AT&T Bell Laboratories Technical Journal
Vol. 63, No. 6, July–August 1984
Printed in U.S.A.

System Sparing for Minicomputer-Based Operations Systems

By D. W. TOLLETH*

(Manuscript received March 29, 1983)

A queueing model was developed to derive guidelines for deploying system spares—complete minicomputer systems used as backup for failed minicomputer-based Operations Systems (OSs). The model is a three-dimensional set of queueing equations incorporating a time-dependent number of repairers, and exponential random variables for the time between failures for each minicomputer system, the time to repair each system, and the time to switch between failed systems and spares. Guidelines derived from numerical solution of the model are being used by the Bell Operating Companies and AT&T Communications to aid planning studies for (1) proving-in system spares, (2) meeting specific OS availability objectives, and (3) improving minicomputer maintenance staff utilization.

I. INTRODUCTION

1.1 Background

Development of minicomputer-based Operations Systems (OSs) to support Bell System operations began in the early 1970s. By the end of 1983 over 5000 such OSs were deployed by the Bell Operating Companies (BOCs) and AT&T Communications.

System spares are complete minicomputer systems that provide backup computing when minicomputer-based OSs fail. Increasing dependence on minicomputer systems to perform daily work functions has led the BOCs and AT&T Communications to deploy system spares

*AT&T Bell Laboratories.

Copyright © 1984 AT&T. Photo reproduction for noncommercial use is permitted without payment of royalty provided that each reproduction is done without alteration and that the Journal reference and copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free by computer-based and other information-service systems without further permission. Permission to reproduce or republish any other portion of this paper must be obtained from the Editor.

for OSs clustered in Minicomputer Maintenance and Operations Centers (MMOCs).

This paper presents a comprehensive solution to the system sparing problem. The BOCs and AT&T Communications are using guidelines based on this solution for determining optimal system sparing levels and for meeting specific availability objectives. The analysis includes transient solutions that examine possible improvements in maintenance staff utilization with sparing.

The work presented here has potential applications in system design, as well as in BOC and AT&T Communications system and staff planning. OSs with stringent availability objectives have usually been designed in duplex or triplex arrangements, with one or two active and one backup processor. For such systems with large deployments in clustered environments, this work shows how to meet stringent availability objectives with fewer backup systems.

1.2 Approach

We achieved the desired results in three steps. First, we characterized the MMOC equipment, staff, and functions pertinent to system sparing, including staffing levels, system failure rates, and mean times to repair. Second, we developed a mathematical model of the sparing process. Third, we obtained system availability data from the model for a range of parameters representing current BOC and AT&T Communications MMOC operations.

1.3 Overview

Section II gives an in-depth discussion of the model and the parameters chosen to characterize the BOC and AT&T Communications MMOCs. Section III presents the mathematical details of the model, including the principal equations and state diagrams. Section IV presents availability data obtained from the model. The presentations of the data are designed to aid planning studies for proving-in spare systems, meeting specific availability objectives, and improving maintenance staff utilization. Section V presents some observations about system sparing.

II. DESCRIPTION OF THE SYSTEM SPARING MODEL

2.1 Model features

How does system sparing affect BOC and AT&T Communications minicomputer users and the MMOCs? To answer this question, we must first characterize the equipment, personnel, and functions involved with system sparing. We can then develop a mathematical model to quantify the effects of sparing.

Minicomputer systems in an MMOC can be located in a single building or distributed throughout a geographic area, with some systems clustered and some standing alone. Our interest in sparing suggests that we break the systems into two groups: (1) a cluster of systems with one or more spares in one location, and (2) the other, nonspared systems in the MMOC (see Fig. 1). The second group may be scattered geographically, clustered at another location, or colocated with the first group. To simplify the model, the second group is sized so that the total number of minicomputers is fixed and corresponds to the typical number of systems for which one repairer would be responsible in the field. This number of systems, 25, is based on current maintenance force staffing levels in BOCs doing self-maintenance, i.e., doing their own hardware maintenance, instead of contracting with a vendor.

Three personnel groups would be affected by system sparing. First, the minicomputer operators, who are located at the minicomputer

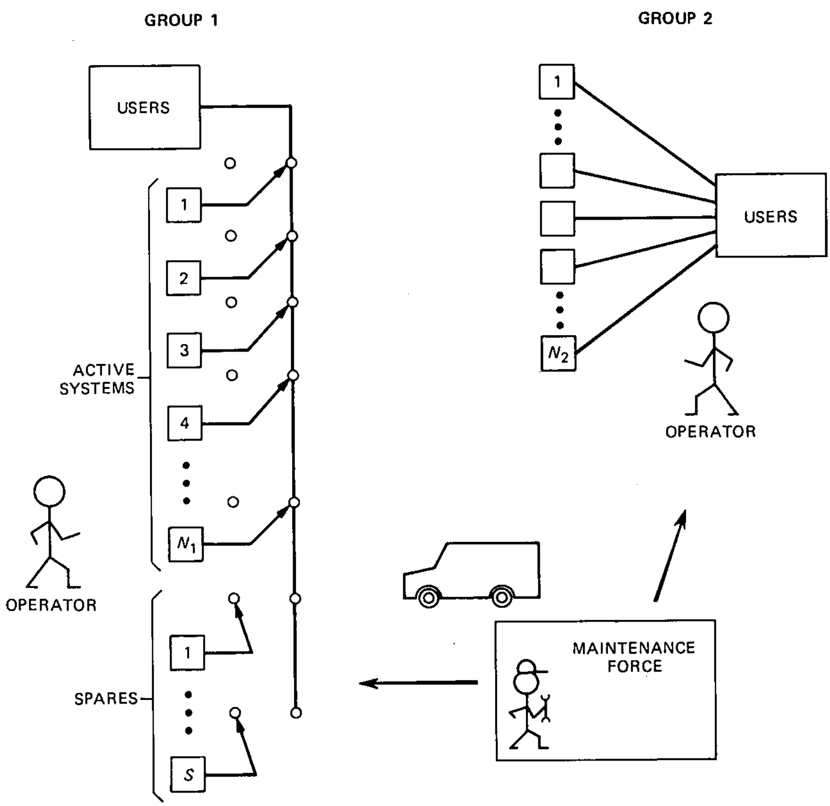

Fig. 1—MMOC systems and personnel.

sites, would be required to do the switching from failed systems to spares and back. Second, the minicomputer maintenance personnel would feel less pressure to repair a failed system if a spare were available. Finally, the minicomputer users would see improved system availability. We will see that availability can improve for all users, including those whose systems are not equipped with a spare.

Users are generally at locations remote from the minicomputer systems and communicate with them via private line circuits. These are the circuits that the operators must switch when moving from a failed system to a spare. Operators switch the data circuits and load the database of the failed system onto the spare.

The effects of sparing on maintenance are broader and more subtle than its effects on operations. When a cluster has a spare, a failed system does not demand immediate attention; the spare can be switched. Thus, when multiple failures occur, priority can be given to the systems for which no spare is available. The effect is more timely repair of the systems without a spare. Of course, users of systems equipped with a spare also enjoy improved availability. Both groups benefit.

In addition to modifying repair priorities, sparing also levels the work load of the maintenance staff: minicomputers can be queued for repair instead of maintenance staff being queued for minicomputer failures. This may allow reduction in the size of the maintenance staff or elimination of nighttime maintenance. The latter will be possible if sparing sufficiently reduces the risk of user outage at the beginning of the morning shift.

For simplicity, the model omits several common practices. First, the model considers only corrective maintenance activities and ignores the value of a spare for preventive maintenance and database management activities. Second, the model omits the common practice of maintenance personnel working overtime at the end of a shift to finish a repair.

The omission of these two features means that the model underestimates the benefits of sparing, because both practices increase the value of sparing. Spare systems can be switched for working systems, allowing preventive maintenance to be performed during the day or evening, rather than during the night shift or on weekends; and overtime at the end of a shift gets systems back on line faster, improving system availability.

The model also omits the step of "switching back" repaired systems for spares. This task is usually performed at night or at some other time when there is little or no demand for a system, so the omission should have little effect on the results.

The system sparing process was modeled mathematically, as described in Section III. The parameters used in the model are discussed below.

2.2 Model parameters—operating conditions in the MMOC

Data collected by the MMOC planning group at AT&T Bell Laboratories indicate that providing 24-hour, 7-day maintenance coverage in a clustered environment requires one repairer per shift for every 25 "COSMOS-like" systems, i.e., nominal *DEC PDP-11/70**-based systems.

Because a repairer covers 25 systems and we want the number of repairers to be an integer, we need to model 25 systems, 50 systems, etc. The minimum set—25 systems, 1 repairer and 1 spare—can be shown to satisfy the objectives of this study.

When a spare is added, the model assumes no increase in the size of the maintenance staff. The repairer who was responsible for 25 systems before sparing is thus responsible for 26 after sparing.

Our data show that systems in the field incur about 100 hours of downtime per year for corrective maintenance. One hundred hours per year corresponds roughly to a mean time between failures of 22 days and a mean time to repair of 6 hours. (The mean time to repair includes both the time to respond to and to fix the trouble.) There is considerable variation in downtime, due to such factors as the environment in which a minicomputer is operated, whether a system is maintained by on-site BOC personnel or off-site vendor personnel, and whether downtime measured is user downtime or system downtime. To allow for these variations, the model was run with 50, 100, and 150 hours of downtime per year.

Switching between a failed system and a spare requires switching the data lines and moving the database. Moving the database, which includes removing the disk packs, moving them to the spare system, and checking the files, is the rate-limiting step, and typically requires about 15 minutes. The time to move the data lines varies, depending on the switching method, but it is usually much less than 15 minutes. The model assumed 15 minutes total for switching.

To permit examination of possible improvements in maintenance staff utilization, the model allows the number of repairers on a shift to change. In particular, the model accommodates no, one, or infinite repairers. No repairers is the case used to model uncovered shifts. One repairer is the nominal case for 25 systems and 1 spare. Infinite repairers is the case used as an approximation for the two-repairer case to examine movement of the night shift personnel to the day

---

* Trademark of Digital Equipment Corporation.

shift. The validity of the infinite repairer approximation is discussed in Section III.

In summary, the model assumes 25 systems; 1 spare; and 0, 1, or ∞ repairers. The time required to switch a spare for a failed system is 15 minutes. Data were produced for systems with 50, 100, and 150 hours of downtime per year.

III. TIME-DEPENDENT QUEUEING MODEL

System sparing increases system availability and allows improvements in maintenance staff efficiency. System availability and the number of repairers required to maintain the systems are, therefore, the principal indicators of the effects of sparing.

We want to develop a mathematical model that predicts average system availabilities with and without sparing, as a function of time. It must keep track of three groups of systems: (1) systems with access to spares (group 1), (2) spare systems, and (3) other systems under the purview of the maintenance staff (group 2). A novel feature of this grouping is the ability to analyze sparing for differing numbers of systems per spare while maintaining a constant work load on the maintenance staff.

The model must account for changes in maintenance force staffing levels, for example at shift changes. Since the mean time to repair (6 hours) is of the order of a shift (8 hours), steady-state solutions will not be attained within a shift. So, we must develop a time-dependent model with a time-dependent number of servers.

We begin by stating assumptions for the model and defining notation. Then, state diagrams and state equations are presented. Finally, we discuss the numerical methods used and checks made on the model.

3.1 Assumptions and notation

We assume that the time between failures for each machine is an independent and identically distributed (i.i.d.) exponential random variable. Switching time and response plus repair time are also i.i.d. exponential random variables. The random variable for switching time is probably closer to deterministic. However, an exponential distribution produces more congestion in the model than a deterministic distribution. So, the model errs on the side of less recovered availability. Our conclusions, thus, are conservative.

States are described by $p(t; n_1, s, n_2)$, the probability that at time t, n_1 systems are failed in group 1, s spares are failed, and n_2 systems are failed in group 2.

We define failure rates λ_1 and λ_2 for systems in group 1 and group 2, respectively. Spares are identical to systems in group 1 and have the same failure rate, λ_1.

The model considers three cases:

1. No repairers—the case to consider for late-night and weekend shifts.

2. One repairer—the nominal case because, for purposes of analysis, the total number of systems modeled (group 1 plus group 2) is sized for one repairer on duty 24 hours per day.

3. Infinite repairers—the case used to approximate the two-repairer case.

There are three different state diagrams and three sets of state equations. The time-dependent parameter is the number of repairers, which can change at 8-hour intervals, corresponding to work shifts. In all cases the time to repair is assumed to be an i.i.d. exponential random variable, with mean $1/\mu$, which is the same for all systems. This time is assumed to include response time. Finally, the time required for an operator to switch a spare for a failed system is also an i.i.d. exponential random variable, with mean $1/\mu_s$.

To summarize, the model is characterized by the following:

$p(t; n_1, s, n_2)$—probability that at time t, n_1, s, and n_2 systems are in the failed state from the total pool of N_1 systems in group 1, S spares, and N_2 systems in group 2, respectively.

λ_1—failure rate for group 1 systems and spare systems.

λ_2—failure rate for group 2 systems.

μ—repair rate for a system (single repairer rate).

μ_s—rate to switch a spare for a failed system.

3.2 State diagrams and state equations

Consider now the three cases for repair.

Case 1: No repairers—Figure 2 shows the state diagram for a general state (n_1, s, n_2) with no repairers. The corresponding state equation describing the probability at time $t + \Delta t$ of being in state (n_1, s, n_2) is, leaving out terms $o(\Delta t)$,

$$\begin{aligned} p(t + \Delta t; n_1, s, n_2) = [1 &- (N_1 - n_1)\lambda_1\Delta t - (S - s)\lambda_1\Delta t \\ &- (N_2 - n_2)\lambda_2\Delta t - \mu_s\Delta t]p(t; n_1, s, n_2) \\ &+ (N_1 - n_1 + 1)\lambda_1\Delta t\, p(t; n_1 - 1, s, n_2) \\ &+ (S - s + 1)\lambda_1\Delta t\, p(t; n_1, s - 1, n_2) \\ &+ (N_2 - n_2 + 1)\lambda_2\Delta t\, p(t; n_1, s, n_2 - 1) \\ &+ \mu_s\Delta t\, p(t; n_1 + 1, s - 1, n_2). \qquad (1) \end{aligned}$$

Switching spares for failed systems in group 1 continues until all

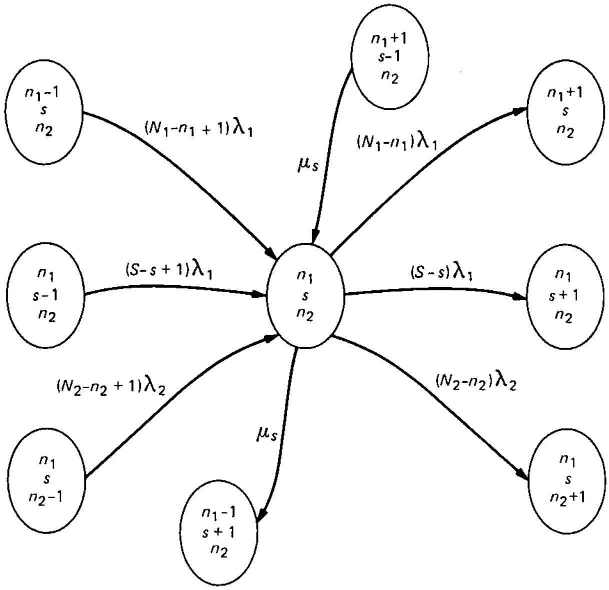

Fig. 2—State diagram for the general state (n_1, s, n_2) with no repairers.

spares are failed ($s = S$). At that point, switching stops. For all states with $s = S$, μ_s is set to zero.

There are two observations that should be made at this point. First, the group 2 part of eq. (1) is completely separable from the part for group 1 and the spares. As we discussed later, separability was used as a check on the numerical solution. We can write

$$p(t;\, n_1,\, s,\, n_2) = p(t;\, n_1,\, s)p(t;\, n_2) \tag{2}$$

and obtain a closed-form solution for $p(t; n_2)$. For ease of programming, this separation was not made in the model. Second, it is clear that the steady-state solution for this case is $p(\infty; N_1, S, N_2) = 1$. That is, with no repairer all systems are failed in the limit $t \rightarrow \infty$.

Case 2: One repairer—Adding a repairer requires four criteria in the model that determine which system the repairer will fix for a given state of the system. First, in the field the repairer probably fixes systems on a first-in first-out basis. Since all failure, repair, and switching times are i.i.d. exponential random variables, the state space is memoryless; it is impossible to determine which minicomputer failed first. For predicting minicomputer system availability with this model, the equivalent of a first-in first-out repair strategy is random repair. The probability that a given system is under repair is defined such that all failed systems have the same opportunity to be repaired. Second, the purpose of spares is to replace failed group 1 systems.

Thus, group 1 systems are not repaired if spares are available; they are switched. Third, spares are not repaired if there are any failed group 2 systems. Finally, when spares are available for failed group 1 systems, priority is given to repair of failed group 2 systems. In summary, the decision criteria for repair are:

1. Randomly repair the failed systems in group 1 and group 2 if all the spares have failed.
2. Never fix group 1 systems if spares are available.
3. Repair spares only when all group 2 systems are working and there are spares available for any failed group 1 systems.
4. Always fix group 2 systems first unless all the spares have failed.

The parameter $\chi(n_1, s, n_2)$ is defined to incorporate these criteria. χ is the probability that either (1) a system in group 1 will be fixed, when $s = S$, or (2) one of the failed spares will be fixed, when $s < S$. The only exception is that χ is the probability that a spare will be fixed, when $s = S$ and $n_1 = 0$. The probability that a system in group 2 will be fixed is $1 - \chi$. The value of χ corresponding to each of the above criteria is:

1. $\chi(n_1, s, n_2) = \dfrac{n_1}{n_1 + n_2}$, for $s = S$, $n_1 + n_2 > 0$.
2. $\chi(n_1, 0, 0) = 0$, for $S > 0$.
3. $\chi(n_1, s, n_2) = 1$, for $s < S$, $n_2 = 0$; or $s = S$, $n_1 = n_2 = 0$.
4. $\chi(n_1, s, n_2) = 0$, for $s < S$, $n_2 > 0$.

Figure 3 shows the state diagram for a general state (n_1, s, n_2) with one repairer. The corresponding state equation describing the probability at time $t + \Delta t$ of being in state (n_1, s, n_2) is, leaving out terms $o(\Delta t)$,

$$\begin{aligned} p(t + \Delta t; n_1, s, n_2) = {} & [1 - (N_1 - n_1)\lambda_1 \Delta t - (S - s)\lambda_1 \Delta t \\ & - (N_2 - n_2)\lambda_2 \Delta t - \mu \Delta t - \mu_s \Delta t] p(t; n_1, s, n_2) \\ & + (N_1 - n_1 + 1)\lambda_1 \Delta t\, p(t; n_1 - 1, s, n_2) \\ & + (S - s + 1)\lambda_1 \Delta t\, p(t; n_1, s - 1, n_2) \\ & + (N_2 - n_2 + 1)\lambda_2 \Delta t\, p(t; n_1, s, n_2 - 1) \\ & + \mu_s \Delta t\, p(t; n_1 + 1, s - 1, n_2) \\ & + \chi(n_1 + 1, s, n_2)\mu \Delta t\, p(t; n_1 + 1, s, n_2)^* \\ & + \chi(n_1, s + 1, n_2)\mu \Delta t\, p(t; n_1, s + 1, n_2)^\dagger \\ & + [1 - \chi(n_1, s, n_2 + 1)]\mu \Delta t\, p(t; n_1, s, n_2 + 1). \quad (3) \end{aligned}$$

* Omit when $s < S$.

† Omit when $s + 1 = S$, $n_1 > 0$.

If there are no spares and $\lambda_1 = \lambda_2$, then there is nothing different about group 1 systems compared with group 2 systems. The model is equivalent to one large group of $(N_1 + N_2)$ minicomputer systems. While the time dependence makes a solution nontrivial, this is a standard, finite source, M/M/1 queue with a well-known steady-state solution.[1] This feature was used to check the model, as discussed later.

Case 3: Infinite repairers—With an infinite number of repairers no criteria are required to determine which system gets fixed next. But, in the spirit of a two-repairer approximation, we impose the following condition on the repair process: Never fix group 1 systems if spares are available.

Figure 4 shows the state diagram for a general state (n_1, s, n_2) with infinite repairers. The corresponding state equation describing the probability at time $t + \Delta t$ of being in state (n_1, s, n_2) is, leaving out terms $o(\Delta t)$,

$$
\begin{aligned}
p(t + \Delta t; n_1, s, n_2) = [1 &- (N_1 - n_1)\lambda_1\Delta t - (S - s)\lambda_1\Delta t \\
&- (N_2 - n_2)\lambda_2\Delta t \\
&- (n_1{}^* + s + n_2)\mu\Delta t - \mu_s\Delta t]p(t; n_1, s, n_2) \\
&+ (N_1 - n_1 + 1)\lambda_1\Delta t\, p(t; n_1 - 1, s, n_2) \\
&+ (S - s + 1)\lambda_1\Delta t\, p(t; n_1, s - 1, n_2) \\
&+ (N_2 - n_2 + 1)\lambda_2\Delta t\, p(t; n_1, s, n_2 - 1) \\
&+ \mu_s\Delta t\, p(t; n_1 + 1, s - 1, n_2) \\
&+ (n_1 + 1)\mu\Delta t\, p(t; n_1 + 1, s, n_2)^* \\
&+ (s + 1)\mu\Delta t\, p(t; n_1, s + 1, n_2) \\
&+ (n_2 + 1)\mu\Delta t\, p(t; n_1, s, n_2 + 1). \qquad (4)
\end{aligned}
$$

The group 2 part of eq. (4) is separable from the parts for group 1 and the spares, as is the case for no repairers [see eq. (2)]. Again, for programming ease this separation was not made.

3.3 Numerical methods and checks of the model

With the three state equations in hand we can solve for minicomputer system availabilities under different sparing and maintenance strategies.

The standard procedure to arrive at analytic solutions for the state

* Omit when $s < S$.

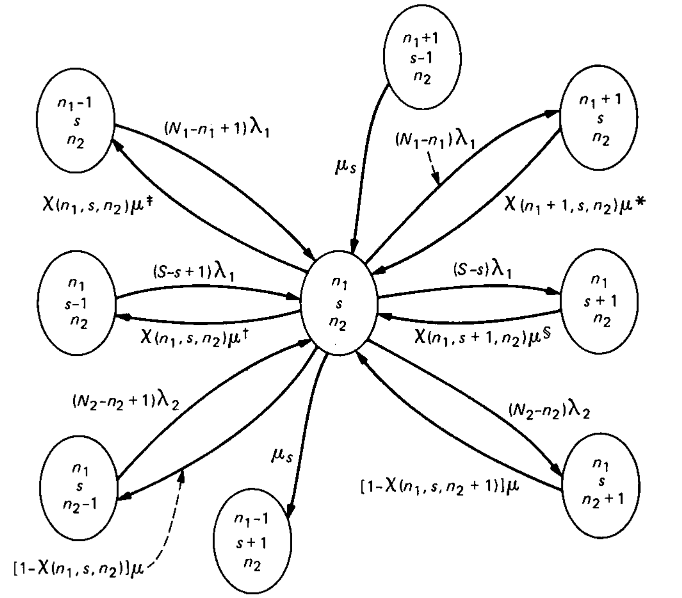

*OMIT WHEN $s < S$.

† OMIT WHEN $s + 1 = S$, $n_1 > 0$.

‡ OMIT WHEN $s < S$; OR $s = S$, $n_1 = 0$.

§ OMIT WHEN $s = S$, $n_1 > 0$.

Fig. 3—State diagram for the general state (n_1, s, n_2) with one repairer.

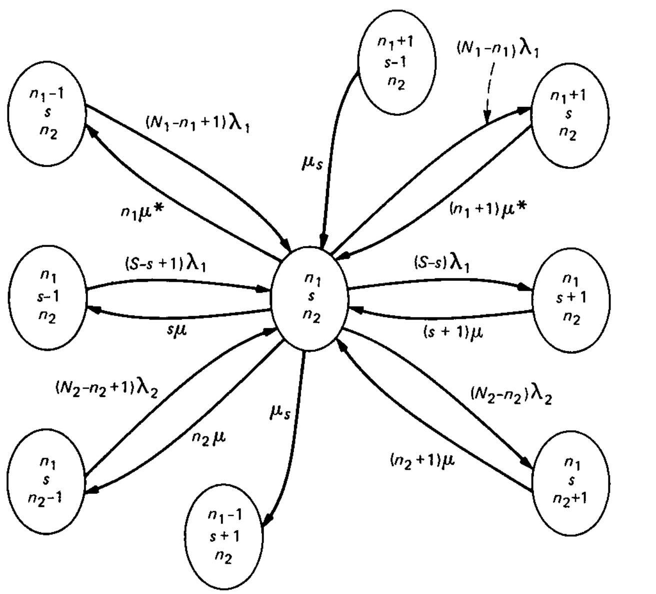

*OMIT WHEN $s < S$.

Fig. 4—State diagram for the general state (n_1, s, n_2) with infinite repairers.

equations is to take the limit as $\Delta t \rightarrow 0$ to obtain differential equations for $p(t; n_1, s, n_2)$. But analytic solutions of the differential equations are intractable.

An alternative is to take the numerical approach and integrate the difference equations (1), (3), and (4). A computer program was written to perform the numerical integration. The computer program takes the probability distribution at some initial time, T_i, and integrates in $(T_e - T_i)/\Delta t$ steps to the desired end time, T_e. The initial distribution $p(T_i; n_1, s, n_2)$ was chosen such that

$$p(T_i; 0, 0, 0) = 1 \tag{5}$$

and for all other values of n_1, s, and n_2

$$p(T_i; n_1, s, n_2) = 0. \tag{6}$$

To avoid nonlinear effects Δt must be small. Even so, numerical roundoff will eventually intrude with the result that

$$\sum_{n_1,s,n_2} p(t; n_1, s, n_2) \neq 1. \tag{7}$$

To compensate for this, the probability distribution was normalized periodically, dividing each probability by the sum of the probabilities, to ensure that the sum remains one.

To escape the influence of the initial conditions, the model must be run for a time period whose length depends on the conditions being modeled. For a constant number of repairers the model must run until the probability distribution becomes constant. For a time-dependent weekly maintenance schedule, the model must run until the probability distribution becomes periodic, repeating from week to week.

Availabilities are calculated from the probability distribution at each hour. The probability that n_1 systems are failed in group 1 at time t is

$$p(t; n_1) = \sum_{s=0}^{S} \sum_{n_2=0}^{N_2} p(t; n_1, s, n_2). \tag{8}$$

The time-dependent, average number of failed group 1 systems is then

$$\langle n_1(t) \rangle = \sum_{n_1=0}^{N_1} n_1 p(t; n_1). \tag{9}$$

From this the average availability of a system in group 1 is

$$n_1 \text{ avail}(t) = 1 - \frac{\langle n_1(t) \rangle}{N_1}. \tag{10}$$

Similar expressions hold for spares and group 2 systems. These expressions produced the data for the figures in this paper.

In addition to availabilities, two parameters were calculated: repairer occupancy (the probability that a repairer is busy) and the probability that three or more systems are failed. The first served as a check of the model. One expects a single repairer's work load to increase by about 4 percent (1 spare/25 systems) when the spare is added. One also expects the work load to be independent of the distribution of the 25 systems between group 1 and group 2. Calculated repairer occupancies agreed with expected values with greater than 0.1 percent accuracy.

The second parameter served to verify the infinite repairer approximation of two repairers. For all cases examined with the maintenance schedule used in the figures, the probability that three or more repairers were busy was less than 0.065 when the number of repairers was infinite. That is, for at most 6.5 percent of the time more than two of the infinite repairers are busy.

Four other checks were made to verify the calculations in the model; three of these were mentioned earlier. First, for no repairers the time dependence for the group 2 systems can be checked explicitly. Using the initial condition that all systems are working at the starting time, T_i, $p(t; n_2)$ from eq. (2) is given by

$$p(t;\, n_2) = \binom{N_2}{n_2} e^{-N_2\lambda_2 t}[e^{\lambda_2 t} - 1]^{n_2}, \tag{11}$$

which gives

$$\langle n_2(t)\rangle = N_2(1 - e^{-\lambda_2 t}). \tag{12}$$

The numeric solution agreed with this relation with greater than 0.1 percent accuracy. It is one of the best checks of the choice for Δt (0.005 hour) and the decision to normalize after every Δt interval, because it checks the time dependence as well as the steady state.

Second, for one repairer the random repair criterion was checked for the case with no spares and $\lambda_1 = \lambda_2$. As we mentioned earlier, in this case the model is equivalent to one large group of $N_1 + N_2$ minicomputers and has a well-known steady-state solution:[1]

$$\langle n\rangle = \frac{\sum_{n=0}^{N} n \dfrac{N!}{(N-n)!}\left(\dfrac{\lambda}{\mu}\right)^n}{\sum_{n=0}^{N} \dfrac{N!}{(N-n)!}\left(\dfrac{\lambda}{\mu}\right)^n}. \tag{13}$$

If we use $N = N_1 + N_2$, the average number of failed systems should be the sum of the average numbers for group 1 and group 2 in the model:

$$\langle n \rangle = \langle n_1 \rangle + \langle n_2 \rangle. \tag{14}$$

In addition, the average availabilities for systems in group 1 and group 2 should be equal and the same as that of the large group. Agreement with this relation not only checked the numerical calculation, it also verified that the random repair criterion is equivalent to first-in first-out, as far as average availabilities are concerned.

Third, for infinite repairers the steady-state solution was checked for the group 2 systems. Equation (4) separates, as we said, allowing an explicit solution for $p(\infty; n_2)$, which agreed with the computer calculations with greater than 0.1 percent accuracy. The steady-state solution gives

$$\langle n_2 \rangle = \frac{\sum_{n_2=0}^{N_2} n_2 \binom{N_2}{n_2} \left(\frac{\lambda_2}{\mu}\right)^{n_2}}{\sum_{n_2=0}^{N_2} \binom{N_2}{n_2} \left(\frac{\lambda_2}{\mu}\right)^{n_2}}. \tag{15}$$

Finally, special conditions exist at the boundaries of the state space. These conditions result, for example, when switching terms drop out of the state equations for $s = S$ or $n_1 = 0$, or when the repair term drops out for $n_1 = s = n_2 = 0$. For one repairer the case with the most complex boundary conditions, an explicit, steady-state solution, was derived for $N_1 = S = N_2 = 1$. This case employs all possible boundary conditions, including those for the decision parameter, χ. The steady-state computer calculation agreed with the explicit solution with greater than 0.1 percent accuracy.

IV. PLANNING FOR SYSTEM SPARING

To provide tools for studies of system sparing, system availability data from the model have been arranged to answer three questions about sparing:

1. What is the minimum number of identical systems needed for proving-in the cost of sparing?
2. Given a minimum system availability objective for a group of identical systems, what is the minimum number of spares required to meet the objective?
3. What are the operational benefits of sparing for the MMOC?

4.1 Proving-in spares

Figure 5 shows the annual recovered downtime, using one spare for the systems in group 1, as a function of the number of systems in group 1, for annual downtimes of 50, 100, and 150 hours per system without a spare, respectively. For each system,

$$\text{Downtime} = (1 - \text{availability}) \times 24 \text{ hrs/day} \times 365 \text{ days/yr.} \quad (16)$$

The solid curves in the figures show the total downtime recovered for all group 1 systems. Nonspared systems, those in group 2, also benefit from sparing, because a repairer can give them higher-priority service when a spare can be switched for a failed group 1 system. The dashed curves in the figures include the additional downtime recovered for group 2 systems.

These curves provide essential data for an economic analysis for proving-in sparing.

4.2 Meeting an availability objective

Figure 6 shows the availability per system for group 1 systems, as a function of the number of group 1 systems. The three curves correspond to annual downtimes of 50, 100, and 150 hours per system without a spare. The data assumes 1 repairer, 24 hours per day, and 1 spare.

These data can be used to determine the maximum number of systems that can be loaded onto one spare and still meet a specified availability objective. For a given availability objective on the vertical axis in the figure, read over to the curve corresponding to the average

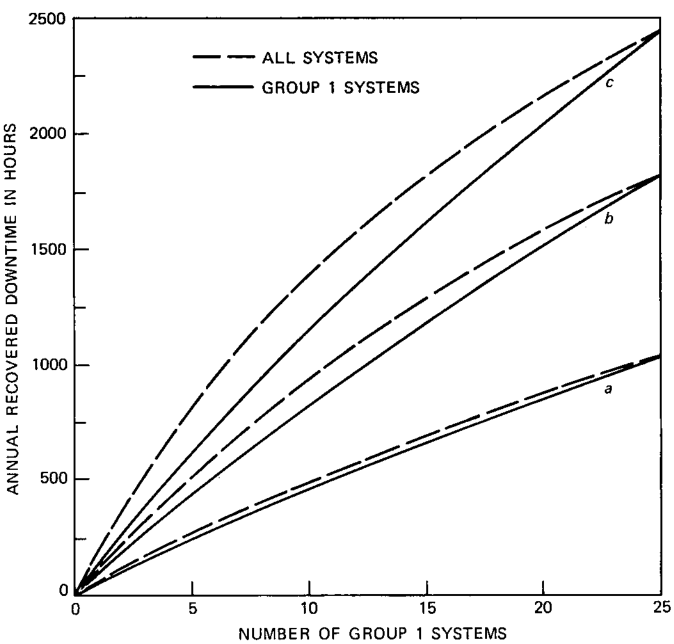

Fig. 5—Annual downtime recovered for group 1 systems and for all 25 systems, group 1 plus group 2, with one repairer, 24 hours per day, and one spare. Annual downtime is (a) 50, (b) 100, and (c) 150 hours per system without a spare.

system downtime without a spare. The maximum number of systems for which the availability objective can be met with one spare is then read on the horizontal axis.

4.3 Improving maintenance staff utilization

The solid curves in Figs. 7 and 8 show, for groups 1 and 2, respectively, system availabilities with sparing, when night shift maintenance personnel are moved to the day shift and weekend maintenance is done on a call-out basis. Such a maintenance schedule could reduce maintenance staff attrition (people do not like to work at night) and would decrease the maintenance staff size by eliminating full-time weekend coverage. These availability data are compared with system availabilities for 1 repairer and 24-hour, 7-day coverage, both with a spare (short-dashed lines in the figures) and without a spare (long-dashed lines). This is one of two potential strategies to improve maintenance staff utilization with sparing. The second is to increase the number of systems per repairer when spares are deployed. The first is recommended for clusters deploying spares. The second can be shown to be inadvisable since the value of lost system availability due to maintenance staff reduction outweighs the savings in staff salaries.

4.4 Interpolations, extrapolations, and sensitivities

Planning studies for system sparing will often require availability data for parameters that differ from those described in Section II. The differences will probably occur in three areas. First, annual system downtimes without sparing will generally not be 50, 100, or 150 hours, but instead will be somewhere within this range. Linear interpolation

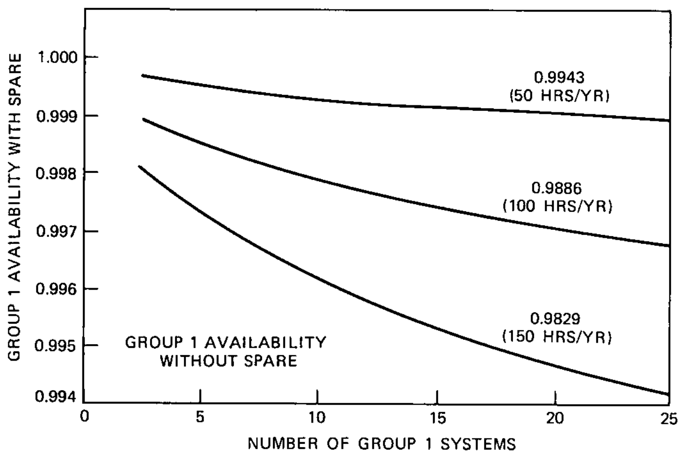

Fig. 6—System availability for minicomputers with sparing (group 1 in the model) as a function of the number of group 1 systems. Curves correspond to annual downtimes in hours per system per year without sparing.

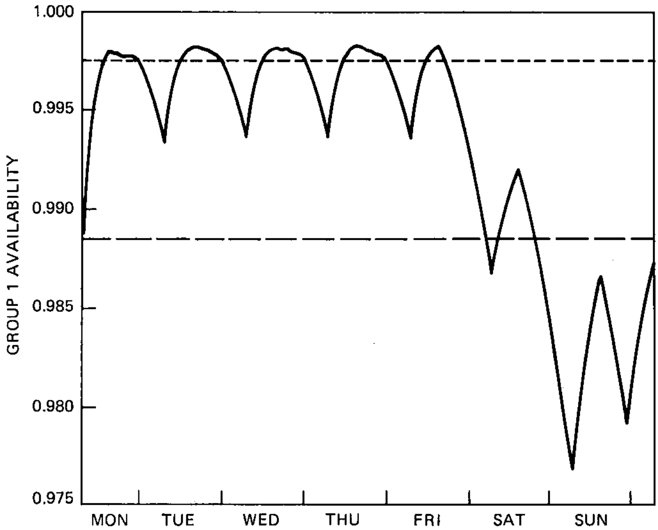

	REPAIRERS PER SHIFT						
	MON	TUE	WED	THU	FRI	SAT	SUN
DAY	2	2	2	2	2	1	1
EVENING	1	1	1	1	0	0	0
NIGHT	0	0	0	0	0	0	1

Fig. 7—Effect of one spare on system availability of 15 group 1 systems is compared for two repairers during the day, one in the evening, and weekend repair done on a call-out basis (solid curve); and for 24-hour, 7-day, single-repairer maintenance (long-dashed line). System availability without sparing is shown for 24-hour, 7-day, single-repairer maintenance coverage (short-dashed line), corresponding to 100 hours of annual down-time per system.

from the data in Figs. 5 and 6 will provide the desired system availabilities.

Second, the number of systems per repairer (group 1 plus group 2) will vary. Changes in this number have less effect on group 1 systems than on group 2 systems. For group 1 systems, as the number of systems per repairer decreases from 25 to 15, recovered downtime increases by 5.5 percent, and as the number increases from 25 to 35, recovered downtime decreases by 3.5 percent. Thus, over a broad range of systems per repairer (15 to 35), the group 1 availability data in Figs. 5 and 6 are accurate to within 5.5 percent. For group 2 systems the effect is more pronounced. For a constant number of group 1 systems, the group 2 contribution to recovered downtime (the region between the solid and dashed curves in Fig. 5) is roughly linearly proportional to the number of group 2 systems. For example, when the number of group 2 systems is one third the value used in the figures, the group 2

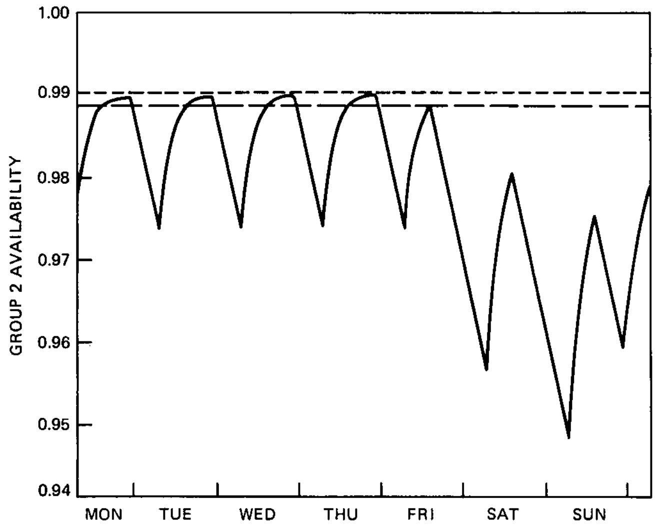

REPAIRERS PER SHIFT	MON	TUE	WED	THU	FRI	SAT	SUN
DAY	2	2	2	2	2	1	1
EVENING	1	1	1	1	0	0	0
NIGHT	0	0	0	0	0	0	1

Fig. 8—Effect of one spare on system availability of 10 group 2 systems is compared for two repairers during the day, one in the evening, and weekend repair done on call-out basis (solid curve); and for 24-hour, 7-day, single-repairer coverage (long-dashed line). System availability without sparing is shown for 24-hour, 7-day, single-repairer maintenance (short-dashed line), corresponding to 100 hours of annual downtime per system.

contribution to recovered downtime is roughly one third the value shown.

Finally, switching time will not be the same for every cluster. The availability data are insensitive to this parameter. Changing the switching time from 15 minutes to 30 minutes produces a negligible change in recovered availability.

V. OBSERVATIONS

5.1 Networking

Increasing dependence on minicomputer-based OSs has led the BOCs and AT&T Communications to deploy system spares for OSs clustered in MMOCs. As networking of systems continues, the impact of system downtime, and thus the value of sparing, will grow. With networking, an out-of-service minicomputer system not only fails to perform its assigned task, it also fails to provide essential data to other systems on the network.

5.2 Commonality

To be economical, system sparing requires a sufficient number of clustered OSs with identical hardware. Studies of system sparing may cause BOCs and AT&T Communications to relocate some systems to clusters with other like systems, or to purchase new systems with older model peripherals to achieve the commonality needed to spare a group of OSs. In anticipation of increasing reliance on high system availability, it is clearly desirable that future OS developments use common hardware.

REFERENCES

1. L. Kleinrock, *Queueing Systems*, New York: Wiley, 1975, Vol. 1: Theory, p. 106.

AUTHOR

David W. Tolleth, B.A. (Physics), 1974, University of California, Irvine; M.S., Ph.D. (Physics), University of California, San Diego, 1976 and 1981, respectively; AT&T Bell Laboratories, 1981—. At AT&T Bell Laboratories, Mr. Tolleth has been planning for effective operation and maintenance of minicomputer-based Operations Systems for the Bell Operating Companies and AT&T Communications.