Minimizing the Time to Troubleshoot a Failed System

E. J. Subelman
Communications Systems Research Section
and
University of California, Los Angeles

This report presents a procedure to determine the order in which to inspect the components of a failed series system so as to minimize the expected time until the failed component is isolated. Our model includes the case in which inspection of a component may falsely indicate the component is functioning.

I. Introduction

In many of the systems and subsystems of the Deep Space Network, one is faced with the problem of troubleshooting a complex system that has failed, that is, with isolating the failed components in order to repair or replace them.

The system's down time can be divided into two phases: troubleshooting and repair. Usually one has little control on the duration of the second phase, except by finding better repair procedures, which may be difficult to do. Furthermore, in the case of modern electronic equipment, repair time consists of the replacement of a circuit board so that the only time to consider is the procurement time for that board.

On the other hand, one usually has a greater degree of control on the troubleshooting procedures, and reducing the total time to detect the fault provides us with a simple way of reducing total system downtime.

Thus the motivation for this paper: to determine the procedures to be followed when troubleshooting a system so as to minimize the total time spent in this phase.

II. The Model

We will consider a series system of n statistically independent components, that is, a system in which the failure of any one component causes the failure of the system. Suppose the system has just failed, so that we know that one component has failed. In order to repair the system, one would like to determine which component has caused the failure, and this usually requires inspection of the components.

Assume that inspection of component i requires an amount of time T_i. Furthermore assume that, if component i is indeed failed and we inspect it, then with probability α_i we will detect the failure, and with probability $\beta_i = 1 - \alpha_i$ we will overlook the failure. Clearly perfect inspection corresponds to $\alpha_i = 1$, and when $\alpha_i < 1$, more than one inspection of component i may be necessary to uncover the fault. We will assume that $T_i > 0$ for all i. Our goal is to find an inspection plan that minimizes the total expected time until the faulty component is detected.

Let $F_i(t) = P$ (life of component $i \leqslant t$); that is, F_i is the distribution of the life of component i. Let $f_i(t) = (d/dt)\, F_i(t)$

and $\overline{F}_i(t) = 1 - F_i(t)$. We also denote by $r_i(t)$ the failure rate

$$r_i(t) = \frac{f_i(t)}{\overline{F}_i(t)}$$

Although the system will initially have all new components, after several repairs have been completed, the ages of the components will be different. We will, however, assume that a repaired component has the same life distribution F_i. This assumption is not necessary but makes our notation simpler.

Let A_i be the age of the component i at the time of the last repair of the system. Let x be the age of the system at the time of failure. Then it can be shown (see Appendix A) that the probability that component i caused the system failure is

$$p_i = \frac{r_i(x + A_i)}{\sum_{j=1}^{n} r_j(x + A_j)} \qquad (1)$$

We remark that for the subsequent development, all we require is the value of p_i, which could be modified to incorporate any additional information available, such as might be obtained by observation of the symptoms of the system's failure, or even subjective estimates. We will also make the assumption that all components must be inspected; that is, even if $p_i = 1$ we must still inspect component i, presumably to pinpoint the cause of failure in order to carry out repairs.

III. The Optimal Policy

With this information we can proceed to determine an optimal inspection plan, which is simply a list of components $i_1, i_2, i_3, \cdots$ with the interpretation: inspect component i_1. If you do not find it failed, proceed to component i_2, etc. stopping when the failed component is found. The list may contain repetitions in the case of imperfect inspection.

If we inspect component j and do not find it failed, then this information can be incorporated by modifying the values of p_i $i = 1, 2, \cdots, n$ by means of Bayes' theorem, to obtain

$$p_i' = \begin{cases} \dfrac{p_i}{1 - \alpha_j p_j} & \text{for all } i \neq j \\[2ex] \dfrac{(1 - \alpha_j)p_j}{1 - \alpha_j p_j} & \text{for } i = j \end{cases} \qquad (2)$$

We can think of the inspection procedure in terms of the set of $\{p_i\}_{i=1}^{n}$. Given the initial set, we select the first component to inspect. If we do not find it failed, we use (2) to modify the probabilities and incorporate this additional information, and then select a new component to inspect; the process continues until we find the failed component.

It is clear that if $p_i > 0$ then the optimal search policy will eventually examine component i since otherwise it would have an infinite expected duration.

The following result characterizes the optimal policy:

Theorem. The optimal policy always inspects the component yielding

$$\min \left\{ \frac{T_i}{p_i \alpha_i} \right\} .$$

Proof. Suppose, without loss of generality, that

$$\frac{T_1}{p_1 \alpha_1} = \min \left\{ \frac{T_i}{p_i \alpha_i} \right\},$$

and suppose component 1 is not inspected first. Assume component 2 is the last component inspected before component 1, and that this is the $(a + 1)$st inspection of component 2.

The inspection sequence is then $j_1, j_2 \cdots j_k, 2, 1, \cdots$ with component 2 appearing a times in the first k entries. We will prove that the sequence $j_1, j_2 \cdots j_k, 1, 2, \cdots$ gives a smaller expected inspection time.

For the first sequence, the expected inspection time is, conditioning on which is the failed component,

$$V_1 = p_1 [M_1 + T_2 + T_1 + N_1]$$

$$+ p_2 [M_2 + \beta_2^a T_2 + \beta_2^{a+1} T_1 + N_2]$$

$$+ (1 - p_1 - p_2) N_3$$

where M_1 and M_2 depend only on $j_1 \cdots j_k$ and N_1, N_2 and N_3 depend on the tail of the sequence. These terms are not affected by the interchange of 1 and 2 and so their form is not important for our purposes.

For the second sequence the expected time is

$$V_2 = p_1[M_1 + T_1 + \beta_1 T_2 + N_1]$$

$$+ p_2[M_2 + \beta_2^a T_1 + \beta_2^a T_2 + N_2]$$

$$+ (1 - p_1 - p_2)N_3$$

Clearly $V_1 \geqslant V_2$ if and only if

$$p_1 \alpha_1 T_2 \geqslant \beta_2^a p_2 \alpha_2 T_1$$

Thus, if

$$\frac{T_1}{p_1\alpha_1} < \frac{T_2}{p_2\alpha_2}$$

then for any $a \geqslant 0$

$$\beta_2^a p_2 \alpha_2 T_1 < p_1 \alpha_1 T_2$$

and therefore

$$V_2 < V_1 .$$

It follows that component 1 should be inspected first.

The procedure for inspection of the system can be summarized as follows:

(1) Initially $A_1 = A_2 = \cdots A_n = 0$ (all new components).

(2) Measure the system uptime until failure, and call it x.

(3) For each $i = 1, 2, \cdots n$ compute p_i using (1) or its equivalent (3) in Appendix A.

(4) Inspect the component giving

$$\min\left\{\frac{T_i}{p_i\alpha_i}\right\},$$

say component j. If it is failed go to step 6.

(5) For each i replace p_i by p_i' obtained by using (2), and go to step 4.

(6) Component j failed. Repair or replace it. Set $A_j = 0$ and $A_i = A_i + x$ for all $i \neq j$. Let the system operate and return to step 2.

In the special case in which inspection of the components yields perfect information on their state ($\alpha_1 = \alpha_2 \cdots = \alpha_n = 1$) the optimal strategy has an even simpler form:

Corollary If $\alpha_i = 1$ for all i then the optimal policy inspects the components in the order of increasing values of T_i/p_i.

Proof By the theorem above, the first component to be inspected is that with the smallest value of T_i/p_i, say component 1. If it is not found failed, then the p_i will be modified using (2) to

$$p_i' = p_i/(1 - p_1) \quad \text{for } i = 2, 3, \cdots n$$

$$p_1' = 0$$

Thus the new ratios will be

$$\infty, T_2(1 - p_1)/p_2, T_3(1 - p_1)/p_3, \cdots T_n(1 - p_1)/p_n$$

Since the new ratios are (excepting that for component 1) the original ratios scaled by $(1 - p_1)$, their ordering remains the same, so that the next component to be inspected should be the one with the second smallest value of T_i/p_i.

IV. An Example

We consider a system consisting of three components, all of them having a Weibull life distribution

$$F(t) = 1 - e^{-(\lambda t)^a}$$

The values of the parameters are

Component	λ	a	Inspection time
1	0.02	1/2	60 min
2	0.01	1	50 min
3	0.008862	2	30 min

It can easily be verified that all three components have a mean lifetime of 100 hours. Component 1 has a decreasing failure rate, component 2 a constant failure rate and component 3 an increasing failure rate.

For the Weibull distribution, the failure rate is

$$r(t) = a\lambda t)^{a-1}$$

Suppose all three components are new initially ($A_1 = A_2 = A_3 = 0$) and the system fails after 40 hours of operation. We can compute the probability that each component was the cause of failure by using (1)

$$r_1(40) = 0.5 \times 0.02 \times (0.02 \times 40)^{0.5-1} = 0.01118$$

$$r_2(40) = \qquad = 0.01$$

$$r_3(40) = \qquad = 0.006283$$

so that

$$p_1 = \frac{0.01118}{0.01118 + 0.01 + 0.006283} = 0.4071$$

$$p_2 = \qquad = 0.3641$$

$$p_3 = \qquad = 0.2288$$

Suppose that inspecting any of the three components always determines whether the component is functioning or failed ($\alpha_1 = \alpha_2 = \alpha_3 = 1$).

We can now see that component 1 is the most likely to be failed, but requires the longest inspection time. Using the corollary to find the order of inspection, we compute

$$\frac{T_1}{\alpha_1 p_1} = \frac{60}{1 \times 0.4071} = 147.38$$

$$\frac{T_2}{\alpha_2 p_2} = \qquad = 137.32$$

$$\frac{T_3}{\alpha_3 p_3} = \qquad = 131.11$$

Thus, according to the corollary, we should inspect the components in the order 3, 2, 1: if component 3 is not found failed we will next inspect component 2 and if it is functioning we will examine component 1.

As a further illustration, we will exhibit the computations necessary for this system assuming the corollary does not apply. If inspection of component 3 finds it functioning, we use (2) to recompute the values of p_1, p_2, p_3

$$p_1 = \frac{0.4071}{1 - 1 \times 0.2288} = 0.5279$$

$$p_2 = \qquad = 0.4721$$

$$p_3 = \frac{(1 - 1) \times 0.2288}{1 - 1 \times 0.2288} = 0$$

These probabilities reflect the fact that we now know component 3 is functioning.

We now find which component to inspect next.

$$\frac{T_1}{\alpha_1 p_1} = \frac{60}{1 \times 0.5279} = 113.65$$

$$\frac{T_2}{\alpha_2 p_2} = \frac{50}{1 \times 0.4721} = 105.90$$

$$\frac{T_3}{\alpha_3 p_3} = \frac{30}{1 \times 0} = \infty$$

Thus we next test component 2, and, if found functioning, we proceed to component 1.

V. Conclusions and Suggestions for Further Work

We have presented a procedure for the determination of the sequence in which to test the components of a failed series system in minimal expected time. The procedure requires an estimate of the lifetime distributions of the different components and of their testing times, and requires a minimal amount of computation to determine the optimal order. Additional information required is the age of each component as well as the age of the system since it was last repaired.

This study suggests the following questions that remain to be analyzed: Under what conditions would it be preferable to replace a component by a new one without inspecting it for failure, and under what conditions would it be preferable to replace a component only after it has been tested and found defective? Furthermore, in what order should the replacements or tests of different components be carried out?

Appendix A
Probability Derivation

We now present a derivation of the probability that component 1 caused the system's failure.

If A_i was the age of component i when the last system repair was completed, then the probability it will function for an additional t units of time is

$$\overline{G}_i(t) = \frac{\overline{F}_i(t + A_i)}{\overline{F}_i(A_i)}$$

Let $G_i(t) = 1 - \overline{G}_i(t)$ and $g_i(t) = (d/dt)\, G_i(t)$.

Assume X_i is the remaining life of component i after the last system repair, and let $X = \min\,\{X_i\}$ be the time between the last repair and the next failure of the system. Clearly

$$P\,[X_i > t] = \overline{G}_i(t) \quad i = 1, 2, \cdots n, \text{ and}$$

$$P\,[X > t] = \prod_{i=1}^{n} \overline{G}_i(t)$$

If it is known that the system failed between times x and $x + \delta$, then the probability that component 1 caused the system's failure is

$$p_1(\delta) = P\,[X = X_1 \,|\, x < X \leqslant x + \delta]$$

$$= \frac{P\,[x < X_1 \leqslant x + \delta, X_2 > X_1, X_3 > X_1 \cdots X_n > X_1]}{P\,[x < X \leqslant x + \delta]}$$

The numerator can be computed by conditioning on the value of X_1, yielding

$$p_1(\delta) =$$

$$\frac{\int_x^{x+\delta} g_1(t)\,\overline{G}_2(t) \cdots \overline{G}_n(t)\,dt}{\overline{G}_1(x)\,\overline{G}_2(x) \cdots \overline{G}_n(x) - \overline{G}_1(x+\delta)\,\overline{G}_2(x+\delta) \cdots \overline{G}_n(x+\delta)}$$

Dividing numerator and denominator by δ and letting δ go to 0 we obtain

$$p_1 = p\,[X = X_1 \,|\, X = x] = \frac{g_1(x)\,\overline{G}_2(x) \cdots \overline{G}_n(x)}{\sum_{k=1}^{n} g_k(x) \prod_{j \neq k} \overline{G}_j(x)}$$

and using the definition of $\overline{G}_i$ and g_i this reduces to

$$p_1 = \frac{f_1(x + A_1) \prod_{j=2}^{n} \overline{F}_j(x + A_j)}{\sum_{k=1}^{n} f_k(x + A_k) \prod_{j \neq k} \overline{F}_j(x + A_j)} \qquad (3)$$

We now recall that $r_i(t) = f_i(t)/\overline{F}_i(t)$ to obtain equation (1):

$$p_i = \frac{r_i(x + A_i)}{\sum_{j=1}^{n} r_j(x + A_j)}$$