Fault Management in Communication Networks: Test Scheduling with a Risk-Sensitive Criterion and Precedence Constraints

(IEEE Conference on Decision and Control, Sidney, Australia, December 2000)

Mark A. Shayman [1]
Electrical & Computer Engineering Department
and Institute for Systems Research
University of Maryland
College Park, MD 20742
Email: shayman@eng.umd.edu

Emmanuel Fernández-Gaucherand [2]
Department of Electrical & Computer
Engineering & Computer Science
University of Cincinnati
Cincinnati, OH 45221
Email: emmanuel@ececs.uc.edu

Abstract

We consider the problem of determining the optimal sequence of tests for the discovery of a faulty component, e.g., in a telecommunications network, where there is a random cost associated with testing a component. A novel feature in our approach is that a risk-sensitive performance criterion is used in order to rank different competing schedules. We characterize optimal schedules both when the testing sequence is not subject to precedence constraints, and when it is subject to such constraints, given by an arbitrary partial order.

1 Introduction

The motivation for the work presented here comes from the problem of fault management for communication networks. Our focus is on *sequential* search techniques for fault detection and diagnosis. The main feature that differentiates our work from the available literature in this field is our choice of optimality criteria used to rank competing schedules. Traditionally, when (random) costs are associated with the fault isolation and test scheduling process, the objective is to obtain scheduling rules that optimize a measure of the *expected* aggregated costs. However, due to variability inherent in the randomess, an "optimal" schedule may produce in real-time operations a performance far inferior to the predicted average performance. Hence, we seek to obtain rules that lead to schedules exhibiting an "optimal" compromise between acceptable variance and expected costs. This is achieved via the use of an exponential utility function, leading to an optimization process that yields risk (i.e., variance) sensitive scheduling rules; see [1], [4], [5], [8], [14], [16], [27].

1.1 The Fault Management Process

We view the fault management process as a hierarchical process, consisting of two phases. In the first phase, the agent searches its domain for evidence of anomalous behavior. If evidence of such behavior is detected in a subdomain, a second phase begins in which the network element(s) in the subdomain are tested to isolate the specific fault.

The phase one problem corresponds to the following generic search problem: Given a set of n components which fail *independently*, each of which may be tested with random testing cost to determine whether or not it is faulty, determine the "optimal" order in which the components should be tested. We refer to this as the *independent fault* (INF) problem. We further assume that once a faulty component is found, the search stops; this has been called a *satisficing search* by Simon and Kadane [23]. The phase two problem corresponds to a different generic search problem: Given a set of n components *with at most one in failure*, determine the "optimal" order in which the components should be tested. We refer to this as the *mutually exclusive fault*

[1] Research supported in part by NSF grant ECS-9626399 and DoD contract MDA90497C3015
[2] Research supported in part by NSF grant NSF-INT 9602939 and by United Engineering Foundation grant 00/ER-09

(MEF) problem.

In the case where the objective is to minimize the expected sum of the testing costs, both INF and MEF are classical problems that have been investigated over several decades. Various aspects of the INF problem have been considered by (among others) Bellman [3], Kadane [11], Garey [6], Simon and Kadane [23], Kadane and Simon [12], Monma and Sidney [15]. Work related to the MEF problem has been reported by (among others) Bellman [3], Staroverov [24], Sidney [21], Stone [25].

Consider the MEF problem and suppose that the probability p_i that each component i is faulty is known–from prior knowledge of the properties of the system together with any observations made before the commencement of testing. Each component i can be tested at a random cost C_i with a binary-valued result indicating whether or not the component is faulty. Suppose that $\{C_i\}_{i=1}^n$ are independent nonnegative random variables and are independent of $\{I_{F_i}\}_{i=1}^n$, where I_{F_i} is the indicator function of fault F_i. (Note that this assumes that the cost of testing component i is independent of whether or not it is faulty.) The components are tested in some sequential order until the (unique) faulty component is discovered at which time the testing terminates. Suppose the goal is to determine the order in which the tests are performed to minimize the expected value of the sum of the testing costs. Since any test for which the cost is 0 (w.p.1) would trivially be scheduled before any other test, there is no loss of generality in excluding the possibility that $C_i = 0$ (w.p.1) for any i. If there are no precedence constraints, an interchange argument [4] can be used to show that it is optimal to schedule the tests in order of increasing ratios $E[C_i]/p_i$. This has been referred to in the literature as the *C/P algorithm* [13].

For communication network fault management, Huard and Lazar [9] have proposed using the the C/P algorithm in conjunction with Bayesian belief networks to determine testing strategies for network fault isolation. The belief networks represent domain knowledge and are used to determine the failure probabilities p_i; see also [10]. (They also discuss the use of dynamic programming for models in which the C/P-rule is not applicable.) Baras, Li, and Mykoniatis [2] have proposed a heuristic based upon which the C/P algorithm can be adapted to the case of multiple faults.

In this paper we focus on the MEF problem. We summarize some results for the situation with no precedence constraints, and then give a more detailed presentation of the situation with precedence constraints. Further results, including a complete treatment of the INF problem, can be found in [18], [19], [20].

1.2 Optimality Criteria

In the above approaches, the objective is to minimize the average sum of the testing costs. This may make sense for diagnostic problems that will be repeated many times under the same conditions–i.e., with the same model–such as the diagnosis of engine failures in a particular brand of motorcycle considered in [13]. This is unlikely to be the case for the diagnosis of communication network faults. The above concerns lead us to consider the problem of optimal sequential fault diagnosis with a *risk-sensitive optimality criterion*. This is achieved via the use of an exponential utility function, leading to an optimization process that yields risk (i.e., variance) sensitive scheduling rules; see [1], [4], [5], [8], [14], [16], [27]. As mentioned previously, it is this choice of optimality criteria that differentiates our work from the available literature.

We therefore assume that risk sensitivity is given by an exponential (dis-) utility function $U_\gamma(x) = (\mathrm{sgn}\gamma)e^{\gamma x}$, $\gamma \in \Re$, $\gamma \neq 0$. The above utility function can be explained as modeling the behavior of a "decision maker" with a constant level of risk aversion, as given by the parameter γ [17], [14]. Another viewpoint is just to consider a Taylor expansion of $E[e^{\gamma \mathcal{C}}]$ about $E[\mathcal{C}]$, where $\mathcal{C}$ represents the sum of aggregated random costs. Up to a first approximation, the latter leads to an objective for optimization composed of the sum of expected costs (the standard or *risk-null* criterion) and the *variance of* $\mathcal{C}$ multiplied by the risk sensitivity parameter γ. If $\gamma > 0$ the decision maker is said to be *risk-averse* in that the objective for the optimization seeks to minimize both variance (uncertainty measuring risks involved) and the expected costs. If $\gamma < 0$, then variance is actually seen as a desirable feature, and the decision maker is said to be *risk-seeking*. Finally, if $\gamma = 0$, the standard expected costs criterion is recovered, which cannot distinguish between two scheduling rules with different variances but equal expected costs, a risk insensitive or *risk-null* situation.

In the sequel, we provide an explicit characterization of the optimal testing sequence in the unconstrained case. In the more general case where the sequence is constrained by an arbitrary partial order, we show that our model leads to modular decompositions of the optimality criterion, and therefore is amenable to study following work by Monma and Sidney [15]. At the root of our results is an interchange argument. That such an argument proves useful is somewhat surprising since scheduling problems that can be solved by interchange arguments in the risk-neutral case are not necessarily amenable to such solution in the case of risk-sensitive criterion [1].

2 Mutually Exclusive Faults Without Precedence Constraints

We begin by summarizing some results for the problem of risk-sensitive sequential diagnosis for MEF without precedence constraints. A complete presentation can be found in [18], [19], [20]. Let C_i be a random variable representing the cost of testing component i given that it is faulty, and let D_i be a random variable representing the cost of testing component i given that it is not faulty.

The exponential utility function used in our risk-sensitive optimality criterion also has some relation to scheduling problems which involve minimizing completion times, and when (exponential) discounting is employed, as illustrated in the following remark.

Remark 1 In the special case where $C_i = D_i$, $\forall i$ and $\gamma < 0$ (risk-seeking decision maker), the MEF problem without precedence constraints is mathematically equivalent to the total weighted exponential completion time problem [7, 15]. To obtain this equivalence, p_i and C_i are redefined to be the weight and processing time, respectively, for job i.

Let $\psi(i;\gamma) := E[e^{\gamma C_i}]$, and let $\overline{\psi}(i;\gamma) := E[e^{\gamma D_i}]$. Thus, $\psi(i;\gamma)$ and $\overline{\psi}(i;\gamma)$ are the moment generating functions of C_i and D_i, respectively. In the sequel, we generally suppress the dependence on γ.

Let $a = (a_1, \ldots, a_m)$ be a permutation of an m-element subset, denoted $\{a\}$, of $\{1, \ldots, n\}$. We refer to a as a schedule and indicate the number of elements in the schedule by $|a|$. If $|a| = n$, we refer to a as a complete schedule. For two (disjoint) schedules $a = (a_1, \ldots, a_m)$ and $b = (b_1, \ldots, b_\ell)$ such that $m + \ell \leq n$, we denote the concatenated schedule as

$$ab := (a_1, \ldots, a_m, b_1, \ldots, b_\ell).$$

To simplify notation, we will generally assume that $\gamma > 0$–i.e., the decision maker is risk-averse. However, we will indicate the analogous results for the case $\gamma < 0$ as well. Denote the random cost as

$$\mathcal{C}_i = I_{F_i} C_i + (1 - I_{F_i}) D_i,$$

and for $a = (a_1, \ldots, a_m)$, the utility of the aggregated cost is computed as

$$\mathcal{C}(a) = exp\{\gamma \sum_{i=1}^{m} \mathcal{C}_{a_i}\}. \tag{1}$$

In the event that $I_{F_{a_i}} = 1$, i.e., the $a_i th$ component is faulty, then

$$E[\mathcal{C}(a) | I_{F_{a_i}} = 1] = \psi(a_i) \prod_{j=1}^{i-1} \overline{\psi}(a_j). \tag{2}$$

Hence, for $|a| = m > 1$, the expected exponential risk-sensitive cost $V(a)$ for the schedule a is given by

$$V(a) = \sum_{i=1}^{m} p_{a_i} \psi(a_i) \prod_{j=1}^{i-1} \overline{\psi}(a_j) + (1 - \sum_{i=1}^{m} p_{a_i}) \prod_{i=1}^{m} \overline{\psi}(a_i), \tag{3}$$

where the first term comes from (2) and the independence assumptions, and the second term accounts for the probability that none of the $|a| = m$ tested components are faulty. In the special case where $|a| = 1$ and $a = (i)$, we have

$$V(i) = p_i \psi(i) + (1 - p_i) \overline{\psi}(i). \tag{4}$$

The following results were proved in [18], [19], [20]. See also the above references for an analysis of cases with small and large values of γ.

Lemma 1 *Suppose that $\{a\} = \{b\}$. Then $V(a) \leq V(b)$ if and only if*

$$\sum_{i=1}^{m} p_{a_i} \psi(a_i) \prod_{j=1}^{i-1} \overline{\psi}(a_j) \leq \sum_{i=1}^{m} p_{b_i} \psi(b_i) \prod_{j=1}^{i-1} \overline{\psi}(b_j). \tag{5}$$

Let

$$\hat{V}(a) := \sum_{i=1}^{m} p_{a_i} \psi(a_i) \prod_{j=1}^{i-1} \overline{\psi}(a_j). \tag{6}$$

As a consequence of Lemma 1, when comparing schedules containing the same set of elements, it is equivalent to base the comparison on the function $\hat{V}$ rather than on V.

Let $\overline{\psi}(a) := \prod_{i=1}^{|a|} \overline{\psi}(a_i)$. It follows easily from the definitions that

$$\overline{\psi}(ab) = \overline{\psi}(a) \overline{\psi}(b), \tag{7}$$

$$\hat{V}(ab) = \hat{V}(a) + \overline{\psi}(a) \hat{V}(b). \tag{8}$$

The following proposition contains the key interchange argument used to obtain the optimal scheduling rules, see [18], [19], [20].

Proposition 1 *If $|i| = |j| = 1$, then*

$$\hat{V}(aijd) \leq \hat{V}(ajid) \Leftrightarrow \frac{\overline{\psi}(i) - 1}{p_i \psi(i)} \leq \frac{\overline{\psi}(j) - 1}{p_j \psi(j)}. \tag{9}$$

Theorem 1 *Let $\gamma \neq 0$. For the MEF problem with no precedence constraints, a complete schedule $t = (t_1, \ldots, t_n)$ is optimal if and only if*

$$\frac{\mathrm{sgn}(\gamma)(\overline{\psi}(t_i) - 1)}{p_{t_i} \psi(t_i)} \leq \frac{\mathrm{sgn}(\gamma)(\overline{\psi}(t_j) - 1)}{p_{t_j} \psi(t_j)}, \tag{10}$$

whenever $i < j$.

3 Mutually Exclusive Faults With Precedence Constraints

Next we extend the results in the previous section to include precedence constraints. Our key result is that we are able to show a decomposition result for our performance objective, and hence we demonstrate that the modular decomposition theory of Monma and Sidney [15] can be applied to our problem. Decomposition results for various scheduling problems under precedence constraints have been described by a number of authors, including (among others) Sidney [21, 22], Kadane and Simon [12], and Glazebrook and Gittens [7].

We begin by recalling the relevant definitions and results from [15]. Let $S = \{1, \ldots, n\}$ be a set of jobs to be sequenced. Let f be a *cost function* that assigns a cost $f(u)$ to the schedule u. The precedence relation constraints are represented by a (directed) *precedence graph* $G = (S, R)$, where the nodes in S correspond to the jobs and an arc $(i, j) \in R$ corresponds to the precedence constraint, denoted $i \rightarrow j$, that job i must precede job j. A complete schedule t is *feasible* if it is consistent with R; it is *optimal* if it minimizes $f(\cdot)$ over the set of all feasible complete schedules. Precedence subgraphs $G_1 = (S_1, R_1)$ and $G_2 = (S_2, R_2)$, with S_1, S_2 disjoint, are *in parallel* if whenever $i \in G_1$ and $j \in G_2$, then $(i, j), (j, i) \notin R$. An *initial* (*terminal*) subset of (S, R) is a subset $S' \subseteq S$ such that if $i \in S'$, then every predecessor (successor) of i is also in S'.

The definitions below provide an extension of the interchange argument approach used in the previous section (see [18], [19], [20]), to the context in [15]. The sequencing function f is said to have the *strong adjacent sequence interchange (strong ASI) property* if there exists a *transitive* "preference" relation $\preceq$ defined on all pairs of schedules such that

$$b \preceq c \Leftrightarrow f(abcd) \leq f(acbd), \ \forall a, d. \tag{11}$$

The sequencing function f is said to have the *strong series network decomposition (strong SND) property* if whenever $\{b\} = \{c\}$, then

$$f(b) \leq f(c) \Rightarrow f(abd) \leq f(acd), \ \forall a, d. \tag{12}$$

$(f, \preceq)$ is said to have the *consistency property* if whenever $\{b\} = \{c\}$, then

$$f(b) \leq f(c) \Rightarrow b \preceq c. \tag{13}$$

The preference relation $\preceq$ is extended to sets by saying that $B \preceq C$ if and only if $b \preceq c$, where b and c are the minimum cost feasible permutations of B and C, respectively. The set $B \subseteq S$ is called *p-minimal* in $G = (S, R)$ if B is a nonempty initial set in G and $B \preceq C$ for every nonempty initial set C in G. A p-minimal set for which no proper subset is p-minimal is called *p^*-minimal.*

Monma and Sidney [15] presented the following algorithm.

Decomposition Algorithm:

Step 0 (Initialize) Set v to be the empty permutation.

Step 1 (Decompose) (a) If $S = \emptyset$ then stop; v is an optimal permutation. (b) Else, find a p^*-minimal set U in $G = (S, R)$.

Step 2 (Sequence) Find an optimal permutation u for U. Set $v := vu$. Delete U from G and go to Step 1.

Theorem 2 *[15] Assume that the strong ASI, strong SND, and consistency properties hold. Then the Decomposition Algorithm produces only optimal permutations, and every optimal permutation can be produced by the Decomposition Theorem.*

Suppose that we use $\hat{V}$ as our sequencing function. It follows that for the strong ASI property to hold, the preference relation would need to be defined so that

$$b \preceq c \Leftrightarrow \frac{\overline{\psi}(b) - 1}{\hat{V}(b)} \leq \frac{\overline{\psi}(c) - 1}{\hat{V}(c)}. \tag{14}$$

However, if the preference relation is so defined, then the conditions $\hat{V}(b) \leq \hat{V}(c)$, $\{b\} = \{c\}$ would imply that $c \preceq b$. This is the opposite of what is required for the consistency property (13) to hold.

Thus, we replace the sequencing function $\hat{V}$ with the sequencing function $\hat{V}'$ defined as follows: For any schedule $a = (a_1, \ldots, a_m)$, define the *reverse* of a, denoted a', by $a' = (a_m, \ldots, a_1)$. Then define

$$\hat{V}'(a) = \hat{V}(a'). \tag{15}$$

Also, given a precedence graph $G = (S, R)$, let $G' = (S, R')$ be the *reverse* of G–i.e., $(i, j) \in R \Leftrightarrow (j, i) \in R'$. Then a complete schedule t satisfies the precedence constraints specified by the precedence graph G if and only if t' satisfies the precedence constraints specified by the reversed precedence graph G'. These definitions lead to the following equivalency result, the proof of which is immediate and therefore omited.

Lemma 2 *An schedule t is the optimal solution to the sequencing problem associated with the sequencing function $\hat{V}$ and precedence graph G if and only if t' is the optimal solution to the problem associated with the sequencing function $\hat{V}'$ and precedence graph G'.*

Thus, from the above result, it suffices to work with $\hat{V}'$ instead of $\hat{V}$. We have that

$$\hat{V}'(ab) = \hat{V}'(b) + \overline{\psi}(b)\hat{V}'(a), \tag{16}$$

$$\hat{V}'(abcd) - \hat{V}'(acbd) = \{\hat{V}'(b)[\overline{\psi}(c) - 1] - \hat{V}'(c)[\overline{\psi}(b) - 1]\}\overline{\psi}(d). \tag{17}$$

It follows that

$$\hat{V}'(abcd) \le \hat{V}'(acbd) \Leftrightarrow \frac{\overline{\psi}(c) - 1}{\hat{V}'(c)} \le \frac{\overline{\psi}(b) - 1}{\hat{V}'(b)}. \tag{18}$$

Thus, the strong ASI property will hold provided we use $\hat{V}'$ as our sequencing function, and we define the preference relation by

$$b \preceq c \Leftrightarrow \frac{\overline{\psi}(c) - 1}{\hat{V}'(c)} \le \frac{\overline{\psi}(b) - 1}{\hat{V}'(b)}. \tag{19}$$

It is straightforward to verify that with this definition, the strong SND property and the consistency property also hold. Note that the numerators in (19) have the same sign as that of the risk-sensitivity parameter γ.

Proposition 2 *For the sequencing function $\hat{V}'$ and associated preference relation (19), each of the strong ASI, strong SND, and consistency properties hold.*

¿From Theorem 2 and Proposition 2 we obtain the following result.

Theorem 3 *For the MEF problem with $\gamma \neq 0$ and precedence graph G, a complete schedule t is optimal if and only if it is the reverse of a schedule produced by the Decomposition Algorithm with precedence graph G' and preference relation (19).*

A more explicit description of the optimal schedule(s) can be obtained by recalling another result of Monma and Sidney. By [15, Lemma 8], the set

$E(G) := \{E : E$ occurs in some sequence of p^*-minimal sets for $G\} = \{E_1, \ldots, E_r\}$,

is well-defined, and is independent of the actual implementation of Step 1 of the Decomposition Algorithm. From this, the following characterization of optimal schedules is obtained.

Theorem 4 *[15] Assume that the strong ASI, strong SND, and consistency properties hold. Then a complete schedule t is optimal for G if and only if t is of the form $(t|E_{\delta(1)}, \ldots, t|E_{\delta(r)})$, where δ is a permutation of $\{1, \ldots, r\}$ and*

(a) *$t|E_k$ is optimal for E_k, $k = 1, \ldots, r$.*

(b) *If $i \to j$ for some $i \in E_p$ and $j \in E_q$, then $\delta^{-1}(p) < \delta^{-1}(q)$.*

(c) *If E_p and E_q are in parallel and $E_p \preceq E_q$, then $\delta^{-1}(p) < \delta^{-1}(q)$.*

Applying Theorem 4 to the MEF problem gives the following result.

Theorem 5 *For the MEF problem with $\gamma \neq 0$ and precedence graph G, let $E(G') := \{E_1, \ldots, E_r\}$ be the collection of p^*-minimal sets generated by the Decomposition Algorithm with constraint relation G' and preference relation (19). A complete schedule t is optimal if and only if its reverse schedule t' is of the form $(t'|E_{\delta(1)}, \ldots, t'|E_{\delta(r)})$, where δ is a permutation of $\{1, \ldots, r\}$ and*

(a) *$t'|E_k$ is optimal for E_k, $k = 1, \ldots, r$.*

(b) *If $i \to j$ (in G') for some $i \in E_p$ and $j \in E_q$, then $\delta^{-1}(p) < \delta^{-1}(q)$.*

(c) *If E_p and E_q are in parallel and $E_p \preceq E_q$, then $\delta^{-1}(p) < \delta^{-1}(q)$.*

Example 1 ([18], [19], [20]) Suppose there are three components with fault probabilities $p_1 = 0.3$, $p_2 = 0.4$, $p_3 = 0.2$. Note that there remains a probability of 0.1 that the failure is not in any of the three suspected components. Suppose that the costs of negative tests are $D_1 = 2$, $D_2 = 3$, $D_3 = 2$ and the costs of positive tests are $C_1 = 1$, $C_2 = 2$, $C_3 = 3$. Thus negative tests are more costly than positive tests for components 1 and 2, but the opposite is true for component 3. Thus, it is more costly to exonerate components 1 and 2 than it is to confirm failure and repair them; the opposite is true for component 3. For example, component 3 might have a substantial repair cost.

Since $D_1/p_1 = 6.67$, $D_2/p_2 = 7.50$, $D_3/p_3 = 10.00$, it follows from Corollary 1 in [19], [20], that for small $|\gamma|$, the optimal schedule is $(1, 2, 3)$. In particular this is the case in the limit as $\gamma \to 0$–i.e., the risk-neutral case. Note that the costs associated with positive tests, and hence repair costs, have no bearing on this conclusion.

Since, $c_1 < c_2 < c_3$, it follows from Corollary 3 in [19], [20], that for sufficiently negative values of γ–i.e., for a sufficiently risk-seeking decision maker–the optimal schedule is $(1, 2, 3)$. In fact, numerical investigation confirms that this is the optimal schedule for all $\gamma \le 0$.

Since $d_3 - c_3 < d_1 - c_1 = d_2 - c_2$ and $p_1 < p_2$, it follows from Corollary 2 in [19], [20], that for sufficiently positive values of γ–i.e., for a sufficiently risk-averse decision maker–the optimal schedule is $(3, 2, 1)$. In fact, numerical investigation based on Theorem 1 reveals that the optimal schedule is $(1,2,3)$ for $-\infty < \gamma \leq 0.1869$, $(1,3,2)$ for $[0.1870, 0.2027]$, $(3,1,2)$ for $[0.2028, 0.2644]$, and $(3,2,1)$ for $[0.2645, \infty)$.

Suppose now we add the precedence constraint $2 \rightarrow 3$– i.e., that component 2 must be tested prior to component 3. Obviously, the schedule $(1,2,3)$ remains optimal for $-\infty < \gamma < 0.1869$. However, it is not obvious which schedule is optimal for larger values of γ since none of $(1,3,2)$, $(3,1,2)$, $(3,2,1)$ are feasible. A direct numerical investigation reveals that $(1,2,3)$ is in fact optimal for $-\infty < \gamma < 0.2119$, and $(2,3,1)$ is optimal for $0.2120 \leq \gamma < \infty$.

References

[1] G. Avila-Godoy and E. Fernandez-Gaucherand. Exponential risk-sensitive optimal scheduling. In *Proceedings of the 36th CDC*, pages 3958–3963, San Diego, CA, December 1997.

[2] J. Baras, H. Li, and G. Mykoniatis. Integrated, distributed fault management for communication networks. Technical Report 98-10, Center for Satellite and Hybrid Communication Networks, University of Maryland, College Park, MD, 1998.

[3] R. Bellman. *Dynamic Programming.* Princeton University Press, Princeton, New Jersey, 1957.

[4] D. Bertsekas. *Dynamic Programming and Optimal Control, Volume 1.* Athena Scientific, Belmont, MA, 1995.

[5] R. Cavazos - Cadena and E. Fernández - Gaucherand. Controlled Markov chains with risk–sensitive criteria: average cost, optimality equations, and optimal solutions. *Mathematical Methods of Operations Research*, 49:299–324, 1999.

[6] M. R. Garey. Optimal sequencing with precedence constraints. *Discrete Mathematics*, 4:37–56, 1973.

[7] K.D. Glazebrook and J.C. Gittens. On single-machine scheduling under precedence constraints and linear or discounted costs. Technical report, Mathematical Institute, University of Oxford, 1979.

[8] A. R. Howard, and J. E. Matheson. Risk–sensitive Markov decision processes. *Management Sciences*, 18:356–369, 1972.

[9] J-F. Huard and A.A. Lazar. Fault isolation based on decision-theoretic troubleshooting. Technical Report 442-96-08, Center for Telecommunications Research, Columbia University, New York, NY, 1996.

[10] J-F. Huard. Probabilistic reasoning for fault management on XUNET. Technical Report, Center for Telecommunications Research, Columbia University, New York, NY, 1994.

[11] J. B. Kadane. Quiz show problems. *Journal of Mathematical Analysis and Applications*, 27:609–623, 1969.

[12] J. B. Kadane and H. A. Simon. Optimal strategies for a class of constrained sequential problems. *Annals of Statistics*, 5(2):237–255, 1977.

[13] J. Kalagnanam and M. Henrion. A comparison of decision analysis and expert rules for sequential diagnosis. In *Uncertainty in Artificial Intelligence 4*, pages 271–281. Elsevier, 1990.

[14] S. I. Marcus, E. Fernández-Gaucherand, D. Hernández-Hernández, S. Coraluppi and P. Fard. Risk sensitive Markov decision processes. In *Systems & Control in the Twenty-First Century.* Series: Progress in Systems and Control, Birkhäuser. Editors: C.I. Byrnes, B.N. Datta, D.S. Gilliam, C.F. Martin, 263–279, 1996.

[15] C. L. Monma and J. B. Sidney. Optimal sequencing via modular decomposition: Characterization of sequencing functions. *Mathematics of Operations Research*, 12(1):22–31, 1987.

[16] M. Pinedo *Scheduling: Theory, Algorithms, and Systems. Prentice Hall*, Englewood Cliffs, N.J., 1995.

[17] J. W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, 1964.

[18] M. A. Shayman and E. Fernández-Gaucherand. Risk–Sensitive decision theoretic troubleshooting. *Proc. 37th Annual Allerton Conference on Communication, Control, and Computing,* September 22-24, 1999.

[19] M. A. Shayman and E. Fernández-Gaucherand. Risk–sensitive decision theoretic troubleshooting. Technical Report, Institute for Systems Research, University of Maryland, College Park, 2000.

[20] M. A. Shayman and E. Fernández-Gaucherand. Risk–sensitive decision theoretic diagnosis. Submitted for publication, April 2000.

[21] J. B. Sidney. Decomposition algorithms for single-machine sequencing with preference relations and deferral costs. *Operations Research*, 23:283–298, 1975.

[22] J. B. Sidney. A decomposition algorithm for sequencing with general precedence constraints. *Mathematics of Operations Research*, 6:190–204, 1981.

[23] H. A. Simon and J. B. Kadane. Optimal problem-solving search: All-or-none solutions. *Artificial Intelligence*, 6:235–247, 1975.

[24] O. Staroverov. On a searching problem. *Theory Probability Appl.*, 8:184–187, 1963.

[25] L. D. Stone. *Theory of Optimal Search.* Academic Press, New York, 1975.

[26] A. Tanenbaum. *Computer Networks.* Prentice Hall PTR, Upper Saddle River, NJ, 1996.

[27] P. Whittle Risk–sensitive Optimal Control. *Wiley*, New York, 1990.