Optimal Design and Use of Retry in Fault-Tolerant Computer Systems

YANN-HANG LEE AND KANG G. SHIN

University of Michigan, Ann Arbor, Michigan

Abstract. In this paper, a new method is presented for (i) determining an *optimal retry policy* and (ii) using retry for *fault characterization*, which is defined as classification of the fault type and determination of fault durations. First, an optimal retry policy is derived for a given fault characteristic, which determines the maximum allowable retry durations so as to minimize the total task completion time. Then, the *combined* fault characterization and retry decision, in which the characteristic of a fault is estimated simultaneously with the determination of the optimal retry policy, are carried out. Two solution approaches are developed: one is based on point estimation and the other on Bayes sequential decision analysis.

Numerical examples are presented in which all the durations associated with faults (i.e., active, benign, and interfailure durations) have monotone hazard rate functions (e.g., exponential Weibull and gamma distributions). These are standard distributions commonly used for modeling and analyses of faults.

Categories and Subject Descriptors: B.2.3 [**Arithmetic and Logic Structures**]: Reliability, Testing, and Fault-Tolerance; G.3 [**Mathematics of Computing**]: Probability and Statistics—*statistical computing*

General Terms: Algorithms, Design, Performance, Reliability, Verification

Additional Key Words and Phrases: Bayes decision problem, estimation, fault characteristic, hypothesis testing, optimal retry

1. *Introduction*

Faults in computer systems are usually classified into three types: *transient*, *intermittent*, and *permanent* [23]. Transient faults die within a certain time of their generation, intermittent faults cycle between being active and inactive, and permanent faults are (as the term indicates) permanent. When an error induced by an existing fault is detected,[1] the system retries to recover from the fault. The executing

[1] Normally, errors are detected but faults are not. However, if an error is defined as incorrectness in the user's program, then the manifestation of a fault captured by built-in detection mechanisms is not the detection of an error, since the fault did not yet induce an error. In order to avoid an endless pedantry, we use the term "failure detection" for detection of an error or manifestation of a fault, and sometimes the term "fault detection" for detection of a fault manifestation.

This work was supported in part by NASA under grant NAG 1-296. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of NASA.

Authors' present addresses: Y.-H. Lee, IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598; K. G. Shin, Division of Computer Science and Engineering, Department of Electrical Engineering and Computer Science, the University of Michigan, Ann Arbor, MI 48109. All correspondence should be addressed to Prof. Kang G. Shin at the above address.

tasks can be continued if retry is successful, or other recovery methods are applied if retry is not successful.

As the term implies, retry consists of restoring the affected process to some fault-free initial state and then rerunning it on the same processor. Retry will become successful after the existing fault disappears. Clearly, retry is only applicable when the manifestation of a fault is confined and the process can be restored to integrity. Although restart and rollback recovery can be viewed as a sort of retry, the retry used in most computer systems means the repetition of microinstruction(s) or an instruction. The latter type of retry is typically hardware controlled [4, 5] and thus has advantages of requiring small time overhead. The time-consuming diagnosis and reconfiguration of the system can be avoided in the case of a successful retry. Moreover, such retries are reported to be highly successful owing to the fact that only a small percentage of faults are permanent [1, 26]. Consequently, we focus in this paper on hardware-controlled retry in which a failure is detected immediately upon its occurrence by certain signal-level detection mechanisms [21]. We assume there are some scratchpad memories used in restoring the process to integrity. Results obtained by Carter et al. [5] indicate that self-checking and retry mechanisms can be incorporated into processors inexpensively and without substantially degrading performance.

Currently, several commerical machines incorporate retry. In the Honeywell 6000 [17], instruction retry is reported to approach an effectiveness rate of 100 percent. Retry in the IBM 360 and 370 series machines is widely used in the peripheral areas (I/O and storage), as well as in the central processor [12]. The UNIVAC 1100/60 uses a hardware timer that goes off after an interval judged to be long enough to allow transient faults to die out, upon which retry can be effected [3]. However, we are provided with no discussion or justification regarding the detailed design of retry, for example, number of retry attempts. This can be seen more clearly when we consider the statement in [4]: "If successful, computer operation proceeds; if unsuccessful the above process [i.e., reexecution of the previous instruction] is usually repeated N times before diagnosis begins."

Clearly, the usefulness of retry mechanisms arises, as we have noted, from (i) the smallness of the proportion of permanent faults in any computer system, and (ii) the fast recovery from nonpermanent faults and thus the small task completion time and recovery overhead. In the case of a permanent fault, to retry a process on the affected processor is worse than useless; it is a waste of time. Thus, the number of retry attempts or retry duration within which retry is applied should be controlled to maximize the difference between the expected gain in performance that results from using retry when the fault is transient or intermittent, and the expected loss that results from using it when the fault is permanent. In this paper, we focus on the determination of the maximum allowable retry duration r^* for the purpose of reducing expected task completion time. If the retry succeeds within this duration, the execution continues. If not, other methods for failure recovery, for example, rollback or restart following the system reconfiguration, must be used. (See Figure 1 for a standard procedure for task execution under the occurrence of fault.)

In addition to the performance gain in case of a successful retry, the characteristic of a fault can be monitored through retries. That is, a retry that succeeds within the retry duration r^* implies that the active duration of the fault following its detection is also less than or equal to r^*. Even when the retry fails, it indicates that the active duration of the fault following its detection is greater than r^*. On the other hand, the detection of a failure gives information regarding the duration between fault occurrences and the benign duration of an intermittent fault. Thus,

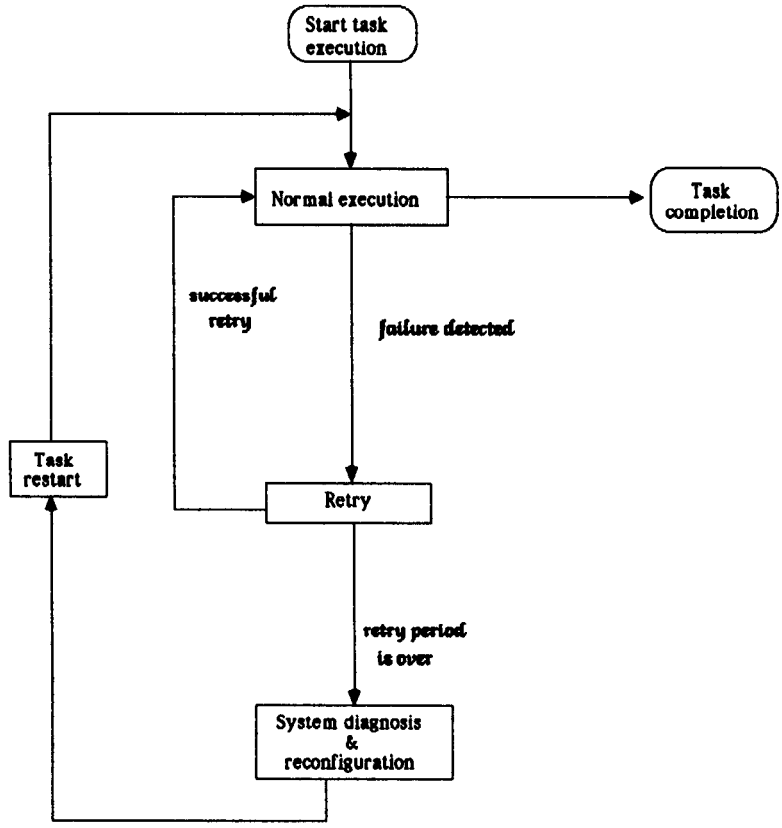

FIG. 1. Standard procedure for handling failures during task execution.

it becomes possible to observe the nature of a fault through both retry and detection mechanisms. Note that due to nonzero fault latency [19–21], the observed nature of a fault may not be the same as its true nature. Although fault latency can be measured [22], we ignore this latency in the rest of the paper, since it has no effect on retry. In the discussion that follows, we treat a fault on the basis of its observed behavior and its effects on the system. For example, we use the term *active duration of a fault* to mean the active duration of a fault following its detection.

Section 2 presents a brief description of the fault model along with necessary assumptions and an informal statement of the problem. It should be obvious that r^* depends on fault behavior, and in Section 3, we show how to derive it, given the fault characteristics. When quantitative descriptions of fault behavior are hard to come by in the real environments, the combination of retry and detection enables us to observe the fault characteristics, while determining the optimal retry policy. We counter this in Section 4 by showing how to use statistical estimation theory to create a system that learns the fault characteristics as it goes, via retry, and therefore becomes increasingly more "optimal" in the sense of minimizing the expected task completion time. In Section 5 we apply Bayes sequential decision analysis to fault characterization and retry decision. The backward induction for testing hypotheses is also presented as an example solution to the sequential decision problem. The paper concludes with Section 6.

In what follows, we use continuous retry instead of the number of retry attempts. A continuous retry can be understood easily when one considers the following two cases: detection mechanisms can monitor the presence of a fault continuously or

retry can be performed instantaneously. Conversion between a continuous retry and its corresponding number of retry attempts is not difficult and is discussed in Section 6.

2. *Fault Model, Assumptions, and Objectives*

Consider the behavior of faults in a computer system. Assume that arrival of faults is a time-invariant Poisson process with rate λ. When a fault occurs, it is assumed to be transient, intermittent, or permanent with probability p_t, p_i, or p_p, respectively. If a permanent fault occurs, it will remain constantly in the system until the component containing the fault is removed. Once a transient fault occurs in the system, it will disappear after an active duration, T_t^a. On the other hand, in case of an intermittent fault, it may become benign after an active duration, T_i^a, and then reappear after a benign duration, T_i^b. That is, an intermittent fault cycles between active and benign states. For simplicity, we assume that T_t^a, T_i^a, and T_i^b are mutually independent random varibles with distribution function F_t^a, F_i^a, and F_i^b and density functions f_t^a, f_i^a, and f_i^b, respectively. Thus, the characteristic of a fault can be represented by a 7-tuple

$$C_f \in \{(p_t, p_i, p_p, \lambda, F_t^a, F_i^a, F_i^b);\ p_t + p_i + p_p = 1\}.$$

Since the interarrival time of faults is usually much larger than any other durations, it is reasonably accurate to assume that there is at most a single fault in the system at any moment. Thus, the above behavior of faults enables us to model the system with a stochastic process shown in Figure 2. Denote the three possible states, namely, nonfaulty, fault-active, and fault-benign by **NF**, **F**, and **FB**. When a fault occurs, the system state changes from **NF** to **F**. The system moves back to **NF** if the fault is transient and disappears. It remains at **F** if the fault is permanent. If the fault is intermittent and becomes benign following an active duration, the system state changes from **F** to **FB**. The system may move back to **F** when this intermittent fault recurs—this is referred to as the *reappearance* of the intermittent fault. Models similar to this have been widely used in reliability analyses and the modeling of faults [14, 21, 25].

When a failure is detected (i.e., the system is in fault-active state), retry is usually applied as a first-step recovery means. Retry will be successful if the fault disappears during the retry period, that is, if the system changes to either nonfaulty or fault-benign state during retry. Otherwise, the system is reconfigured and the executing task is migrated to a nonfaulty component and then recovered via the other means such as rollback or restart. The advantages of a successful retry are twofold. One is the avoidance of complicated, time-consuming recovery actions, such as fault-isolation, system reconfiguration, and task recovery. The other attractive gain from retry is to rescue an executing task. Consider a practical case in which a system (i) becomes faulty once and gets back to normal during execution of a task, and (ii) never becomes faulty again before the task is completed. In such a case, it is possible to avoid the overhead of migrating and restarting the task by means of a successful retry, leading to a faster completion of the task.

In what follows, we derive an optimal retry policy that minimizes the expected task completion time when failures are detected during the task execution. We assume that initially no task is started on any faulty or potentially faulty module (having a benign intermittent fault) and that the system has enough redundancy so it can be reconfigured and made operational again when retry recovery fails. In such a case, the associated task is restarted following system reconfiguration.

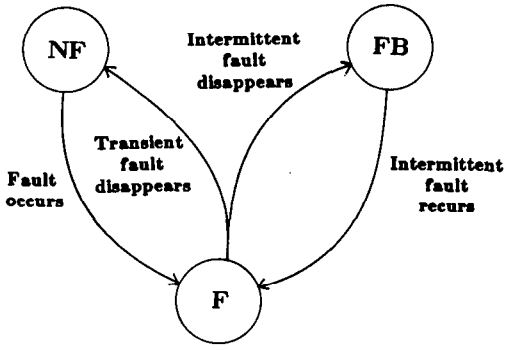

FIG. 2. The fault model. NF = nonfaulty, F = faulty, FB = fault-benign.

When the system enters fault-active state, there is no way to determine whether the fault is transient, intermittent, or permanent. If retry becomes successful, the fault is either transient or intermittent. Under the assumption that the benign duration of an intermittent fault is much smaller than the fault interarrival time, the fault is declared to be intermittent if the system fails again within a short period after the disappearance of the previous fault. Thus, a retry policy should specify two maximum allowable retry durations: one for a new fault and the other for an old recurring intermittent fault.

Our problem is to derive an optimal retry policy that minimizes task completion time under the occurrence of faults. When the characteristics of faults (e.g., the durations associated with the above fault model) are not known a priori, the problem calls for an adaptive optimization in which the system decides a retry policy to minimize the task completion time while learning about the fault characteristics via retry. We solve this problem beginning with a simple case in Section 3 where the fault characteristics are known, and then the general case in Sections 4 and 5 where the fault characteristics are unknown.

3. *Optimal Retry Policy for Given C_f*

3.1 EXPECTED TASK COMPLETION TIME. Let x_0 denote the computation time initially needed to complete the task under a fault-free condition. A failure may be detected when the amount of computation remaining to complete the task, that is, *residual computation*, is reduced to x, where $0 < x < x_0$. A retry policy is defined as $R = \{(r_1(x, C_f), r_2(x, C_f)); 0 < x < x_0\}$, where the maximum retry durations are $r_1(x, C_f)$ and $r_2(x, C_f)$, respectively, for the detection of a *new* fault and an *old* intermittent fault when the residual computation is x and the fault characteristic is C_f. For notational simplicity, we shall use r_i, whenever convenient, in the sequel, to represent $r_i(x, C_f)$, $i = 1, 2$.

Let the expected times needed to complete the residual computation x be denoted by $V_1(x, C_f, R)$, $V_2(x, C_f, R)$, $V_3(x, C_f, R)$, and $V_4(x, C_f, R)$ when the retry policy R is adopted and the system is in the following situations: execution starts/resumes on a nonfaulty system, a new fault is detected, an old intermittent fault is detected again, and execution continues following a successful retry for an intermittent fault, respectively. Based on transitions among these situations, one can derive the following recursive equations for $V_1(x, C_f, R)$ through $V_4(x, C_f, R)$. For example, when the residual computation is x, the task completion time needed following the task resumption would be x if no new fault occurs within the duration x or would

be the sum of t and the time needed for completing the computation if a new fault occurs at the residual time $x - t$, that is, $V_2(x - t, C_f, R)$.

$$V_1(x, C_f, R) = e^{-\lambda x}x + \int_0^x \{t + V_2(x - t, C_f, R)\}\lambda e^{-\lambda t}\, dt, \tag{1}$$

$$\begin{aligned} V_2(x, C_f, R) = {} & p_t \int_0^{r_1} \{t + V_1(x, C_f, R)\}\, dF_t^a(t) \\ & + p_i \int_0^{r_1} \{t + V_4(x, C_f, R)\}\, dF_i^a(t) \\ & + \{1 - p_t F_t^a(r_1) - p_i F_i^a(r_1)\}\{V_1(x_0, C_f, R) + r_1 + t_s\}, \end{aligned} \tag{2}$$

$$\begin{aligned} V_3(x, C_f, R) = {} & \{1 - F_i^a(r_2)\}\{V_1(x_0, C_f, R) + r_2 + t_s\} \\ & + \int_0^{r_2} \{t + V_4(x, C_f, R)\}\, dF_i^a(t), \end{aligned} \tag{3}$$

$$V_4(x, C_f, R) = \{1 - F_i^b(x)\}x + \int_0^x \{t + V_3(x - t, C_f, R)\}\, dF_i^b(t), \tag{4}$$

where t_s is the set-up time necessary for system reconfiguration and task reinitialization. The optimization problem can be defined as follows: Find an optimal retry policy,

$$R^* = \{(r_1^*(x, C_f), r_2^*(x, C_f)); 0 < x < x_0\},$$

such that $\forall x\ V_2(x, C_f, R^*) = \min_R V_2(x, C_f, R)$ and $V_3(x, C_f, R^*) = \min_R V_3(x, C_f, R)$. Obviously, this optimal policy also minimizes $V_1(x, C_f, R)$ and $V_4(x, C_f, R)$.

Since the mean time between failures is usually much longer than the other durations, $V_1(x, C_f, R)$ can be accurately approximated by x. The bounds of $V_1(x, C_f, R)$ under the optimal retry policy are derived in the Appendix A. Note that the difference between the upper and lower bounds of $V_1(x, C_f, R^*)$ is negligible. Thus, the approximation is used throughout the rest of the paper.

In general, there are no closed-form solutions for $r_1^*(x, C_f)$ and $r_2^*(x, C_f)$. However, these optimal retry durations can be calculated numerically as explained below. Let Δx be an arbitrarily small positive value. With the initial condition $V_4(0, C_f, R) = 0$, $V_3(k\Delta x, C_f, R)$ with $k = 0$ and any r_2 can be computed. Then, r_2^* is chosen to minimize $V_3(k\Delta x, C_f, R)$ for the residual computation $k\Delta x$, $k = 1$. By incrementing k we can recursively compute $V_4(k\Delta x, C_f, R)$, $V_3(k\Delta x, C_f, R)$, and r_2^* for the residual computation $k\Delta x$. Once $V_4(x, C_f, R)$ is known, one can compute $V_2(x, C_f, R)$ for any r_1 and can therefore determine $r_1^*(x, C_f)$.

Define the *recovery overhead* as the total time required to resume normal operation following the detection of a failure. When the recovery overhead in place of the task completion time is to be minimized, $r_2^*(x, C_f) = 0$, $\forall x \in (0, x_0)$, because the recovery overhead will accrue during reappearances of an intermittent fault. In this case, the recovery overhead can be expressed as $V_2(x, C_f, R) - V_1(x, C_f, R)$, which is the time spent to restore the system to its state immediately before the failure is detected. The optimal retry duration $r_1^*(x, C_f)$ can be determined from eq. (2) just as we can compute that for minimizing the task completion time.

3.2 Fault Active Durations with Monotone Hazard Rate Functions. Since T_i^a is a continuous random variable, one can assume that $f_i^a(t)$ is continuous in $[0, \infty)$. The *hazard rate function* of the active duration of an intermittent fault

is defined by $u_i^a(t) \equiv f_i^a(t)/(1 - F_i^a(t))$. When the hazard rate function of the active duration of an intermittent fault is monotonically increasing, constant, or monotonically decreasing, the optimal retry duration r_2^* exhibits interesting properties. These properties play a significant role in determining the optimal retry policy, since the time durations associated with faults are usually modeled to have monotone hazard rate functions. Typical distributions with monotonically increasing hazard rate functions include the gamma and the Weibull distributions with the shape parameters greater than 1. When their shape parameters are less than 1, they have monotonically decreasing hazard rate functions. The exponential distribution has a constant hazard rate. Consider first the nondecreasing hazard rate function that leads to the following theorem.

THEOREM 1. *When $u_i^a(t)$ is monotonically nondecreasing in t and $V_1(x, C_f, R) \doteq x$, $r_2^* = 0$ or $r_2^* = \infty$.*

PROOF. Using the approximation $V_1(x, C_f, R) \doteq x$ and differentiating eq. (3) with respect to r_2, we obtain

$$\frac{\partial V_3(x, C_f, R)}{\partial r_2} = f_i^a(r_2)\left[V_4(x, C_f, R) + \frac{1}{u_i^a(r_2)} - \{V_1(x_0, C_f, R) + t_s\}\right]$$

$$\doteq f_i^a(r_2)\left\{V_4(x, C_f, R) + \frac{1}{u_i^a(r_2)} - (x_0 + t_s)\right\}. \tag{5}$$

Since $V_4(x, C_f, R)$ is independent of the past and current retry durations[2] $r_2(y, C_f)$, where $y \geq x$, $V_4(x, C_f, R) + (1/u_i^a(r_2)) - (x_0 + t_s)$ is nonincreasing in $r_2(x, C_f)$. Thus, $V_3(x, C_f, R)$ is a concave function of r_2. The optimal retry duration r_2^* is then equal to 0 or ∞. □

Following the definition of $r_2^*(x, C_f)$, $r_2^*(x, C_f) = 0$ implies that no retry be attempted for reappearing intermittent faults, whereas $r_2^*(x, C_f) = \infty$ means that retry should be applied until the intermittent fault becomes benign.

COROLLARY 1. *When $u_i^a(t)$ is monotonically nondecreasing in t, and if there exists an x_2^* such that $x_0 + t_s - x_2^* - R_i^b(x_2^*)E[T_i^a] = 0$, where $R_i^b(x)$ is the renewal function* [7] *corresponding to the distribution $F_i^b(t)$, then $r_2^*(x, C_f) = \infty$ if $x \leq x_2^*$ and $r_2^*(x, C_f) = 0$, otherwise.*

PROOF. From Theorem 1, $r_2^*(x, C_f)$ is either ∞ or 0. When $r_2^*(x, C_f) = \infty$, there exists an r such that the integral $\int_0^r (\partial V_3(x, C_f, R)/\partial r_2)\, dr_2$ becomes negative. Since $V_4(x, C_f, R^*)$ is a monotonically nondecreasing function of x, there also exists an r such that the integral $\int_0^r (\partial V_3(y, C_f, R)/\partial r_2)\, dr_2$ becomes negative when $y \leq x$. Thus, $r_2^*(y, C_f) = \infty$ $\forall y \leq x$. Using the assumption that the active and benign durations are mutually independent, we get $V_4(x, C_f, R^*) = x + \{E[N(x)] - 1\}E[T_i^a]$, where $N(x)$ is the number of reappearances of the intermittent fault during the residual computation x, namely, $N(x) = \inf\{n; \sum_{k=1}^n T_{i,k}^b \geq x\}$, where $T_{i,k}^b$ is the benign duration following the kth occurrence of the intermittent fault. The expected value of $N(x)$, $E[N(x)]$, is equivalent to the renewal function

[2] Note that the probability of having a zero benign duration of an intermittent fault should be zero, that is, $\Pr(T_i^b = 0) = 0$. Otherwise, no useful computation can be done.

$R_i^b(x)$ corresponding to the distribution $F_i^b(t)$. Also, $V_3(x, C_f, R)|_{r_2=\infty} \leq V_3(x, C_f, R)|_{r_2=0}$ if $r_2^*(x, C_f) = \infty$, that is,

$$\int_0^\infty (t + V_4(x, C_f, R^*))\, dF_i^a(t) = \mathrm{E}[T_i^a] + V_4(x, C_f, R^*) \leq a(x_0) + t_s.$$

From the equality in the right-hand side of the above equation, we obtain x_2^* and thus the corollary is proved. □

Theorem 1 can also be viewed as below using the concept of *stochastic ordering* between two random variables. A random variable X is said to be *stochastically larger than* the other random variable Y if $\Pr(X > t) \geq \Pr(Y > t)\ \forall t$ [18]. Let $T_i^a(|\, r)$ be the remaining life of the intermittent fault after retry has been applied for the duration r. When the hazard rate function is nondecreasing, $T_i^a(|\, r)$ is stochastically larger than $T_i^a(|\, s)$, provided $r \leq s$. Thus, $\forall s \geq r$; if it is worth continuing retry beyond the retry duration r (in the sense of minimizing the task completion time), then we should continue the retry even after the retry duration s. Consequently, the retry continues until the intermittent fault disappears.

Note that when the hazard rate function is nondecreasing, x_2^* is determined by the mean active duration and is independent of the shape of the distribution. x_2^* could become negative when $\mathrm{E}[T_i^a]$ is large, that is, intermittent faults have a long active duration. In such a case, Corollary 1 implies that no retry be applied for intermittent faults. On the other hand, if the set-up overhead t_s is large, x_2^* could be even larger than x_0, implying that retry be used as a sole means of recovering from an intermittent fault.

When the hazard rate $u_i^a(t)$ is decreasing, the nice properties stated in both Theorem 1 and Corollary 1 do not exist. However, there exists at most one root of eq. (5) that minimizes V_3. In such a case, since there is no closed-form expression of $V_4(x, C_f, R^*)$, we have to resort to numerical techniques for determining both $r_2^*(x, C_f)$ and $r_1^*(x, C_f)$, as was previously mentioned.

Several numerical examples are shown in Figures 3–5, where the durations are normalized with respect to x_0, and the active duration of the intermittent fault is assumed to have the gamma or Weibull distribution. Figure 3 presents x_2^*'s when the shape parameter α's of the gamma and Weibull distributions, respectively, are greater than or equal to 1. Figures 4 and 5 show the optimal retry duration $r_2^*(x, C_f)$; the solid lines for $\alpha < 1$ and the dashed lines for $\alpha = 1$. Note that for the gamma distribution $u_i^a(t)$ approaches $1/\beta$ as $t \to \infty$, where β is the scale parameter. Thus, it is possible for the derivative of V_3 to be negative, (i.e., eq. (5) becomes negative), implying $r_2^*(x, C_f) = \infty$. For the Weibull distribution with $\alpha < 1$, r_2^* never becomes ∞ since $u_i^a(\infty) = 0$.

Consider the case where T_t, T_i^a, and T_i^b are all exponentially distributed with the parameters τ, μ, ν for the transient fault disappearance rate, the intermittent fault disappearance and reappearance rates, respectively. Since $f_i^b(t) = \nu e^{-\nu t}$, the renewal function $R_i^b(x)$ becomes $1 + \nu x$. From Corollary 1, we have $r_2^*(x, C_f) = \infty$ if $x \leq x_2^*$ and

$$r_2^*(x, C_f) = 0 \qquad \text{if} \quad x > x_2^*,$$

where

$$x_2^* = \frac{\mu}{\mu + \nu}\left(x_0 + t_s - \frac{1}{\mu}\right).$$

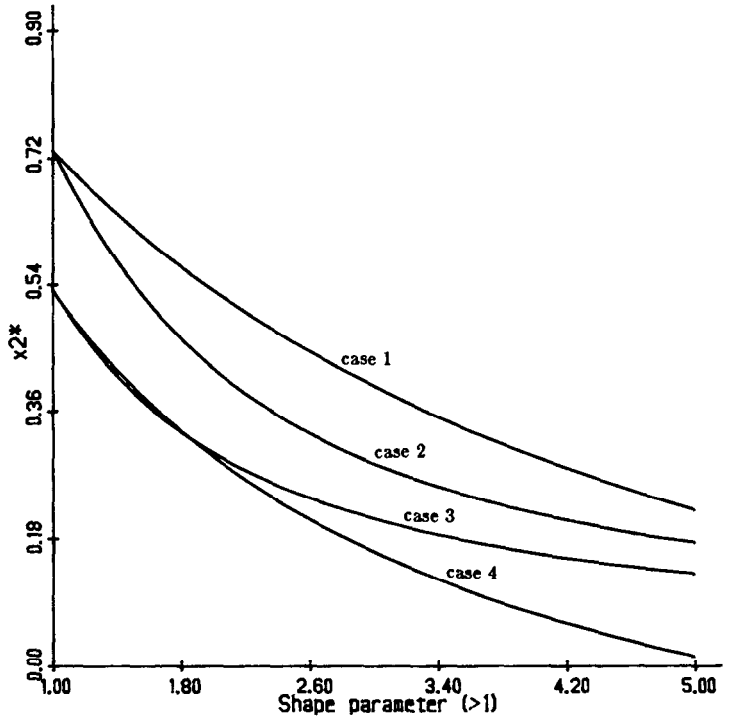

FIG. 3. $x_2^*(x, C_f)$ versus the shape parameter α when the hazard rate is increasing and $f_t^b(t) = 3e^{-3t}$. Case 1: Gamma distribution, $\beta = 0.1$. Case 2: Weibull distribution, $\beta = 0.1$. Case 3: Weibull distribution, $\beta = 0.2$. Case 4: Gamma distribution, $\beta = 0.2$.

$V_4(x, C_f, R^*)$ then becomes

$$V_4(x, C_f, R^*) = \begin{cases} \left(1 + \left(\frac{\nu}{\mu}\right)\right)x & \text{if } x \le x_2^*, \\ x_0 + t_s + \frac{1}{\nu} - \left(\frac{1}{\mu} + \frac{1}{\nu}\right)\exp(-\nu(x - x_2^*)) & \text{if } x > x_2^*. \end{cases} \tag{6}$$

The derivative of $V_2(x, C_f, R)$ with respect to r_1 becomes

$$\frac{\partial V_2(x, C_f, R)}{\partial r_1} = p_p + p_t\exp(-\tau r_1)\{1 - (x_0 + t_s - x)\tau\} + p_i\exp(-\mu r_1)[1 - \{x_0 + t_s - V_4(x, C_f, R)\}\mu]. \tag{7}$$

With $r_2^*(x, C_f)$ determined as in Corollary 1 and $V_4(x, C_f, R^*)$ as in eq. (6), eq. (7) can have at most two roots. The optimal retry duration $r_1^*(x, C_f)$ can be obtained by examining $V_2(x, C_f, R)$ at the boundaries, $r_1 = 0$ and $r_1 = \infty$, and the roots of eq. (7). Note that r_1^* cannot be infinite as long as $p_p < 0$. Unlike r_2^*, r_1^* does not have to be zero when $x > x_2^*$. Several cases of $V_2(x, C_f, R)$ as a function of r_1 are shown in Figure 6 where all parameters are normalized with respect to x_0. The case 2 in Figure 6 shows an example for which two positive roots of eq. (7) exist.

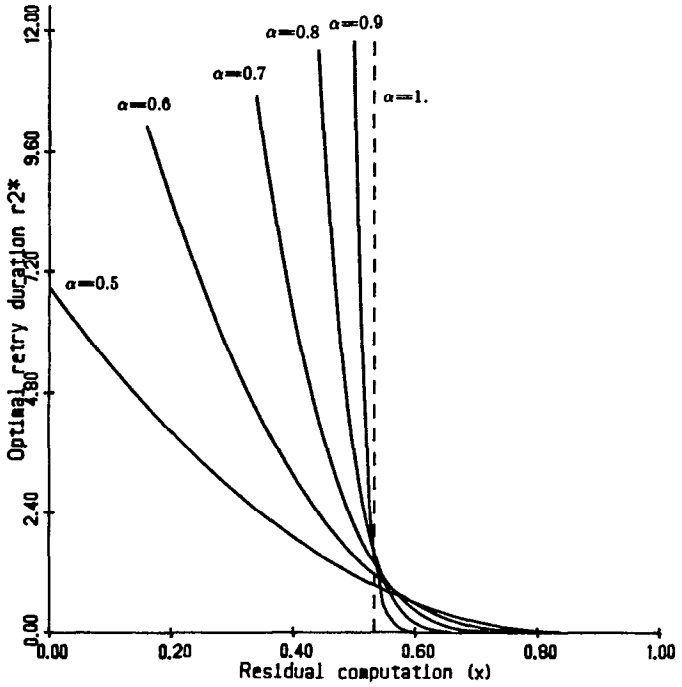

FIG. 4. The optimal retry duration $r_2^*(x, C_f)$ for Weibull distributions with increasing hazard rate. The density function of the active duration $f_i^a(t) = \alpha/\beta\, t^{\alpha-1}\exp(-t^\alpha/\beta)$. $\beta = 0.2; f_i^b(t) = 3e^{-3t}$.

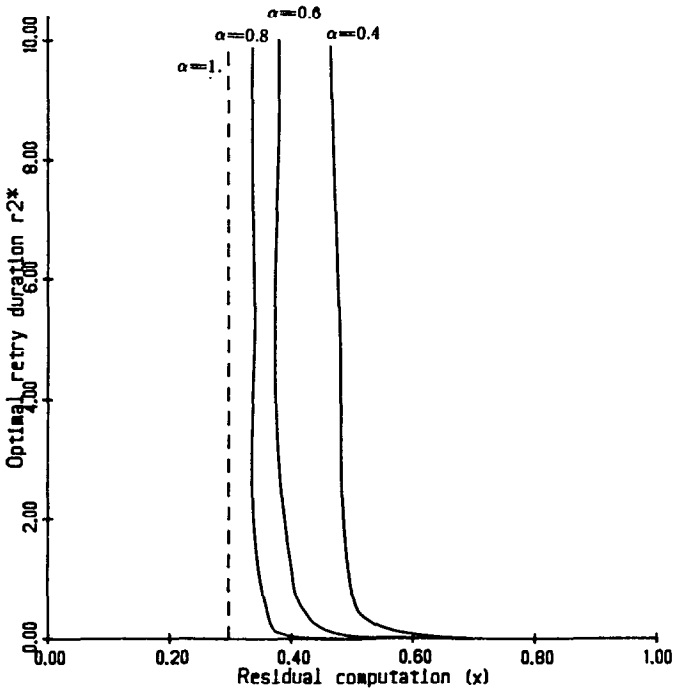

FIG. 5. The optimal retry duration $r_2^*(x, C_f)$ for Gamma distributions with decreasing hazard rate. The density function of active duration $f_i^a(t)\ 1/(\beta^\alpha\Gamma(\alpha))t^{\alpha-1}\exp(-t/\beta)$. $\beta = 0.4; f_i^b(t) = 3e^{-3t}$.

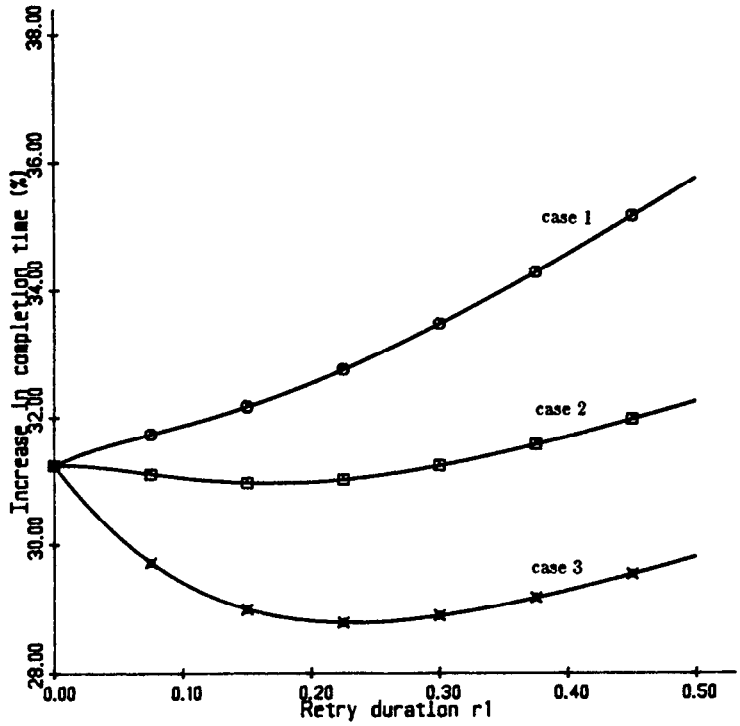

FIG. 6. $V_2(x, C_f, R) - x$ versus the retry during r_1 for exponentially distributed durations, $x = 0.8$, $\tau = 10$, $\mu = 12$. Case 1: $p_p = 0.1$, $p_t = 0.5$, $p_i = 0.4$, $\nu = 6$. *Case* 2: $p_p = 0.05$, $p_t = 0.55$, $p_i = 0.4$, $\nu = 12$. Case 3: $p_p = 0.05$, $p_t = 0.6$, $p_i = 0.35$, $\nu = 6$.

Figure 7 presents some numerical results on $r_1^*(x, C_f)$ as a function of x. Note that x_2^* depends upon the ratio of ν to μ, whereas r_1^* varies as p_t, p_i, and p_p change.

4. *Optimal Retry Policy and Parameter Estimation*

In Section 3, we have derived an optimal retry policy for a given fault characteristic C_f. It is, however, very difficult in practice to know a priori the fault characteristic. Even if the fault characteristic is measured during device manufacture, it may well vary as the execution environment and the executing tasks change. Another factor that makes the fault characteristics time variant is the aging of components, for example, the bathtub curve of the failure rate as a function of time [23]. Thus, it is important to determine an optimal retry policy for uncertain fault characteristics.

Detection mechanisms can be useful in collecting data of the duration between two successive fault occurrences or the benign duration of an intermittent fault for characterizing the behavior of fault occurrence and reappearance. Retry may lead to an indication of the active duration of a transient or intermittent fault, which is, on the other hand, affected by the retry policy applied. More specifically, when C_f is unknown, C_f has to be estimated first with the observation of system state transitions (as shown in the model of Figure 2) with retry and detection mechanisms. Then, the retry duration will be determined based on an estimated C_f. In such a case, the computer system has to adjust its retry policy using the information on the fault behavior collected during its past and current retries. See Figure 8 for a block diagram of such an adaptive optimization.

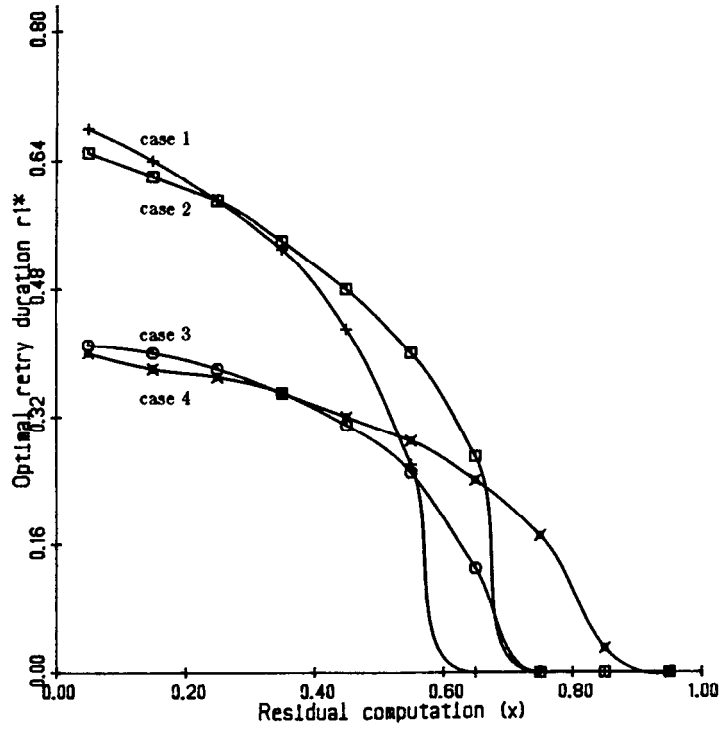

FIG. 7. The optimal retry duration $r_1^*(x, C_f)$ for exponentially distributed durations. Case 1: $p_p = 0.1$, $p_t = 0.3$, $p_i = 0.6$, $\tau = 6$, $\mu = 5$, $\nu = 3$. Case 2: $p_p = 0.1$, $p_t = 0.6$, $p_i = 0.3$, $\tau = 6$, $\mu = 5$, $\nu = 3$. Case 3: $p_p = 0.1$, $p_t = 0.3$, $p_i = 0.6$, $\tau = 12$, $\mu = 10$, $\nu = 6$. Case 4: $p_p = 0.1$, $p_t = 0.6$, $p_i = 0.3$, $\tau = 12$, $\mu = 10$, $\nu = 6$.

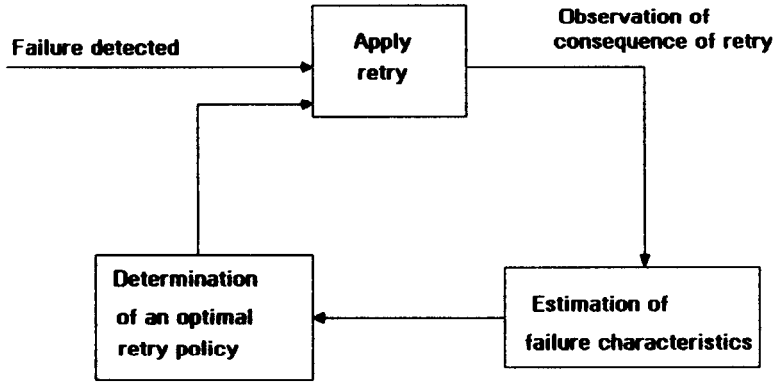

FIG. 8. Block diagram of the adaptive optimization.

In what follows, we shall limit the discussion only to the characterization of the active duration of intermittent faults and the simultaneous determination of an optimal retry policy that minimizes the task-completion time. The reasons for such a limit are

(1) When T_f, T_t^a, T_i^a, and T_i^b are mutually independent, the estimation of these durations can be treated separately and individually with the same approach.

(2) Unlike the probabilities of having different types of fault and the fault occurrence interval, the behavior of an intermittent fault is difficult to measure.
(3) The fault occurrence interval and the benign duration of an intermittent fault can be observed completely via detection mechanisms, whereas the information collected through retry may become *incomplete* when the retry stops unsuccessfully.
(4) When retry is applied, the current sample has effects on the current retry decision during its collection. Thus, simultaneous estimation and decision must be performed continuously during retry.

Let the active duration of an intermittent fault have the density function form $f_i^a(t \mid \theta)$ with the parameter θ unknown. (θ could be a vector if there are two or more parameters, for example, the shape parameter α and the scale parameter β for the Weibull and gamma distributions.) If the form of the density function is not known a priori, we can assume several possible forms and perform a test of fit, for example, a chi-square test of fit, or hypotheses testing, as is discussed in the next section.

The samples obtained for the active duration can be represented by a 2-tuple (I, t) where I is a single-bit flag and t indicates a duration. $I = 0$ represents a successful retry, and hence t indicates the active duration of the fault. On the other hand, when a retry fail, $I = 1$ and t is the retry duration. Let (I_i, t_i), $i = 1, 2, \ldots n$, denote the past samples related to the active duration of an intermittent fault. These resulting samples are *type I progressively censored*, following Cohen's definition in [8] with continuous censoring times. There are several different types of estimators conceivable for estimating the parameter θ on the basis of these progressively censored samples. For the Weibull and gamma distributions, the maximum likelihood estimators have been widely studied as in [8–10], [16], and [28] when the samples are progressively censored.

When the fault is still active even after the current retry duration r, we shall have collected an additional sample $(1, r)$ via the current retry. Let $\hat{\theta}(r)$ be the maximum likelihood estimator of θ which is based on the samples up to and including the current sample $(1, r)$. Note that the dependency of $\hat{\theta}$ on the current sample $(1, r)$ is expressed explicitly, since the current retry duration will depend on this estimated $\hat{\theta}$.

Let the optimal retry durations based on the estimated $\hat{\theta}(r)$ be denoted by $r_k^*(x, \hat{\theta}(r))$, $k = 1, 2$, for a newly detected fault and an old intermittent fault, respectively. Use the notation $C_f(\hat{\theta}(r))$ to indicate that the active durations of intermittent faults have the density function $f_i^a(t \mid \hat{\theta}(r))$, and let $R(r)$ denote the policy that the maximum allowable retry duration for the current retry is r. Then, the direct solution of the optimal retry duration is to find the minimum of $V_k(x, C_f(\hat{\theta}(r)), R(r))$, $k = 1, 2, 3, 4$. Notice that the retry duration r not only appears in the integral equations (2) and (3), but also affects the fault characteristic C_f.

Under certain conditions, it can be proved that $r_k^*(x, \hat{\theta}(r))$ is a nonincreasing function of r. We first derive the results under such conditions and then discuss its application later in this section.

THEOREM 2. *When (a) the active duration of an intermittent fault has the density function $f_i^a(t \mid \theta)$, and (b) for $t_j \geq t_k$ the ratio $(f_i^a(t \mid \hat{\theta}(t_j)))/(f_i^a(t \mid \hat{\theta}(t_k)))$—a likelihood ratio* [15]*—is nondecreasing in t, then the optimal retry duration is determined by*

$$r_k^* = \inf\{r;\ r_k^*(x, \hat{\theta}(r)) \leq r\}. \tag{8}$$

To prove this theorem, we need the following three lemmas.

LEMMA 1. *Under the same conditions as in Theorem* 2, *let* T_j *and* T_k *be random variables with the density functions* $f_i^a(t \mid \hat{\theta}(t_j))$ *and* $f_i^a(t \mid \hat{\theta}(t_k))$, *respectively, and let* $\Psi(t)$ *be a nondecreasing function of* t; *then* $E[\Psi(T_j)] \geq E[\Psi(T_k)]$, *provided* $t_j \geq t_k$.

PROOF. Proof of this lemma follows immediately from Lemma 2 of Chapter 3 in [15]. □

Let $u_i^a(t \mid \hat{\theta}(t_j))$ be the hazard rate function when the density function of T_i^a is $f_i^a(t \mid \hat{\theta}(t_j))$. The following lemma gives the ordering of $u_i^a(t \mid \hat{\theta}(t_j))$ with respect to t_j.

LEMMA 2. *Under the same conditions as in Theorem* 2, $u_i^a(t \mid \hat{\theta}(t_j))$ *is a nonincreasing function of* t_j *for every fixed* t.

PROOF. For $t_j \geq t_k$, we have

$$\frac{f_i^a(t \mid \hat{\theta}(t_j))}{f_i^a(t \mid \hat{\theta}(t_k))} \leq \frac{f_i^a(s \mid \hat{\theta}(t_j))}{f_i^a(s \mid \hat{\theta}(t_k))}$$

for all $s \geq t$. This inequality implies that

$$\frac{f_i^a(t \mid \hat{\theta}(t_j))}{f_i^a(t \mid \hat{\theta}(t_k))} \leq \frac{\int_t^\infty f_i^a(u \mid \hat{\theta}(t_j))\, du}{\int_t^\infty f_i^a(u \mid \hat{\theta}(t_k))\, du} = \frac{1 - F_i^a(t \mid \hat{\theta}(t_j))}{1 - F_i^a(t \mid \hat{\theta}(t_k))}.$$

Thus, $u_i^a(t \mid \hat{\theta}(t_j)) \leq u_i^a(t \mid \hat{\theta}(t_k))$ if $t_j \geq t_k$. □

Let $V_k^*(x, \hat{\theta}(t)) = \min_R V_k(x, \hat{\theta}(t), R)$, $k = 2, 3, 4$, where $\hat{\theta}(t)$ is used in place of $C_f(\hat{\theta}(t))$. Note in this case that the active duration of the intermittent fault is distributed with the parameter $\hat{\theta}(t)$ and that all the other distributions are known.

LEMMA 3. *Under the same conditions defined as in Theorem* 2, *if* $t_1 > t_2$, *then*

(*i*) $V_k^*(x, \hat{\theta}(t_1)) \geq V_k^*(x, \hat{\theta}(t_2))$, $k = 2, 3, 4$,

(*ii*) $r_k^*(x, \hat{\theta}(t_1)) \leq r_k^*(x, \hat{\theta}(t_2))$, $k = 1, 2$.

PROOF. The proof for $k = 3, 4$ is done by mathematical induction. Let $V_{k,n}(x, \hat{\theta}(t_j), r_2(n, j))$, $k = 3, 4$, be the expected times needed to complete the residual computation x when there are at most n retries to be attempted following the current one, and let $r_2(n, j)$ be the maximum retry duration allowed. Also, let the optimal retry duration to achieve the minimum $V_{k,n}^*(x, \hat{\theta}(t_j))$ be $r_2^*(n, j)$. For $n = 0$, $V_{4,0}(x, \hat{\theta}(t_j)) = x$ and

$$V_{3,0}^*(x, \hat{\theta}(t_1)) - V_{3,0}(x, \hat{\theta}(t_2), r_2^*(0, 1)) = \int_0^\infty \Psi(t, x, r_2^*(0, 1))\{f_i^a(t \mid \hat{\theta}(t_1)) - f_i^a(t \mid \hat{\theta}(t_2))\}\, dt,$$

where $\Psi(t, x, y) = t + x$ when $t \leq y$ and $y + x_0 + t_s$ when $t > y$. Since $\Psi(t, x, r_1^0)$ is nondecreasing in t, the right-hand side of the above equation is nonnegative as a result of Lemma 2. Also, since $V_{3,0}^*(x, \hat{\theta}(t_2))$ is the minimum when the active duration of the intermittent fault has the density function $f_i^a(t \mid \hat{\theta}(t_2))$, we get

$$V_{3,0}^*(x, \hat{\theta}(t_1)) \geq V_{3,0}(x, \hat{\theta}(t_2), r_2^*(0, 1)) \geq V_{3,0}^*(x, \hat{\theta}(t_2)).$$

Suppose that $V_{3,n}^*(x, \hat{\theta}(t_1)) \geq V_{3,n}^*(x, \hat{\theta}(t_2))$ and $V_{4,n}^*(x, \hat{\theta}(t_1)) \geq V_{4,n}^*(x, \hat{\theta}(t_2))$ $\forall x$, provided $t_1 \geq t_2$. It is obvious to see from eq. (4) that $V_{4,n+1}^*(x, \hat{\theta}(t_1)) \geq$

$V^*_{4,n+1}(x, \hat{\theta}(t_2))$ $\forall x$. Thus,

$$V^*_{3,n+1}(x, \hat{\theta}(t_1)) - V_{3,n+1}(x, \hat{\theta}(t_2), r^*_2(n+1, 1)) \geq \int_0^\infty \Psi(t, V^*_{4,n+1}(x, \hat{\theta}(t_1)), r^*_2(n+1, 1))\{f^a_i(t \mid \hat{\theta}(t_1)) - f^a_i(t \mid \hat{\theta}(t_2))\}\, dt.$$

$\Psi(t, V^*_{4,n+1}(x, \hat{\theta}(t_1)), r^*_2(n+1, 1))$ is nondecreasing in $t \leq r^*_2(n+1, 1)$. Also, since $r^*_2(n+1, 1)$ is the optimal retry duration, $V^*_{4,n+1}(x, \hat{\theta}(t_1)) \leq x_0 + t_s$. Hence, $\Psi(t, V^*_{4,n+1}(x, \hat{\theta}(t_1)), r^*_2(n+1, 1))$ is always nondecreasing in t. The right-hand side of the above equation becomes nonnegative, resulting in $V^*_{3,n+1}(x, \hat{\theta}(t_1)) \geq V_{3,n+1}(x, \hat{\theta}(t_2), r^*_2(n+1, 1)) \geq V^*_{3,n+1}(x, \hat{\theta}(t_2))$. By mathematical induction, we have $V^*_k(x, \hat{\theta}(t_1)) \geq V^*_k(x, \hat{\theta}(t_2))$ for $k = 3, 4$.

To prove $r^*_2(x, \hat{\theta}(t_1)) \leq r^*_2(x, \hat{\theta}(t_2))$, the following cases are examined. When $r^*_2(x, \hat{\theta}(t_1)) = 0$, the relation is always true. When $r^*_2(x, \hat{\theta}(t_1)) > 0$, using Lemma 2 and the first part of this proof, the derivative of $V_3(x, \hat{\theta}(t_j), R)$ with respect to the retry duration r has the following ordering relationship $\forall r$ and $t_1 \geq t_2$.

$$V^*_4(x, \hat{\theta}(t_1)) + \frac{1}{u^a_i(r \mid \hat{\theta}(t_1))} - (x_0 + t_s) \geq V^*_4(x, \hat{\theta}(t_2)) + \frac{1}{u^a_i(r \mid \hat{\theta}(t_2))} - (x_0 + t_s),$$

where all retries after the current one are assumed to employ the optimal policy. Thus, for $t_1 \geq t_2$, $r_2(x, \hat{\theta}(t_2)) = \infty$ when $r^*_2(x, \hat{\theta}(t_1)) = \infty$, and $r_2(x, \hat{\theta}(t_2)) \geq r^*_2(x, \hat{\theta}(t_1))$ when $r^*_2(x, \hat{\theta}(t_1))$ is finite.

For the case of $k = 2$, it is easy to see that $V_2(x, \hat{\theta}(t_j), R)$ is a linear combination of the effects of both transient and intermittent faults. Thus, $V_2(x, \hat{\theta}(t_j), R^*) \geq V_2(x, \hat{\theta}(t_k), R^*)$. Also, the handling of V_2 with respect to r_1 has the same ordering relationship as that of V_3 with respect to r_2. Thus, $V_2(x, \hat{\theta}(t_j), R^*) \geq V_2(x, \hat{\theta}(t_k), R^*)$, and $r^*_1(\hat{\theta}(t_j)) \leq r^*_1(\hat{\theta}(t_k))$ when $t_j \geq t_k$. □

Lemma 3 shows that $r^*_k(x, \hat{\theta}(t_j))$, $k = 1, 2$, is nonincreasing in t_j. Thus, there exists an r such that $r \geq r^*_k(x, \hat{\theta}(r))$. The proof of Theorem 2 is given as follows:

PROOF OF THEOREM 2. Suppose that the retry has been applied for the period r but the fault is still active. When $r^*_k(x, \hat{\theta}(r)) > r$, the retry should be continued since it decreases the expected task completion time. Thus, $r^*_k(x) > \sup\{r;\ r^*_k(x, \hat{\theta}(r)) > r\}$. Suppose there is an $r_j \in \{r;\ r^*_k(x, \hat{\theta}(r)) \leq r\}$. Then $V_{k'}(x, \hat{\theta}(r_j), r_j) \geq V^*_{k'}(x, \hat{\theta}(r_j)) \geq V^*_{k'}(x, \hat{\theta}(r^*_k))$, where r^*_k is defined as in eq. (8) and $k' = k + 1$. Thus, the theorem follows. □

If the maximum likelihood estimator is chosen for $\hat{\theta}(r)$, it maximizes the likelihood function:

$$L(\theta) = \left\{\prod_{i=1}^{n} \eta(I_i, t_i, \theta)\right\}\eta(1, r, \theta), \tag{9}$$

where $\eta(I, t, \theta)$ is defined as

$$\eta(I, t, \theta) = \begin{cases} f^a_i(t \mid \theta) & \text{if } I = 0, \\ 1 - F^a_i(t \mid \theta) & \text{if } I = 1. \end{cases}$$

For the same example in Section 3.2, suppose the active duration of an intermittent fault is exponentially distributed with an unknown disappearance rate μ. Using a method similar to the Cohen's derivation in [8], the maximum likelihood estimator

$\hat{\mu}(r)$ for an exponential distribution—which maximizes $\log L(\mu)$—is obtained as

$$\hat{\mu}(r) = \left(n - \sum_{i=1}^{n} I_i\right) \frac{1}{\sum_{i=1}^{n} t_i + r}. \tag{10}$$

Theorem 2 gives the optimal stopping time for the current retry. Note that the true value of μ is unknown and its maximum likelihood estimator is to determine the optimal retry duration. In the case of retry for a reappearing intermittent fault, the optimal retry duration for a given μ is either 0 or ∞ as shown in Corollary 1. Using Theorem 2 and eq. (10), we get the optimal retry duration as follows

$$r_2^* = \max\left[0, \left\{\left(n - \sum_{i=1}^{n} I_i\right) \frac{x_0 + t_s - x}{1 + \nu x} - \sum_{i=1}^{n} t_i\right\}\right]. \tag{11}$$

Note that the gamma distribution has a nondecreasing likelihood ratio for both α and β [15]. Furthermore, the estimators provided by Cohen [10] show that both the estimated α and β are increasing in the current retry period r. Thus, Theorem 2 can be applied directly when the active duration of the intermittent fault has the gamma distribution. When the distribution of the active duration is Weibull, Theorem 2 cannot be applied directly due to the fact that the Weibull distribution has a nondecreasing likelihood ratio with respect to its scale parameter only. A reasonably good approximation can be obtained by assuming that α is constant during the current retry and β is estimated using both the past and current samples as discussed above.

There are some shortcomings when the maximum likelihood estimator is used for the progressively censored samples. Particularly, the estimator is biased when the samples are censored. Also, in the case of the exponential distribution, $\hat{\mu}$ does not contain sufficient statistics of μ when the samples are censored and incomplete, that is, when there exists at least one sample (I_i, t_i) with $I_i = 1$. These shortcomings can be seen easily from a trivial example: $\hat{\mu}$ becomes zero when $I_i = 1$ for all $i = 1, 2, \ldots, n$. In fact, as shown by van Zwet [27], for most practical cases it is impossible to obtain unbiased estimators when the samples are Type I censored in a semi-infinite interval. Note, however, that there is no restriction about which estimator to be used in the foregoing determination of the optimal retry policy, meaning that estimators other than the maximum likelihood estimator can be used without altering our method described thus far.

5. *Bayes Sequential Analysis and Optimal Retry*

In the previous section, the unknown parameters of a distribution are estimated first, and the optimal retry policy is then determined using the estimated results. Notice that there could be more subsequent retries for the same intermittent fault before the task completion. Since the estimated parameter, $\hat{\theta}$, changes with the samples obtained via these retries, the usage of constant $\hat{\theta}(r)$ for further retries throughout the rest of the computation to determine r_2^* is not accurate. In other words, the point estimation approach treated in Section 4 does not include the possible variation of $\hat{\theta}$ during the subsequent retries in determining r_2^*. In this section, we shall take the Bayes approach to remedy this problem.

5.1 Optimal Retry and Bayes Decision. Let the distribution of T_i^a be governed by some unknown parameters W_i. The a priori information concerning W_i is expressed in terms of a probability distribution function defined on Ω. Let

the density function of W_i be $\xi_i(w)$. Denote further the fault characteristics, given w_i and the prior density function ξ_i, by $C_{f|w_i}$ and $C_{f|\xi_i}$, respectively.

To apply the Bayes decision theory for the retry of an intermittent fault, we define the risk with a retry policy R, given ξ_i and the residual computation x, as follows:

$$\rho_k(x, \xi_i, R) \equiv \int_\Omega V_k(x, C_{f|w_i}, R)\xi_i(w)\, dw, \qquad k = 3, 4. \tag{12}$$

Thus, the (optimal) *Bayes risk* is given as

$$\rho_k^*(x, \xi_i) \equiv \inf_R \rho_k(x, \xi_i, R), \qquad k = 3, 4. \tag{13}$$

The optimal retry duration in case of the detection of an old intermittent fault, $r_2^*(x, C_{f|\xi_i})$, abbreviated by $r_2^*(x, \xi_i)$, yields the Bayes risk $\rho_3^*(x, \xi_i)$. Similarly, the Bayes risk of the retry for a newly detected fault can be defined by eqs. (12) and (13). However, the determination of $r_1^*(x, \xi_i)$ is a one-stage Bayes decision problem. Once $\rho_4(x, \xi_i)$ and $r_2^*(x, \xi_i)$ are obtained, the normal form of analysis [2] can be applied directly for the solution of $r_1^*(x, \xi_i)$.

Following a retry attempt for an intermittent fault, regardless of whether it fails or succeeds, an event related to the fault active duration T_i^a is observed. The event observed during a retry of the duration r is either "success" or "fail." The "success" event, denoted by $e^s(t)$, occurs when the detected fault disappears after the retry duration t, which is less than or equal to the maximum allowable retry duration r. The "fail" event, denoted by $e^f(r)$, occurs when the detected fault does not disappear by the end of the retry duration r. Let $S(r) = \{e^s(t);\ t \le r\} \cup \{e^f(r)\}$. With the prior density function $\xi_i(w)$, the posterior density function following the observation of $e \in S(r)$, denoted by $\xi_i(w \mid e)$, $i = 1, 2$, becomes

$$\xi_i(w \mid e) = \frac{g(e \mid w)\xi_i(w)}{\int_\Omega g(e \mid w)\xi_i(w)\, dw}, \tag{14}$$

where $g(e \mid w)$ is the generalized conditional density function for the event e as in [11], that is,

$$g(e \mid w) = \begin{cases} \text{the density function of } T_i^a \text{ at } t & \text{if } e = e^s(t) \text{ and } t \le r, \\ \Pr(e^f(r)) & \text{if } e = e^f(r). \end{cases} \tag{15}$$

This posterior density function will become the prior density for the next retry. Consequently, the system's behavior is similar to a sequential decision procedure which determines first a retry policy and then observes the resulting sample. The procedure will be repeated with a new prior distribution which is determined on the basis of the new sample observed and the old prior distribution. The decision on retry and the sampling for fault characterization will continue as long as there is an occurrence of fault.

The problem of selecting the optimal retry policy can also be treated as the optimal stopping problem with continuous observations [13]. Suppose an intermittent fault is detected again when the residual computation is x. Then, retry is applied for a specified stopping time r. The task will be continued, without applying recovery methods other than retry, if the fault disappears during the retry period r. Otherwise, it has to be restarted from the beginning.[3] The posterior density function

[3] For simplicity, it is assumed that there is only one alternative to the retry recovery, that is, restart.

of w_i becomes $\xi_i(w \mid e^s(t))$ or $\xi_i(w \mid e^f(r))$, depending on the outcome of retry. The cost of an observation is the amount of time used for monitoring the fault until its disappearance (i.e., $c(e^s(t)) = t$) or until the end of retry (i.e., $c(e^f(r)) = r$). The costs associated with the termination of retry are defined as the amount of time necessary to complete the residual computation x as follows:

$$L(x, r, \xi_i \mid e^s(t)) = \rho_4^*(x, \xi_i(w \mid e^s(t))),$$
$$L(x, r, \xi_i \mid e^f(r)) = x_0 + t_s.$$

The expected loss for the stopping time r_2^* is the same as the Bayes risk defined in eq. (13). According to the theory presented by Irle and Schmitz in [13], there always exists an optimal stopping time, $r_2^* \in [0, \infty)$, satisfying eq. (13).

In the next section we solve the sequential decision problem using the backward induction [11] for testing hypotheses where the prior and posterior information is described by discrete probability distributions. Note that the minimax method in [2] cannot be used to solve eqs. (12) and (13), since the decision space—which consists of all possible maximum retry durations—is neither countable nor finite.

5.2 Optimal Retry and Hypotheses Testing. Suppose that there are a primary and some alternative hypotheses concerning the active duration of an intermittent fault. Consider the sequential testing of these hypotheses and the simultaneous determination of the optimal retry policy; this is not difficult to solve since both the prior and posterior probabilities lie in the same unit interval (0, 1). For given hypotheses, the initial prior distribution can be assumed to be equally likely among the hypotheses.

To be more specific, consider an example in which the active duration of an intermittent fault is assumed to be exponentially distributed with an unknown parameter μ. Let there be two hypotheses on μ, H_0 and H_1 for $\mu = \mu_0$ and $\mu = \mu_1$, respectively, and let $\mu_0 > \mu_1$. The uncertainty associated with these hypotheses can be represented by the probability h of having $\mu = \mu_0$. We first determine the optimal retry policy $\forall h \in (0, 1)$. Then, we consider the problem of testing hypotheses as well as estimating the expected sample size to reach a certain significance level under the optimal retry policy.

Consider the optimal retry duration $r_2^*(x, h)$ upon detection of an old intermittent fault. In this case, we get the posterior probabilities given the events $e^s(t)$ and $e^f(r)$, denoted by $h(t)$ and $\bar{h}(r)$, respectively, as follows:

$$h(t) = \frac{h\mu_0 e^{-\mu_0 t}}{h\mu_0 e^{-\mu_0 t} + (1 - h)\mu_1 e^{-\mu_1 t}}, \qquad \text{where} \quad t \le r, \tag{16}$$

$$\bar{h}(r) = \frac{he^{-\mu_0 r}}{he^{-\mu_0 r} + (1 - h)e^{-\mu_1 r}}. \tag{17}$$

As was discussed in Section 3.2, we can compute x_2^* for a given μ_i, denoted by $x_2^*(\mu_i)$ $i = 0, 1$, such that (i) $r_2^*(x, 1) = \infty$ if $x \le x_2^*(\mu_0)$, or 0 otherwise, and (ii) $r_2^*(x, 0) = \infty$ if $x \le x_2^*(\mu_1)$, or 0 otherwise. Since $x_2^*(\mu_0) > x_2^*\mu_1)$, $r_2^*(x, h) = \infty$ if $x \le x_2^*(\mu_1)$, and $r_2^*(x, h) = 0$ if $x \ge x_2^*(\mu_0)$. Note that the above represents extreme cases of retry, that is, retries of duration zero or infinite.

For the nonextreme case, that is, the case of $x_2^*(\mu_1) < x < x_2^*(\mu_0)$, let $h^* = \sup\{h; r_2^*(x, h) = 0\}$. Since $r_2^*(x, 1) = \infty$ and $r_2^*(x, 0) = 0$ for $x_2^*(\mu_0) < x < x_2^*(\mu_1)$, we get $0 \le h^* < 1$. For all $h > h^*$, $r_2^*(x, h) > 0$, that is, retry must be applied upon detection of a failure. Suppose retry has been applied for a small duration $\delta r < r_2^*(x, h)$. Then, the memoryless property of the exponential distribution leads

to the following equation:

$$\rho_3^*(x, h) = (1 - F_i^a(\delta r \mid h))(\delta r + \rho_3^*(x, \bar{h}(\delta r))) + \int_0^{\delta r} f_i^a(t \mid h)(t + \rho_4(x, h(t))\, dt. \tag{18a}$$

By letting $\delta r \to 0$ and changing variables, eq. (18a) becomes

$$\frac{d\rho_3^*(x, h)}{dh} = \{-f_i^a(0 \mid h)(\rho_4^*(x, h(0)) - \rho_3^*(x, h))\}\left\{\left.\frac{d\bar{h}(r)}{dr}\right|_{r=0}\right\}^{-1} = \frac{h\mu_0 + (1-h)\mu_1}{h(1-h)(\mu_0 - \mu_1)}\left\{\rho_4^*(x, h(0)) - \rho_3^*(x, h) + \frac{1}{h\mu_0 + (1-h)\mu_1}\right\}. \tag{18b}$$

On the other hand, $\rho_3^*(x, h) = x_0 + t_s\ \forall h \le h^*$. Using the same approach as in Theorem 1, we can prove that h^* satisfies the following equation:

$$\rho_4^*(x, h^*(0)) = x_0 + t_s - \frac{1}{h^*\mu_0 + (1-h^*)\mu_1}. \tag{19}$$

From eq. (4) and the definition of ρ_4^* in eq. (13), $\rho_4^*(x, h)$ is expressed as

$$\rho_4^*(x, h) = \frac{1}{\nu}(1 - e^{-\nu x}) + e^{-\nu x}\int_0^x \nu e^{\nu y}\rho_3^*(y, h)\, dy. \tag{20}$$

With the initial conditions $r_2^*(x, 1) = \infty$, $\rho_3^*(x, h)$ and $\rho_4^*(x, h)$ for $x \le x_2^*(\mu_1)$, and eqs. (18)–(20), we can calculate $\rho_k^*(x, h)$ $k = 3, 4$ $\forall x \in (x_2^*(\mu_1), x_2^*(\mu_0))$ with the following numerical algorithm:

A1. Set $h = 1$.
A2. Calculate $\rho_3^*(x, 1)$ and $\rho_4^*(x, 1)$ $\forall x \in (x_2^*(\mu_1), x_2^*(\mu_0))$.
A3. Calculate $d\rho_3^*(x, h)/dh$ using eq. (18) and $\rho_3^*(x, h - \delta h)$ $\forall x \in (x_2^*(\mu_1), x_2^*(\mu_0))$. (Note $\rho_3^*(x, h)$ and $\rho_4^*(x, h(0))$ are both known.)
A4. Calculate $\rho_4^*(x, h - \delta h)$ using eq. (20) $\forall x \in (x_2^*(\mu_1), x_2^*(\mu_0))$. (Note $\rho_3^*(x, h - \delta h)$ is known $\forall x$.)
A5. Set $h = h - \delta h$. If $h \le 0$, terminate the algorithm.
A6. If $\rho_4^*(x, h(0)) < x_0 + t_s - (1/(h\mu_0 + (1-h)\mu_1))$, go to A3. Otherwise, set $\rho_3^*(x, h - \delta h) = x_0 + t_s$ and go to A4.

From the test at A6, one can determine h^* $\forall x \in (x_2^*(\mu_1), x_2^*(\mu_0))$ so as to satisfy eq. (19). Owing to the memoryless property of the exponential distribution, $r_2^*(x, h) = 0$ when $h \le h^*$ or satisfies eq. (17) with $\bar{h}(r) = h^*$ if $h > h^*$. In Figure 9, r_2^* versus the prior probability h is plotted for various values of the residual computation x. Intersections of the curves in Figure 9 with the horizontal axis give the values of h^* for different values of x.

Remark 1. In case the active duration of an intermittent fault has a general distribution (instead of an exponential distribution), a differential equation similar to eq. (18b) cannot be obtained. In such a case, the original integral equation of $\rho_3(x, \xi_i)$, i.e., the combination of eqs. (3) and (12), has to be used instead.

From the foregoing discussion we can determine the optimal retry policy that is based on the prior probabilty h. Under this optimal retry policy, we can also determine trajectories of the posterior probabilities after a large number of occur-

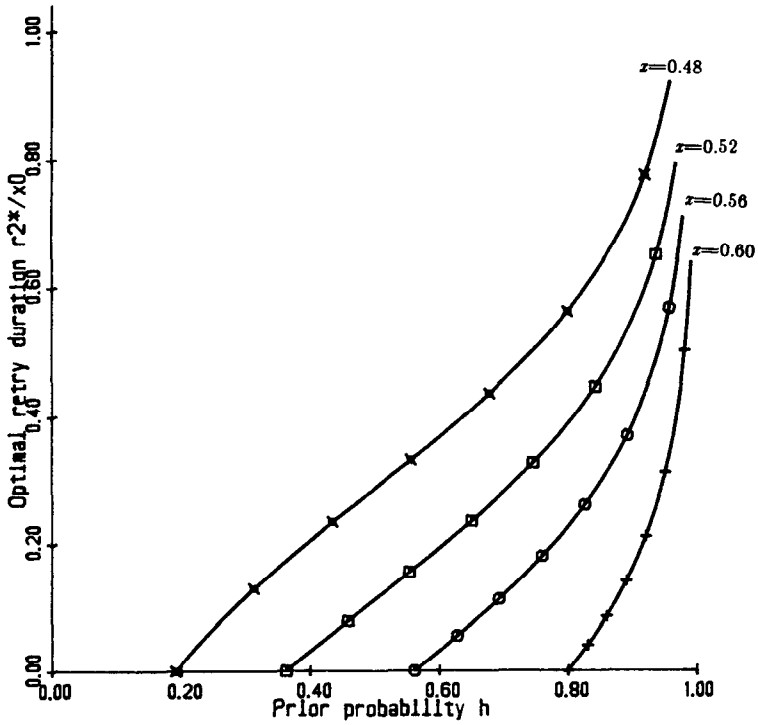

FIG. 9. The optimal retry duration $r_2^*(x, h)$ versus the prior probability h with $\mu_0 = 10$, $\mu_1 = 5$, and $\nu = 5$. $x_2^*(\mu_0) = 0.63$. $x_2^*(\mu_1) = 0.43$.

rences and reappearances of intermittent faults have been observed. Let each retry be numbered by a two-tuple (m, n_m) on the basis of occurrences and reappearances of intermittent faults. The (m, n_m)th retry is used to recover from the mth occurrence of fault in case of $n_m = 0$ or from the n_mth reappearance of the mth intermittent fault if $n_m \neq 0$. For the hypotheses H_i $i = 0, 1$, let $h_i(m, n_m)$ represent the posterior probability after the (m, n_m)th retry is applied. Also, let n'_m be the total number of reappearances of an intermittent fault during the execution of a task and $h_i(m)$ be the prior probability before the mth occurrence of fault, which is equal to $h_i(m-1, n'_{m-1})$ by definition. There are now two main problems to be addressed: (i) Will $h_i(m)$ converge to either 0 or 1, namely to the true fault characteristic as $m \to \infty$?; (ii) If converges, how fast will it converge? For convergence, we get the following theorem:

THEOREM 3. *Let* $M = \inf\{m;\ h_i(m) > 1 - \epsilon,\ or\ h_i(m) < \epsilon\}$, *where* $0 < \epsilon < 1$. *If* $0 < h_i(0) < 1$, *and* $x_0 + t_s - (1/\mu_i) > 0$ *for all hypotheses* H_i *and all tasks, then* $Pr(M < \infty) = 1$ *and* $E[M] < \infty$.

PROOF. Let $S_i(m) = \log(h_i(m)/h_j(m))$ for $j \neq i$. Thus, M can be defined as $\inf\{m;\ |S_i(m)| > K\}$, where $K = \log((1-\epsilon)/\epsilon)$. Let

$$z_i(m, n_m) = \log \frac{g[e(m, n_m) \mid \mu_i]}{g[e(m, n_m) \mid \mu_j)]},$$

where $e(m, n_m)$ is the event observed at the (m, n_m)th retry and $g(e \mid \mu_j)$ is the generalized conditional density function defined in eq. (15). (When the retry duration defined by a retry policy is zero, $e(m, n_m)$ is null and $z_i(m, n_m) = 0$. Also,

when $n_m = 0$, the retry duration is r_1^* since the fault type is not known at its first occurrence.) From Bayes theorem, we have

$$S_i(m') = S_i(m' - 1) + \sum_{n=0}^{n'_{m'}} z_i(m', n)$$

$$= \sum_{m=1}^{m'} \sum_{n=0}^{n'_{m'}} z_i(m, n) + \log \frac{h_i(0)}{h_j(0)}.$$

Let $y_i(m) = \sum_{n=0}^{n'_m} z_i(m, n)$ under the optimal retry policy. $S_i(m)$ becomes the sum of independent random variables. After an event is observed, the expected value of $z_i(m, n_m)$ is the Kullback-Leibler information number and is greater than or equal to zero when H_i is true [6]. In this case, $\mathrm{E}[z_i(m, n_m)] = 0$ if and only if the prior probability before the (m, n_m)th retry is 0 or 1. Since $x_0 + t_s - (1/\mu_i) > 0$ for all hypotheses H_i and all tasks executed, $\Pr(r_i \neq 0) > 0$ $i = 1, 2$. Hence, $\Pr[\, y_i(m) = 0] < 1$ $\forall m < M$. Following the proof in [24] that the sampling of a sequential probability-ratio test (SPRT) terminates with probability 1, $\Pr(M < \infty) = 1$ and $\mathrm{E}[M] < \infty$ are obtained. □

Remark 2. Since the tasks affected by intermittent faults do not have to be identical, the random variables $y_i(1)$, $y_i(2)$, are independently but not identically distributed. Moreover, for a fixed m, $z_i(m, n_m)$'s are dependent on one another because the events observed are controlled by the retry duration that is in turn a function of the moment of reappearance. However, all $z_i(m, n_m) \geq 0$ when H_i is true. The condition, $x_0 + t_s - (1/\mu_i) > 0$ for all hypotheses H_i and all tasks executed, indicates that retry is always a useful recovery when an intermittent fault is detected. In fact, this condition is not necessarily true for all tasks, but Theorem 3 holds as long as $\Pr(r_i^* > 0) \neq 0$.

Theorem 3 shows that the expected number of faults observed—that makes the posterior probability reach either ϵ or $1 - \epsilon$—is finite. This also holds for other distributions and retry policies as long as $r_1 \neq 0$ and $r_2 \neq 0$ for some x. However, it does not provide the average sample size, $\mathrm{E}[M \mid H_i]$ that is necessary to reach these termination boundaries K and $-K$. Also, one has to justify whether or not the posterior probability at the termination implies the true fault characteristic. In other words, it is important to know the error probability, $\Pr(S_i(M) < -K \mid H_i)$.

There are two difficult aspects in the evaluation of $\mathrm{E}[M \mid H_i]$ and $\Pr(S_i(M) < -K \mid H_i)$; one is that $y_i(m)$'s are not identically distributed, and the other is the nonexistence of closed-form solutions for both r_1^* and r_2^*. If the same task is executed repeatedly under the condition $x_0 + t_s - (1/\mu_i) > 0$ for all hypotheses, then $y_i(m)$'s become independently and identically distributed. Assume further that initially, both hypotheses are equally likely, that is, $h_0(0) = h_1(0) = 0.5$. Using the characteristics of SPRT in [11], the error probability is approximated by

$$\Pr(S_i(M) < -K \mid H_i) \approx \frac{1 - e^{-K}}{e^K - e^{-K}} \approx e^{-K}.$$

Even if the same task is executed repeatedly, it is difficult to obtain an exact solution for $\mathrm{E}[\, y_i]$ because of the dependency between the optimal retry durations and the observed samples of the active durations. This fact in turn makes it impossible to obtain the exact solution of $\mathrm{E}[M \mid H_i]$. Owing to the above difficulties, in what follows, we shall derive upper and lower bounds of $\mathrm{E}[M \mid H_i]$ instead of an exact solution.

Suppose there are two retry policies R^0 and R^1 with the retry durations (r_1^0, r_2^0) and (r_1^1, r_2^1), respectively. $r_1^0(x, h)$ and $r_1^1(x, h)$ are defined the same as $r_1^*(x, h)$.

$r_2^j(x, h)$ is equal to ∞ if $x \leq x_2^*(\mu_j)$ and 0 otherwise for $j = 0, 1$. Let $y_i^j(m)$ and M^j be $\sum_{n=0}^{n_m^j} z_i(m, n)$ and the number of faults observed to reach the termination boundaries under the retry policy R^j, respectively. Then, (i) $\Pr(M^j < \infty) = 1$ and $\mathrm{E}[M^j] < \infty$, and (ii) $\mathrm{E}[y_i^1] \leq \mathrm{E}[y_i] \leq \mathrm{E}[y_i^0]$. (Note that the indices m are omitted because of the distributions being identical.) Once $\mathrm{E}[y_i^j \mid H_i]\ j = 0, 1$ is calculated as in the Appendix 2, the expected sample size to reach the boundaries $1 - \epsilon$ and ϵ is bounded by

$$\mathrm{E}[M^0 \mid H_i] \leq \mathrm{E}[M \mid H_i] \leq \mathrm{E}[M^1 \mid H_i]$$

where

$$\mathrm{E}[M^j \mid H_i] \approx \frac{K}{\mathrm{E}[y_i^j \mid H_i]}, \qquad j = 0, 1$$

(see DeGroot [11] for more on this).

The above equations give the error probability and the bounds of the expected sample size when a certain level of significance is to be achieved. These bounds of $\mathrm{E}[M \mid H_i]$ become tight when the difference between μ_0 and μ_1 is small. Of course, the expected sample size under the optimal retry policy is larger than that for the case when the complete information about active duration is observed, that is, $r_1 = r_2 = \infty$.

Thus far, we have discussed solutions to the problem of sequential retry decision and hypotheses testing only for the case of exponentially distributed durations. Notice, however, that (i) the same method, with little modification, can be applied to the cases with any other kind of distributions, and (ii) Theorem 3 holds as long as $\Pr[y_i(m) = 0] < 1$. Moreover, the method can be extended to the testing of multiple alternative hypotheses by specifying the prior and posterior probabilities as a vector, each element of which represents the probability that the corresponding hypothesis is true.

6. *Conclusion*

In this paper, we have investigated optimal retry policies and demonstrated the use of retry to estimate the unknown fault characteristics. Although the data obtaining from retries are censored, they are the only significant means of monitoring the fault characteristics. By combining the estimation of fault characteristics and the decision of retry, the computer system performs an adaptive optimization of task completion times.

In the discussion of retry policies, retries are assumed to be continuously applied. In fact, the retry durations should be discrete since the time required for repeated execution of an operation cannot be cascaded into a single continuous duration. Since the expected risk is a continuous function of the retry duration, it is not difficult to find the optimal retry policy that is specified as a number of retry attempts.

As was pointed out in the discussion of the expected sample size for reaching a certain level of confidence in hypothesis testing, the test under the optimal retry policy turns out to be inefficient in the sense of maximizing the information observed. This is due to the fact that the optimal retry policy is defined to minimize the total completion time of the task affected by the occurrence of fault. Thus, the retry policy is a local optimum, that is, "optimal" only for the task involved. Clearly, the retry policy that gives complete maximum information should have infinite retry durations, although such a retry policy is totally unacceptable in

reality. It would be interesting to examine the trade-off between the two extreme objectives, that is, minimizing the local task completion time and maximizing the information to be collected. This problem can be formulated as the minimization of the asymptotically accumulated risk, $\lim_{m\to\infty}(1/m)\sum_{j=1}^{m} E[\rho_k(x, C_f^j)]$, where j and m are used to number the successive retries and C_f^j is the measured fault characteristic at the jth retry. It also indicates that the global optimal retry policy should collect more information (it is definitely not complete though) from the beginning to speed up the estimation of the true fault characteristics and then implement the local optimal retry policy once the true characteristics are obtained.

Another important aspect is the choice of an accurate model for the fault behavior. As was discussed in Sections 4 and 5, the optimal retry policy and the measurement of the fault characteristics are dependent on the family of density functions that are initially selected. The suitability of chosen models can be validated through goodness-of-fit tests, for example, chi-square goodness-of-fit. Although sometimes the expected task completion time may not be minimized because of the poor choice of model, the information collected via retries can still be used to check the suitability of the model. Thus, after a sufficiently large number of samples have been obtained, it is possible to select an appropriate form of density function and then achieve the minimum task completion time. The other approach is to begin with hypotheses of various forms of density functions. As sampling progresses, the parameters associated with the density function forms are estimated and then the hypotheses are tested.

The work presented in this paper is to incorporate the capability of on-line estimation (of the fault characteristics) and decision (on optimal retry policies) into the computer system. The results are a self-adjustable (thus intelligent) system and a powerful measurement of the fault characteristics. This idea can also be extended to other applications, for example, the measurement of program behavior and the simultaneous decision of system configuration or scheduling. Such extensions would be significant contributions toward the construction of highly intelligent computer systems.

Appendix A. Bounds of $V_1(x, C_f, R^*)$

Clearly, x is a lower bound of $V_1(x, C_f, R^*)$. To find an upper bound of $V_1(x, C_f, R^*)$, consider a policy R' under which only restart recovery is used upon detection of a failure. Since $r_1 = 0$ under R', we get

$$V_1(x, C_f, R^*) \le V_1(x, C_f, R') \le x + F_f(x)V_1(x_0, C_f, R')$$
$$= x + F_f(x)\frac{x_0}{(1 - F_f(x_0))}.$$

These bounds are very tight. For example, given that the fault occurrence process is exponentially distributed with MTBF = one week and the task execution time = one minute, the largest difference between the upper and lower bounds is less than 0.0001 minute.

Appendix B. The Expression of $E[y_i^j \mid H_i]$

The retry duration r_2^j under the retry policy R^j is equal to ∞ if $x \le x_2^*(\mu_j)$, and 0 otherwise. Thus, the complete information will be gathered if an old intermittent fault is detected again at $x \le x_2^*(\mu_j)$ and no information will be obtained if it is detected at $x > x_2^*(\mu_j)$. Hence, if the retry for a newly detected intermittent fault when the residual computation is x succeeds, we expect to collect information

from the successive retries before the task completion as follows:

$$\mathrm{E}[I_2 \mid x] = \mathrm{E}\left[\sum_{n=1}^{n_m^j} z_i(m, n) \mid x\right]$$

$$= \begin{cases} \nu x\left(\log \dfrac{\mu_i}{\mu_j} - \dfrac{\mu_i - \mu_j}{\mu_i}\right) & \text{if } x \leqslant x_2^*(\mu_j), \\ \exp[-\nu(x - x_2^*(\mu_j))]\nu x_2^*(\mu_j)\left(\log \dfrac{\mu_i}{\mu_j} - \dfrac{\mu_i - \mu_j}{\mu_i}\right) & \text{otherwise.} \end{cases}$$

Let the maximum retry duration for a newly detected fault be $r_1^*(x)$ when the residual computation is x. Also, let $\psi(x_d)$ be the density function of the detection time of a new intermittent fault, x_d, given that it is detected during the task execution. Then, $\psi(x_d) = \lambda_i \exp[-\lambda_i x_d]/(1 - \exp[-\lambda_i x_0])$, where $\lambda_i = p_i\lambda$. Thus, we have $\mathrm{E}[y_i^j \mid H_i]$ as follows:

$$\mathrm{E}[y_i^j \mid H_i] = \int_0^{x_0} \psi(x_0 - x)\exp[-\mu_i r_1(x)] I_1^f(r_1^*(x))\, dx + \int_0^{x_0} \int_0^{r_1^*(x)} \mu_i \exp(-\mu_i t)\psi(x_0 - x)\{I_1^s(t) + \mathrm{E}[I_2 \mid x]\}\, dt\, dx,$$

where $I_1^f(r) = -(\mu_i - \mu_j)r$ is the information collected from an unsuccessful retry of the maximum retry duration r and $I_1^s(r) = \log(\mu_i/\mu_j) - [(\mu_i - \mu_j)/\mu_i]r$ is the resulting information when the retry succeeds after the duration r.

ACKNOWLEDGMENTS. The authors are grateful to C. M. Krishna and anonymous referees for their careful comments on this paper, and to both R. Butler and M. Holt of NASA Langley Research Center for their technical assistance.

REFERENCES

1. Ball, M., and Hardie, F. Effects and detection of intermittent failures in digital systems. In *Proceedings of AFIPS Fall Joint Computer Conference*, vol. 35. AFIPS Press, Reston, Va., 1969, pp. 329–335.
2. Berger, I. O. *Statistical Decision Theory.* Springer-Verlag, New York, 1980.
3. Boone, L. A., Liebergot, H. L., and Sedmak, R. M. Availability, reliability, and maintainability aspects of the SPERRY UNIVAC 1100/60. In *Proceedings of the 10th Annual International Symposium on Fault-Tolerant Computing* (Kyoto, Japan). IEEE, New York, 1980, pp. 3–9.
4. Carter, W. C. A short survey of some aspects of hardware design techniques for fault tolerance. IBM Research Rep. RC-10811. IBM, Yorktown Heights, N.Y., 1984.
5. Carter, W. C., Putzolu, G. R., Wadia, A. B., Bouricius, W. G., Jessep, D. C., Hsieh, E. P., and Tan, C. J. Cost effectiveness of self checking computer design. In *Proceedings of the 7th Annual International Symposium on Fault-Tolerant Computing* (Los Angeles, Calif.). IEEE New York,1977, pp. 117–123.
6. Chernoff, H. *Sequential Analysis and Optimal Design.* SIAM, Philadelphia, Pa., 1972.
7. Cinlar, E. *Introduction to Stochastic Processes.* Prentice-Hall, New York, 1975.
8. Cohen, A. C. Progressively censored samples in life testing. *Technometrics 5*, 3 (Aug. 1963), 327–339.
9. Cohen, A. C. Multi-censored sampling in the three parameter Weibull distribution. *Technometrics 17*, 3 (Aug. 1975), 347–351.
10. Cohen, A. C. Progressively censored sampling in the three-parameter gamma distribution. *Technometrics 19*, 3 (Aug. 1977), 333–340.
11. DeGroot, M. H. *Optimal Statistical Decision.* McGraw-Hill, New York, 1970.
12. Droulette, D. L. Recovery through programming system/360-System/370. In *Proceedings of the AFIPS Spring Computer Conference*, vol. 38. AFIPS Press, Reston, Va., 1971, pp. 467–476.

13. IRLE, A., AND SCHMITZ, N. Decision theory for continuous observations I: Bayes solutions. In *Transactions of the 7th Prague Conference on Information Theory, Statistical Decision Functions and Random Processes.* 1974, pp. 209–221.
14. KOREN, I., AND SU, S. Y. H. Reliability analysis of N-modular redundancy systems with intermittent and permanent faults. *IEEE Trans. Comput. 28*, 7 (July 1979), 514–520.
15. LEHMANN, E. L. *Testing Statistical Hypotheses.* Wiley, New York, 1959.
16. LEMON, G. H. Maximum likelihood estimation for the three parameter Weibull distribution based on censored samples. *Technometrics 17*, 2, 1975, 247–254.
17. MAESTRI, G. H. The Retryable Processor. In *Proceedings of AFIPS Fall Joint Computer Conference* vol. 41. AFIPS Press, Reston, Va., 1972, pp. 273–277.
18. ROSS, S. M. *Stochastic Processes.* Wiley, New York, 1983.
19. SHEDLETSKY, J. J. The error latency of a fault in a sequential digital circuit. *IEEE Trans. Comput. C-25*, 6 (June 1976), 655–659.
20. SHEDLETSKY, J. J., AND MCCLUSKY, E. J. The error latency of a fault in a combinational digital circuit. In *Proceedings of the 5th Symposium on Fault-Tolerant Computing* (Paris, France). IEEE, New York, 1975, pp. 210–214.
21. SHIN, K. G., AND LEE, Y. H. Error detection process: Model, design, and its impact on computer performance. *IEEE Trans. Comput. C-33*, 6 (June 1984), 529–540.
22. SHIN, K. G., AND LEE, Y. H. Measurement and application of fault latency. *IEEE Trans. Comput. C-35*, 4 (Apr. 1986), 370–375.
23. SIEWIOREK, D. P., AND SWARZ, R. S. *The Theory and Practice of Reliable System Design.* Digital Press, Educational Services, Digital Equipment Corporation, Bedford, Mass., 1982.
24. STEIN, C. A note on cumulative sums. *Ann. Math. Stat. 17* (1946), 489–499.
25. STIFFLER, J. J., AND BRYANT L. A. CARE III phase report—Mathematical description. NASA Rep. 3566. NASA, Washington, D.C., Nov. 1982.
26. TASAR, O., AND TASAR, V. A study of intermittent faults in digital computers. In *Proceedings of AFIPS National Computer Conference*, vol. 46. AFIPS Press, Reston, Va., 1977, pp. 807–811.
27. VAN ZWET, W. R. Bias in estimation from type I censored samples. *Statist. Neerlandica 20* (1966), 143–148.
28. WINGO, D. R. Solution of the three-parameter Weibull equations by constrained modified quasilinearization (progressively censored samples). *IEEE Trans. Reliability R-22*, 2 (June 1973), 96–102.

RECEIVED MAY 1984; REVISED MARCH 1985; JULY 1986; ACCEPTED JANUARY 1987