(0)$, $P_\infty(1 \mid x)$, and $P_\infty(4 \mid x)$. The steady-state probability of being at the inherent state is

$$P_\infty(0) = \frac{\lambda(2)}{\lambda(1)+\lambda(2)}.$$

and in a degraded state

$$P_\infty(1 \mid x) = \frac{\lambda'(1 \mid x)}{\lambda(1)+\lambda(2)}.$$

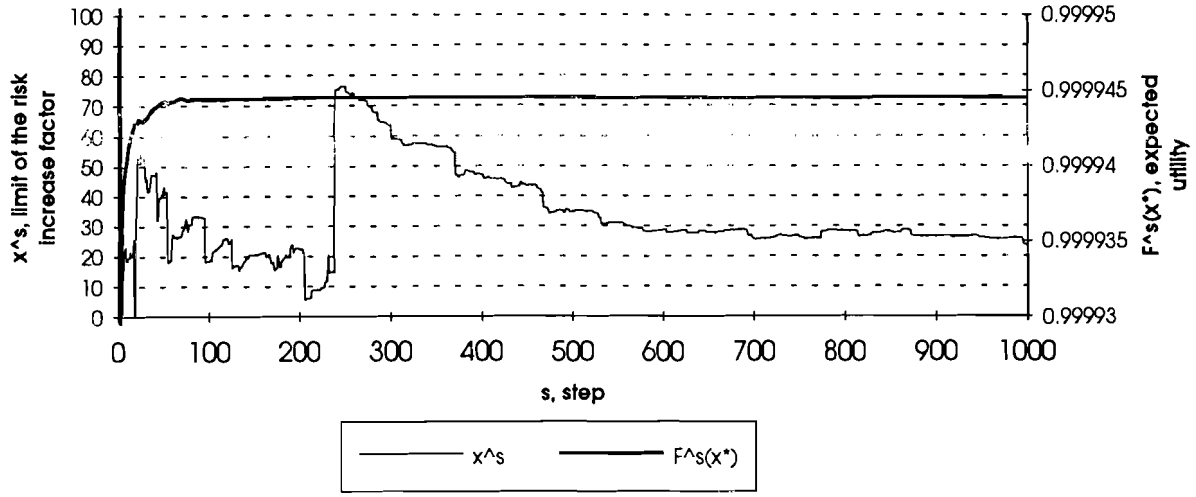

Figure 6: Optimization of the limit of the risk increase factor by the stochastic quasi-gradient algorithm.

Then the nominal safety system failure probability (when the process is operated) is

$$c^n(0 \mid x) = (p_1 + (1 - p_1)q^n(x))p_2, \tag{35}$$

where

$$q^n(x) = \frac{P_\infty(1 \mid x)}{P_\infty(0) + P_\infty(1 \mid x)} \frac{1}{2} \frac{p_1(x-1)}{1 - p_1}.$$

The nominal accident hazard rate is

$$\lambda^n(x) = \lambda(0)c^n(0 \mid x) + \lambda(4 \mid x)c(4)\frac{P_\infty(0)}{P_\infty(0) + P_\infty(1 \mid x)}, \tag{36}$$

when the process is operated and 0 when it is shut down.

We have to take into account the production losses due to shut down periods. The average availability of the process is

$$\zeta(x) = P_\infty(0) + P_\infty(1 \mid x) = \frac{\lambda(2) + \lambda'(1 \mid x)}{\lambda(1) + \lambda(2)}. \tag{37}$$

The expected utility can now be evaluated by (20). The derivative of this function is somewhat complicated, but the maximum can easily be found numerically. The optimal solution is $x^* \approx 26$ yielding nominal accident hazard rate $\lambda^n(x^*) \approx 5.99 \cdot 10^{-6}$ 1/a, process availability $\zeta(x^*) = 99.03\%$, and expected lifetime utility $E[u(J_{TL}) \mid x^*] \approx 0,9999445$ which is a little bit better than the extreme strategy to always shut down for repair times.

3.6.2 Stochastic quasi-gradient algorithm

Figure 6 shows one run of a stochastic quasi-gradient procedure. After 1000 steps, the expected utility has increased to $E[u(J_{TL}) \mid x^*] \approx 0,999944$. The value of the optimal control still keeps changing. At $s = 1000$, it is $x^{1000} \approx 24$. The results are quite comparable with our approximative results.

3.7 Limited repair time

In order to avoid long periods of being in a degraded condition, we apply as a decision function an indicator function

$$V(a, t^r, x) = 1_{\{t^r \leq x\}} 1_{\{a=a^c\}} + 1_{\{t^r > x\}} 1_{\{a=a^s\}}, \quad a \in A(1).$$

The control variable x expresses here the maximum allowed repair time.

It should be noted that, in practice, very short temporary shutdowns are not possible, and the model should be made more realistic by accounting the minimum possible shutdown period.

3.7.1 Approximative analytical solution

The conditional repair time, given that it is shorter than x, is

$$\mathbf{E}[t^r \mid t^r \leq x] = \frac{1}{\lambda(2)} - x\frac{e^{-\lambda(2)x}}{1 - e^{-\lambda(2)x}},$$

and, given that it is longer than x, is

$$\mathbf{E}[t^r \mid t^r > x] = \frac{1}{\lambda(2)} + x.$$

Since the probability of getting into an increased risk state given a failure is $1 - \lambda(2)x$, we have the following relation of the steady-state probabilities,

$$\begin{aligned} \frac{P_\infty(1|x)}{P_\infty(1|x)+P_\infty(4|x)} &= \frac{(1 - e^{-\lambda(2)x})\mathbf{E}[t^r \mid t^r \leq x]}{(1 - e^{-\lambda(2)x})\mathbf{E}[t^r \mid t^r \leq x] + e^{-\lambda(2)x}\mathbf{E}[t^r \mid t^r > x]} \\ &= 1 - e^{-\lambda(2)x} - x\lambda(2)e^{-\lambda(2)x}. \end{aligned}$$

Note that $P_\infty(0)$ is the same as before. Since $P_\infty(1 \mid x) + P_\infty(4 \mid x) = 1 - P_\infty(0)$, we have

$$P_\infty(1 \mid x) = (1 - e^{-\lambda(2)}(1 + \lambda(2)x))\frac{\lambda(1)}{\lambda(1) + \lambda(2)},$$

and

$$P_\infty(4 \mid x) = 1 - \zeta(x) = 1 - e^{-\lambda(2)}(1 + \lambda(2)x)\frac{\lambda(1)}{\lambda(1) + \lambda(2)}.$$

The rate of a shutdown is

$$\lambda(4 \mid x) = \lambda(1)e^{-\lambda(2)x},$$

and the rate of degraded conditions is

$$\lambda'(1 \mid x) = \lambda(1)(1 - e^{-\lambda(2)x}).$$

Since the degradation degrees are independent of repair times and they are not controlled in any way, we have

$$q^n(x) = \frac{1}{2}\frac{P_\infty(1 \mid x)}{P_\infty(1 \mid x) + P_\infty(2)}.$$

The nominal accident hazard rate can be calculated using (35) for $c^n(0 \mid x)$ and then (36) for $\lambda^n(x)$.

The optimal solution is $x^* \approx 3.7$ h yielding nominal accident hazard rate $\lambda^n(x^*) \approx 5.98 \cdot 10^{-6}$ 1/a, process availability $\zeta(x^*) \approx 99.01\%$, and expected lifetime utility $E[u(J_{TL}) \mid x^*] \approx 0.999945$ which is about the same as with the limited instantaneous accident hazard rate strategy.

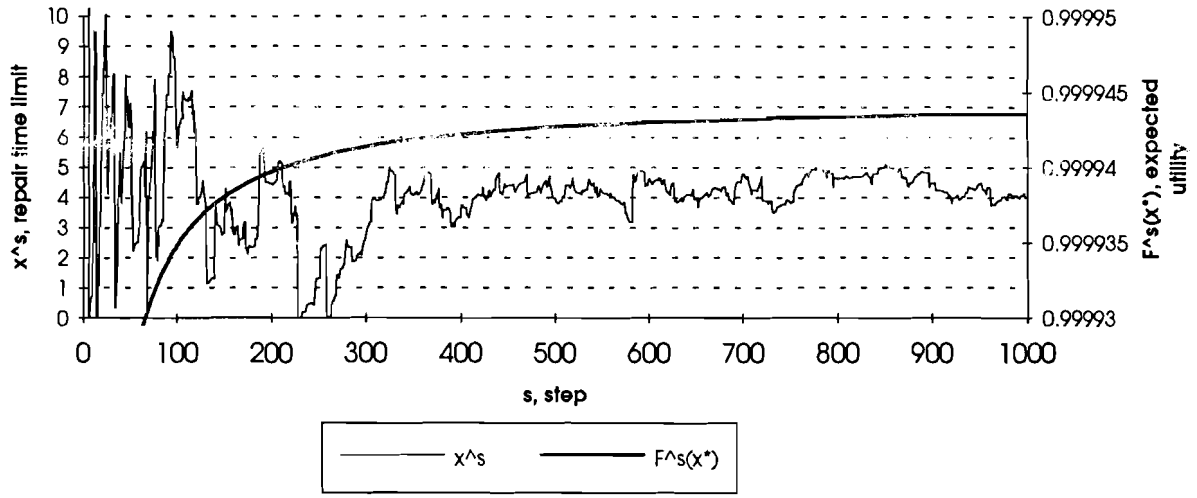

Figure 7: Optimization of the repair time limit by the stochastic quasi-gradient algorithm.

Table 4: Summary of the results.

Control	$\mathbf{E}[u(J_{TL}) \mid x^*]$	x^*	P(accident)	$\zeta(x^*)$
Always shut down	0.9999444	a^s	$2.97 \cdot 10^{-4}$	99.01%
Always continue operation	0.999797	a^c	$1.5 \cdot 10^{-3}$	100%
Limited risk increase factor	0.999945	26 [1] / 24 [2]	$2.97 \cdot 10^{-4}$	99.04%
Limited repair time	0.999945	3.6 / 4.0	$2.96 \cdot 10^{-4}$	99.01%

[1] approximative optimum

[2] estimated by the stochastic quasi-gradient method

3.7.2 Stochastic quasi-gradient algorithm

Figure 7 shows one run of our stochastic quasi-gradient procedure. As in the optimization of the risk increase factor limit, the expected utility has increased to $E[u(J_{TL}) \mid x^*] \approx 0,999944$ after 1000 steps, and the value of the optimal control still keeps changing. At $s = 1000$, it is $x^{1000} \approx 4$. The results are quite comparable with our approximative results.

3.8 Summary of the analysis

The results are summarized in Table 4. The global optimum seems to be close to the strategy of always shutting down the process during the repair time. The stochastic quasi-gradient approach and approximative, analytical equations provide similar results.

We can notice that the objective function is flat around the optimum, which makes it difficult to find and determine the optimal solution. However, we do not need to know the exact answer. In practice a risk model includes a lot of uncertainties, and it is sufficient for us to know where approximately the optimal solution is.

Presumably, a combination of the risk increase factor limit and repair time limit would improve the result. It could also be beneficial to have a time-dependent control variable. On the other hand, the simplicity of the decision rules is always a preferable feature, and the risk increase factor limit as well as the repair time limit are simple decision rules.

4 Discussion

The problem formulation presented in this paper provides a general approach to model decision problems related to stochastic processes. The model incorporates a short-term decision analysis into the optimization of a long-term objective function. Therefore, the two-level decision model can be applied to modeling problems of risk management which, from their very nature, consist of a complex mixture of inter-related problems. In this model, the event specific (short-term) decision functions are functions of the operating history and long-term control variables.

To find an optimal control for the stochastic process model is a computationally demanding problem. We have not studied various possibilities for doing it, but we suggest the use of stochastic quasi-gradient procedures.

In order to represent the risk aversion for large accidents, we have introduced a utility function in the decision model. After that, the objective function is no longer linear with respect to profits from the operation of the process. Our case study demonstrates that the strategy of optimizing problems individually, based on up-to-date knowledge without a long-term control, does not lead to an optimum.

We defined the utility function based on the acceptance of the accident hazard rate. It is a kind of operative boundary condition for the process. The exponential utility function turns out to be convenient with this approach, because by defining the acceptable hazard rate we get a condition for acceptable values of the parameter of the exponential utility function.

The next question is the definition of the cost function. This will be different if it is for the responsible company or the society. In the case of technological processes, we find it more natural to adopt the company's point of view. A difficulty remains in how to count the costs and profits over the process lifetime.

From the methodological viewpoint, to deal with several accident categories should not cause any troubles. It, however, affects the formulation of the utility function. In addition, we could take into account other than monetary outcomes and apply a multi-criteria utility function.

In the case study, we control the allowed downtimes of components important for safety. In practice, the limits of allowed downtimes for safety systems at nuclear power plants usually depend on the degree of lost redundancy. They are deterministic rules. Risk-based rules are applied only at few plants, e.g., at Heysham 2 in the United Kingdom (Horne 1991). At present, the development of dynamic risk models, called *living PSA*, for nuclear power plants (Johanson and Holmberg 1994), have initiated the discussion of risk-based rules. Most of the suggested risk-based rules limit only the nominal or instantaneous accident hazard rates (see e.g. IAEA 1991, IAEA 1993). Our decision analytic approach is based on a utility function over the plant lifetime profits and costs.

Next, the case study could be extended by assuming values for the unknown model parameters. We can define prior distributions for them and update the distributions based on the operating experience. We could apply the same kind of decision rules as in this study, but maybe it would be wiser to incorporate the operating experience in the short-term decision function. We could also study the optimization of the test intervals by assuming that the degradation failures occur latently. The latent failures could be detected by surveillance tests. In order to optimize the test interval, we should add the costs of testing the profit function. At each test or end of repair epoch, we have a decision problem of choosing the next test epoch. Then the short-term decision space, $A(Z_n)$, is continuous.

For future research, the feasibility of our approach could be studied with a proper risk model. The complexity (size) of the risk model will probably not cause limitations since it is rather easy to build a simulation model that generates sample process histories. The computational complexity depends on the number and type of decision problems we try to solve simultaneously. The problem of temporary shutdowns alone is simple, but linking this task to the test interval optimization or even to design modification problems may be too difficult for analytical approaches.

References

1. Anon. (1989). Ministry of Housing, Spatial Planning and the Environment, Premises for Risk Management; Risk Limits in the Context of Environmental Policy, Second Chamber of the States General, 1988–1989 session, 21137, no. 5, The Hague.

2. Arjas, E. and Holmberg, J. (1995). Marked point process framework for living probabilistic safety assessment and risk follow-up, *Reliability Engineering and System Safety*, Vol. 49, 59–73.

3. Bremaud, P. (1981). *Point processes and queues, martingale dynamics.* Springer-Verlag, New York.

4. Ermoliev, Y., Norkin, V.I. and Wets, R.J.B. (1995). The minimization of semicontinuous functions: mollifier subgradients. *SIAM J. Control and Optimization*, Vol. 33, No. 1, 149–167.

5. Ermoliev, Y. and Wets, R.J.B. (eds.) (1988). *Numerical Techniques for Stochastic Optimization.* Springer Verlag.

6. Holmberg, J., Johanson, G. and Niemelä, I. (1993). Risk measures in living probabilistic safety assessment, VTT Publications 146, Technical Research Centre of Finland, Espoo, 59 p. + app. 10 p.

7. Horne, B.E. (1991). The use of probabilistic safety analysis methods for planning maintenance and testing unavailabilities of essential plant at Heysham 2 AGR power station, In "Use of probabilistic safety assessment to evaluate nuclear power plant technical specifications", Report IAEA-TECDOC-599 of a Technical Committee meeting in Vienna, June 18–22, 1990, International Atomic Energy Agency, Vienna, pp. 165–175.

8. IAEA. (1989). Status, Experience and Future Prospects for the Development of Probabilistic Safety Criteria, IAEA-TECDOC-524, International Atomic Energy Agency, Vienna.

9. IAEA. (1991). Use of probabilistic safety assessment to evaluate nuclear power plant technical specifications, Report IAEA-TECDOC-599 of a Technical Committee meeting in Vienna, June 18–22, 1990, International Atomic Energy Agency, Vienna.

10. IAEA. (1993). Risk based optimization of technical specifications for operation of nuclear power plants, Report IAEA-TECDOC-729, International Atomic Energy Agency, Vienna, 145 p.

11. Johanson, G. and Holmberg, J. (eds.) (1994). Safety evaluation by living probabilistic safety assessment. Procedures and Applications for Planning of Operational Activities and Analysis of Operating Experience, SKI Report 94:2, Swedish Nuclear Power Inspectorate, Stockholm, 108 p. + app.

12. Oortmarssen van, G. and Ermoliev, Y. (1994). Stochastic Optimization of Screening Strategies for Preventing Irreversible Changes, Report WP-94-124, International Institute for Applied Systems Analysis, Laxenburg, 23 p.

13. Peroggi, G.E.G. and Wallace, W.E. (1994). Operational Risk Management: A New Paradigm for Decision Making, *IEEE Transactions of Systems, Man, and Cybernetics* Vol. 24, No. 10.

14. Pulkkinen, U. and Uryas'ev, S. (1990). Optimal operational strategies for an inspected component — statement of the problem, Report WP-90-62, International Institute for Applied Systems Analysis, Laxenburg, 22 p.

15. Vesely, W.E., Davis, T.C., Denning, R.S. and Saltos, N. (1983). Measures of risk importance and their applications, Report NUREG/CR-3385, Battelle Columbus Laboratories, Columbus, 107 p.

16. Wahlström, B. (1992). Avoiding Technological Risks. The Dilemma of Complexity. *Technological Forecasting and Social Change* 42, 351–365.