Risk-Sensitive Decision-Theoretic Troubleshooting

Mark A. Shayman *
Electrical & Computer Eng. Dept. and
Institute for Systems Research
University of Maryland
College Park, MD 20742
shayman@eng.umd.edu

Emmanuel Fernández-Gaucherand †
Systems & Industrial Engineering Dept.
University of Arizona
Tucson, AZ 85721
emmanuel@sie.arizona.edu

Abstract

We study the problem of sequencing 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.

1 Introduction

Consider the problem of fault management for communication networks. As networks increase in size and complexity, the frequency of faults increases. At the same time, the tremendous increase in network speed means that the cost per unit time of failures is much greater than in the recent past. Network faults may be difficult to predict or diagnose, and may result in huge costs.

Our focus in this paper is on *sequential* search techniques for fault detection and diagnosis. Thus, it is concerned with the optimal scheduling of monitoring/testing tasks for a single intelligent agent. 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 [2], [5], [7], [15], [23], [27], [36].

*Research Supported in part by NSF grant ECS-9626399 and DoD contract MDA90497C3015
†Research supported in part by NSF grant INT-9602939 and by United Engineering Foundation grant 00/ER-09

1.1 The Fault Management Process

The fault management process can be viewed 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 following generic search problem corresponds to the phase one 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. 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* (MEF) problem.

When 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 [4], [25] Dean [8], Kadane [20], Joyce [19], Garey [10], Simon and Kadane [32], Kadane and Simon [21], Monma and Sidney [26]. Work related to the MEF problem has been reported by (among others) Bellman [4], Staroverov [33], Greenberg [12], Matula [24], Black [6], Denby [9], Horn [14], Adolphson and Hu [1], Sidney [30], Stone [34], Hall [13].

In this paper we concentrate on the MEF problem, when there are no constraints on the admissible testing sequence. A more comprehensive study is presented elsewhere [29], where both the INF and the MEF problems are treated, with and without constraints on the admissible testing sequence.

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 [5] 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* [22]. To motivate our presentation in the sequel, we proceed with a simple derivation of the C/P algorithm: Let $q_i := 1 - p_i$, and assume that $(t_1, t_2, \ldots, t_k, i, j, \ldots, t_n)$ is an optimal schedule; the expected cost associated with this schedule is

$$E[C_{t_1} + q_{t_1} C_{t_2} + \ldots + (\Pi_{\ell=1}^k q_{t_\ell}) \cdot C_i + (\Pi_{\ell=1}^k q_{t_\ell}) \cdot q_i C_j + \ldots].$$

Comparing the optimal (minimal) cost above with that yielded by the (suboptimal) schedule obtained by interchanging the i and j tests, we obtain the necessary condition for optimality:

$$E[C_i] + q_i E[C_j] \leq E[C_j] + q_j E[C_i],$$

which is equivalent to the C/P rule:

$$\frac{E[C_i]}{p_i} \leq \frac{E[C_j]}{p_j}.$$

It is also easy to conclude that an optimal policy for this problem is an (open-loop) schedule, and hence the C/P rule is actually necessary and sufficient for optimality [5].

Huard and Lazar [17] have proposed using the the C/P algorithm in conjunction with Bayesian belief networks to determine testing strategies for communication network fault isolation. The belief networks represent domain knowledge and are used to determine the failure probabilities p_i; see also [16]. (They also discuss the use of dynamic programming for models in which the C/P-rule is not applicable.) Baras, Li, and Mykoniatis [3] have proposed a heuristic based upon which the C/P algorithm can be adapted to the case of multiple faults.

1.2 Optimality Criteria

In the literature mentioned previously, 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. This is unlikely to be the case for the diagnosis of communication network faults. This concern leads 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 [2], [5], [15], [23], [27], [36].

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 [28], [23]. 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. Hence, 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, for the MEF problem without constraints, we are able to obtain explicit characterization of the optimal testing sequence in the unconstrained case. 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 [2].

2 Mutually Exclusive Faults Without Precedence Constraints

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 relationship between C_i and D_i–e.g., the relative magnitudes of $E[C_i]$ and $E[D_i]$ is application-dependent. We assume that conditional on the state of component i, the testing cost of component i is independent of the states and testing costs of all other components. We allow for the possibility that the true fault may not be one of the n hypothesized faults; thus $\sum_{i=1}^{n} p_i$ may be less than 1. If there is a faulty component, the testing sequence terminates when that component is tested; otherwise, every component is tested.

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 [11, 26]. 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

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

If we want to emphasize the dependence on the risk-sensitivity parameter γ, we will use the notation $V(a;\gamma)$ instead of $V(a)$.

The following result plays a crucial role in the sequel.

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

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

Proof: The result follows directly from (3). Let $m = |a| = |b|$. Since $\{a\} = \{b\}$,

$$(1 - \sum_{i=1}^{m} p_{a_i})\prod_{i=1}^{m}\overline{\psi}(a_i) = (1 - \sum_{i=1}^{m} p_{b_i})\prod_{i=1}^{m}\overline{\psi}(b_i). \tag{6}$$

■

Let

$$\hat{V}(a) := \sum_{i=1}^{m} p_{a_i}\psi(a_i)\prod_{j=1}^{i-1}\overline{\psi}(a_j). \tag{7}$$

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. This greatly simplifies the analysis.

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

$$\overline{\psi}(ab) = \overline{\psi}(a)\overline{\psi}(b), \tag{8}$$
$$\hat{V}(ab) = \hat{V}(a) + \overline{\psi}(a)\hat{V}(b). \tag{9}$$

The following proposition contains the key interchange argument used subsequently to obtain the optimal scheduling rules.

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

$$\hat{V}(aijd) \le \hat{V}(ajid) \Leftrightarrow \frac{\overline{\psi}(i) - 1}{p_i\psi(i)} \le \frac{\overline{\psi}(j) - 1}{p_j\psi(j)}. \tag{10}$$

Proof: Iterating (9), we obtain

$$\hat{V}(abcd) = \hat{V}(a) + \overline{\psi}(a)\hat{V}(b) + \overline{\psi}(a)\overline{\psi}(b)\hat{V}(c) + \overline{\psi}(a)\overline{\psi}(b)\overline{\psi}(c)\hat{V}(d). \tag{11}$$

Hence, using (8) and (11), we obtain

$$\hat{V}(abcd) - \hat{V}(acbd) = \overline{\psi}(a)\{\hat{V}(c)[\overline{\psi}(b) - 1] - \hat{V}(b)[\overline{\psi}(c) - 1]\}, \tag{12}$$

and thus

$$\hat{V}(abcd) \leq \hat{V}(acbd) \quad \Leftrightarrow \quad \frac{\hat{V}(c)}{\overline{\psi}(c)-1} \leq \frac{\hat{V}(b)}{\overline{\psi}(b)-1} \quad \Leftrightarrow \quad \frac{\overline{\psi}(b)-1}{\hat{V}(b)} \leq \frac{\overline{\psi}(c)-1}{\hat{V}(c)},$$

from which the result follows with $b = i$ and $c = j$. ■

Remark 2 That an interchange argument leads to optimal "policies" in the present situation is clearly true, since no information gathering is done while a schedule is being implemented, and thus the optimal policy is an open-loop schedule [5]. This observation leads to the following scheduling rule.

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*

$$\frac{\operatorname{sgn}(\gamma)(\overline{\psi}(t_i)-1)}{p_{t_i}\psi(t_i)} \leq \frac{\operatorname{sgn}(\gamma)(\overline{\psi}(t_j)-1)}{p_{t_j}\psi(t_j)}, \tag{13}$$

whenever $i < j$.

Proof: For the case $\gamma > 0$, the result follows immediately by applying an interchange argument, based on Proposition 1. The analogous result for the case $\gamma < 0$ is derived similarly. Note that if $\gamma < 0$, then $\psi(i) < 1$ and $\overline{\psi}(i) < 1$. ■

Thus, from Theorem 1, we conclude that the components should be tested in increasing order of the ratios

$$\frac{\operatorname{sgn}(\gamma)(E[e^{\gamma D_i}]-1)}{p_i E[e^{\gamma C_i}]}.$$

The following result shows that the classical ratio test for the risk-neutral problem is recovered when the absolute value of the risk-sensitivity parameter is small.

Corollary 1 *Suppose that the ratios $\{E[D_i]/p_i\}_{i=1}^n$ are distinct. Then there exists $\epsilon > 0$ such that if $|\gamma| < \epsilon$ then there is a unique optimal complete schedule $t = (t_1, \ldots, t_n)$ for the MEF problem, and it is specified by the condition that whenever $i < j$, then*

$$\frac{E[D_{t_i}]}{p_{t_i}} < \frac{E[D_{t_j}]}{p_{t_j}}. \tag{14}$$

Proof: For small $|\gamma|$, $\psi(i) \approx 1 + \gamma E[C_i]$ and $\overline{\psi}(i) \approx 1 + \gamma E[D_i]$. Using these approximations, the assertion follows from Theorem 1. ■

Remark 3 From Corollary 1 we see that if the risk-sensitivity parameter is sufficiently small, the costs $\{C_i\}$ associated with *positive* tests are irrelevant to the determination of the optimal testing order. In particular this is the case in the risk-neutral problem–i.e., in the limit as $\gamma \to 0$. However, from Theorem 1, it follows that for the general risk-sensitive problem the costs associated with positive tests as well as costs associated with negative tests are relevant to the optimal testing order, a rather important distinction with respect to the risk-neutral case, i.e., the C/P algorithm.

■

Next, we consider the structure of the optimal schedule for $|\gamma| >> 0$. For this discussion, we make the simplifying assumption that for each i, the random variables C_i and D_i are constants with values c_i and d_i, respectively.

Corollary 2 *Suppose that $C_i = c_i$ and $D_i = d_i$, $\forall i$. Suppose that the complete schedule $t = (t_1, \ldots, t_n)$ is such that if $i < j$, then one of the following conditions holds:*

(a) $d_{t_i} - c_{t_i} < d_{t_j} - c_{t_j}$.

(b) $d_{t_i} - c_{t_i} = d_{t_j} - c_{t_j}$ *and* $p_{t_j} < p_{t_i}$.

(c) $d_{t_i} - c_{t_i} = d_{t_j} - c_{t_j}$ *and* $p_{t_j} = p_{t_i}$ *and* $c_{t_i} < c_{t_j}$.

Then for all sufficiently large $\gamma > 0$, t is the unique optimal complete schedule for the MEF problem.

Proof: From Proposition 1 it follows that $\hat{V}(aijb) < \hat{V}(ajib)$ if and only if

$$p_j e^{\gamma c_j}(e^{\gamma d_i} - 1) < p_i e^{\gamma c_i}(e^{\gamma d_j} - 1). \tag{15}$$

By considering (15) for $\gamma >> 0$ and using an interchange argument, the result follows. ■

In particular, this result shows that if the differences $\{d_i - c_i\}$ are distinct, then a very risk-averse decision maker should order the tests so that the differences between the cost of a negative test and positive test are increasing. On the other hand, if there are two tests for which these differences are equal, the test with higher probability of success (i.e., testing the component with higher failure rate) should be scheduled first. If both the differences and probabilities are equal, the test with the smaller cost should be scheduled first. In the special case in which the cost of negative tests is the same as that of positive tests, and in which the success probabilities are all distinct, the tests should be scheduled in order of decreasing success probability. The behavior exhibited by a very risk-averse decision maker is therefore "myopic."

Corollary 3 *Suppose that $C_i = c_i$ and $D_i = d_i$, $\forall i$. Suppose that the complete schedule $t = (t_1, \ldots, t_n)$ is such that if $i < j$, then one of the following conditions holds:*

(a) $c_{t_i} < c_{t_j}$.

(b) $c_{t_i} = c_{t_j}$ *and* $p_{t_j} < p_{t_i}$.

(c) $c_{t_i} = c_{t_j}$ *and* $p_{t_j} = p_{t_i}$ *and* $d_{t_i} < d_{t_j}$.

Then for all $\gamma < 0$ with $|\gamma|$ sufficiently large, t is the unique optimal complete schedule for the MEF problem.

Proof: Follows easily from Proposition 1 and Theorem 1, e.g., reversing the inequality in (15). ■

Example 1 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 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 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 $\gamma \leq 0$.

Since $d_3 - c_3 < d_1 - c_1 = d_2 - c_2$ and $p_1 < p_2$, it follows from Corollary 2 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 $[0.2645, \infty)$.

■

References

[1] D. Adolphson and T. C. Hu. Optimal linear ordering. *SIAM J. Appl. Math.*, 25:403–423, 1973.

[2] 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.

[3] 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.

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

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

[6] W. Black. Discrete sequential search. *Information and Control*, 8:156–162, 1965.

[7] 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.

[8] B. V. Dean. Stochastic networks in research planning. In Yovits et al., editor, *Research Program Effectiveness*, chapter 12. Gordon and Breach, New York, 1966.

[9] D. C. Denby. Minimum downtime as a function of reliability and priority in a component repair. *Journal of Industrial Engineering*, 18:436–439, 1967.

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

[11] 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.

[12] H. Greenberg. Optimum test procedure under stress. *Operations Research*, 12:689–692, 1964.

[13] G. J. Hall. Sequential search with random overlook probabilities. *Annals of Statistics*, 4:807–816, 1976.

[14] W. A. Horn. Single-machine job sequencing with treelike precedence ordering and linear delay penalties. *SIAM J. Appl. Math.*, 23:189–202, 1972.

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

[16] J-F. Huard. Probabilistic reasoning for fault management on xunet. Technical report, Summer Report, AT&T Bell Labs, Murray Hill, NJ, 1994.

[17] 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.

[18] F.V. Jensen. *An Introduction to Bayesian Networks*. Springer, New York, 1996.

[19] W. B. Joyce. Organization of unsuccessful r&d programs. *IEEE Transactions on Engineering Management*, 18:57–65, 1971.

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

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

[22] 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.

[23] S. I. Marcus, E. Fernández-Gaucherand, D. Hernández-Hernández, S. Coraluppi, and P. Fard. Risk sensitive markov decision processes. In C. I. Byrnes, B. N. Datta, D. S. Gilliam, and C. F. Martin, editors, *Systems & Control in the Twenty-First Century*, pages 263–279. Birkhauser, 1996.

[24] D. Matula. A periodic optimal search. *Amer. Math. Monthly*, 71:15–21, 1964.

[25] L. G. Mitten. An analytic solution to the least cost testing sequence problem. *Journal of Industrial Engineering*, 11:17, 1960.

[26] 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.

[27] M. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Prentice Hall, Englewood Cliffs, NJ, 1995.

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

[29] M. A. Shayman and E. Fernández-Gaucherand. Risk-sensitive decision-theoretic troubleshooting. Technical Report, July 8, 1999.

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

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

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

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

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

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

[36] P. Whittle. *Risk-sensitive Optimal Control*. Wiley, New York, 1990.