Dynamic Spares Provisioning for the DSN

I. Eisenberger and G. Lorden
Communications Systems Research Section

A method is proposed to maintain cost-effective spares stockpiles for an entire DSN station or complex as new modules or subsystems are continuously added. Levels of spares are calculated individually for each new module type so as to conform to an established standard for the entire spares pool. Finding the spares levels is computationally simplified and in many cases is reduced to consulting a Down-Time Ratio Chart. A simple modification permits taking into account a criticality factor.

I. The Spares Provisioning Algorithm

Cost-effective sparing algorithms have been proposed by the authors in Refs. 1 through 4. These were designed to provide optimal spares packages for systems or subsystems by taking into account for each module type the mean-time between failures (MTBF), unit cost, number operating, and number redundant. Each spares package is optimal in the sense that it minimizes the cost of attaining its up-time ratio (UTR). The choice of a suitable package among optimal ones is made by trading off cost against UTR.

The benefits of using such an algorithm flow from the fact that every dollar spent on spares buys the right choice of module type to maximize the UTR. A difficulty arises, however, when new subsystems or parts of subsystems are introduced and the algorithm is used to determine the spares needed. A spares package with cost C and up-time ratio U might be chosen this year and another "optimal" package with cost C' and UTR U' may be chosen next year for new equipment. The combination of the two will generally be suboptimal and could be improved upon by considering the two procurements together, thus spending the right relative amounts on the two to maximize the *overall* UTR. Since it is obviously impractical to make all new procurements at the same time or to recompute the entire spares stockpile of a station every time a new type is added, a method of dynamic spares provisioning is needed to insure that overall cost-effectiveness is maintained.

To outline such a method it is convenient to work with the down-time-ratio ($DTR = 1 - UTR$) of a system, the fraction of the time that the system is not operable due to unavailability of spares. Let $A > 0$ be chosen as the "sparing" standard" (discussed below). Then when a new module type with unit cost C_i is added to the spares pool, the number of spares it requires is the *smallest* number such that

$$DTR_i < A \cdot C_i, \tag{1}$$

where DTR_i denotes the down-time ratio of the ith module type. One computes DTR_i for 0 spares, 1 spare, 2 spares, etc., until a value smaller than $A \cdot C_i$ is obtained. (The computation of up-time, hence down-time, ratios is easily made by the algorithm in Ref. 1). Thus, $A \cdot C_i$ is a "ceiling" for the down-time ratio of module type i.

The advantage of the method using relation (1) is that it is applied to each module type without regard for which others are determined at the same time, the resulting spares being the same as if the entire spares complement of a station (or complex) were recomputed each time a new module type is added.

The disadvantages of the method are:

(1) It is only approximately optimal.

(2) The value of A must be chosen in advance (as described in the next section).

The nonoptimality results from the intentionally simplified character of relation (1), not from considering each module type separately. Investigation of actual procurement examples previously computed using the optimal package approach indicates that the sacrifice of strict optimality has negligible effect: the method based on relation (1) often yields the optimal packages and, even when it does not, incurs minimal degradation in cost-effectiveness.

II. Choice of a Sparing Standard

The role of the sparing standard A in determining the spares complement is basically one of establishing the level of the tradeoff between total spares cost and overall UTR. A rough idea of the meaning of the A-value can be obtained from the following upper bounds on the overall DTR:

$$DTR = 1 - UTR = 1 - \prod_i UTR_i = 1 - \prod_i (1 - DTR_i)$$

$$< \sum_i DTR_i < A \sum_i C_i,$$

using relation (1) for the last inequality.

Thus,

$$A > \frac{DTR}{\sum_i C_i}, \qquad (2)$$

where $\sum_i C_i$ is the cost of a "unit spares package" (one spare of each module type). In typical cases, the two sides of inequality (2) are within a factor of 2 or so; thus, A is roughly the ratio of the overall DTR to the cost of a unit spares package. (Note that the use of this rough interpretation is conservative; i.e., actual DTR's are smaller than the interpretation suggests.)

Although the rough interpretation of the sparing standard A is helpful, it is not accurate enough to be a completely satisfactory basis for choosing the value of A. That choice should be based upon a fully informed tradeoff between cost and up-time ratio for an entire spares pool. For this purpose one has to compute spares for all modules using a range of values of A and then compare the resulting costs and UTR's. This computation can be performed piecemeal, taking modules singly or in convenient groups. Once A is chosen, the process need not be repeated for some time.

Periodically it may be necessary or desirable to set a new value for the sparing standard in order to adapt to changes in performance requirements or cost constraints. Also, if the amount of equipment increases significantly, then the overall UTR will go down and it may be decided that more spares should be procured using a smaller value of A.

Without changing A, one can recalculate the sparing needs of individual modules from time to time based upon better estimates of their MTBF's. If, for example, a module's failure history indicates a lower MTBF than specified by the manufacturer, then an appropriate number of additional spares can be determined by recalculating DTR's using the more realistic MTBF.

III. Calculation of Down-Time Ratios

The actual determination of spares levels using the method of Section I reduces to the calculation of down-time ratios for given numbers of spares. The algorithm in Ref. 3 calculates DTR's with or without redundancies among operating modules. In the case of multistation spares pools, the UTR's for the various stations will generally *not* be equal and the geometric mean of the station UTR's is the most convenient measure of overall performance. Thus, one defines the "average DTR" by

$$DTR = 1 - [UTR(1) \times \ldots \times UTR(s)]^{1/s}$$

where s = number of stations and $UTR(j)$ = up-time ratio at station j. A similar definition is given for DTR_i, the average DTR of module type i. It is easily verified that the method of sparing based on relation (1) yields an overall DTR satisfying the upper bound in inequality (2).

The use of pooled sparing or redundancy, while highly cost-effective, does not lead to calculations that can be summarized in graphs or tables. In the simplest sparing situation, however, where a module type is operating without redundancy and without pooling, a simple computation easily expressed in graphs, or tables, can be used to determine the down-time ratios and, hence, the number of spares required for a prescribed A value.

In terms of the *operating ratio*

$$V = \text{no. of modules operating} \times \frac{MTTR}{MTBF} \text{(in same units)}$$

The *reciprocal* DTR is expressible as follows:

No. spares	$\frac{1}{DTR}$
0	$1 + \frac{1}{V}$
1	$\left(1 + \frac{1}{V}\right)\frac{2}{V} + 1$
2	$\left[\left(1 + \frac{1}{V}\right)\frac{2}{V} + 1\right]\frac{3}{V} + 1$
:	:
.	.
n	$\sum_{k=0}^{n+1} \frac{n+1}{(n+1-k)} v^{-k} = \left\{\left[\left(1 + \frac{1}{v}\right)\frac{2}{v} + 1\right]\frac{3}{v} + \ldots\right\}\frac{n+1}{v} + 1$

It is easy to carry out these computations on a pocket calculator. At each stage simply multiply the last answer by $n + 1$, divide by V, and add 1. When the result exceeds $1/(A \cdot C_i)$, the number of spares required according to relation (1) is n.

Alternatively, one can use a Down-Time Ratio Chart (like the one in Fig. 1) that plots $1/DTR$ as a function of

$$V^* = \text{no. operating} \times \frac{MTTR \text{ (weeks)}}{MTBF \text{ (k-hr)}} = 5.952\ V$$

rather than V for greater convenience of computation. Locate the point with coordinates V^* on the horizontal axis and $1/(A \cdot C_i)$ on the vertical axis. The number of spares required is determined by the region in which that point falls. For example, the point with $V^* = 0.4$ and $1/(A \cdot C_i) = 3000$, shown as P in Fig. 1, lies in the region above the $1/DTR$ curve for 1 spare and below the curve for 2 spares. Hence, the goal for $1/DTR$, which is 3000, is not reached by using 1 spare but is exceeded by using 2 spares. Thus, the region in which P falls is labeled "2 spares" and the other regions between curves are labeled similarly.

IV. An Example

Table 1 lists the results of computations for a pooled sparing example with no redundancy. The data on module types were obtained by selection from a spares calculation on actual DSN equipment. Note in Table 2 that the configurations of the stations differ and the DTR's differ correspondingly. The value of $A = 6.42 \times 10^{-7}$ was chosen based on inequality (2) with a DTR goal of 1%. The actual average DTR for the three stations came out 0.48%, which is roughly a factor of 2 smaller (as is typical; see the discussion in Section II). The spares package obtained in this example is actually an optimal one as was verified by using the algorithm of Ref. 3. A repair time of two weeks was assumed for all modules.

Note that module type 5, requiring the most spares, has a unit cost in the middle of the range and an MTBF roughly as good as three of the other modules. However, the number operating is high for a module with this MTBF level, so that the number operating/MTBF (= total number of failures per 1000 hours of operation) is relatively high. Type 4, which requires the second highest number of spares, is distinguished by very low cost. The relatively high spares level determined for type 4 produces an extremely low DTR (0.003%), but these spares are still "needed" because of their cost-effectiveness in maximizing system UTR. Type 2, by contrast, is quite expensive and is spared at a level which gives it a relatively high DTR.

V. Criticality Factors

A useful sparing algorithm should be capable of taking into account differences in criticality of different module types. Down-time due to shortage of spares may be much less critical for some modules than for the most essential modules.

To represent such differences mathematically, let each module type be assigned a *criticality factor* K_i, where $0 < K_i < 1$. The interpretation is that down-time of module type i is assigned a weighting factor K_i on a scale from 0 to 1, with higher K_i's thus denoting higher criticality. The modified algorithm replaces relation (1) by

$$K_i \cdot DTR_i < A \cdot C_i \tag{3}$$

This modification essentially minimizes $\sum_i K_i \cdot DTR_i$ rather than $\sum_i DTR_i$ for the dollars spent.

Another interpretation of relation (3) is that one follows the algorithm in relation (1) but with C_i, the unit cost, replaced by C_i/K_i, the *cost-criticality ratio*. Modules of low criticality wind up being treated as if their cost were much higher, whereas the costs of modules of maximum criticality, 1, are unchanged.

This method can greatly enhance the flexibility of the sparing algorithm in dealing with diverse types of equipment. One way of implementing it would be to classify all module types into, say, two or three categories with appropriately specified criticalities.

VI. Conclusions

The method of sparing proposed in this article offers both simplicity and cost-effectiveness. It is easy to understand and flexible in applications, permitting consideration of criticality as well as redundancy and pooled sparing. It also enables sparing determinations to be made for new modules without reconsidering the entire spares stockpile. With suitable software implementation, this method should prove highly useful in DSN spares provisioning.

References

1. Eisenberger, I., Lorden, G., and Maiocco, F., "A Preliminary Study of Spares Provisioning for the Deep Space Network," *DSN Progress Report*, TR 32-1526, Vol. XVIII, Jet Propulsion Laboratory, Pasadena, Calif., Dec. 15, 1973.
2. Eisenberger, I., Lorden, G., and Maiocco, F., "Cost-Effective Spares Provisioning for The Deep Space Network," *DSN Progress Report 42-20*, Jet Propulsion Laboratory, Pasadena, Calif., Apr. 15, 1974.
3. Eisenberger, I., and Lorden, G., "Cost-Effectiveness of Pooled Spares in the Deep Space Network," *DSN Progress Report 42-38*, Jet Propulsion Laboratory, Pasadena, Calif., Aug. 15, 1975.
4. Eisenberger, I., and Lorden, G., "Priority Repair Schemes in the Deep Space Network", *DSN Progress Report 42-30*, Jet Propulsion Laboratory, Pasadena, Calif., Dec. 15, 1975.

Table 1. Data for sparing example

No.	No. operating[a]	MTBF, k-hr	No. op. MTBF	Unit cost, $	No. spares	Avg.[a] DTR_i, %
1	21	13	1.62	4000	3	.0834
2	51	50	1.02	5200	2	.1830
3	21	15	1.40	2150	3	.0494
4	63	49	1.29	236	4	.0030
5	42	10	4.20	2000	5	.1254
6	6	10	.60	2000	2	.0397
Total				15586	19	.4833

[a]See Table 2 for individual stations.

Table 2. Breakdown by station

Module No.	No. operating Station 1	Station 2	Station 3	DTR_i, % Station 1	Station 2	Station 3
1	8	9	4	.0952	.1065	.0486
2	17	14	20	.1831	.1516	.2144
3	9	8	4	.0632	.0564	.0287
4	25	25	13	.0036	.0036	.0019
5	8	16	18	.0744	.1428	.1590
6	4	2	0	.0789	.0401	0
Overall	71	74	59	.4976	.5002	.4520

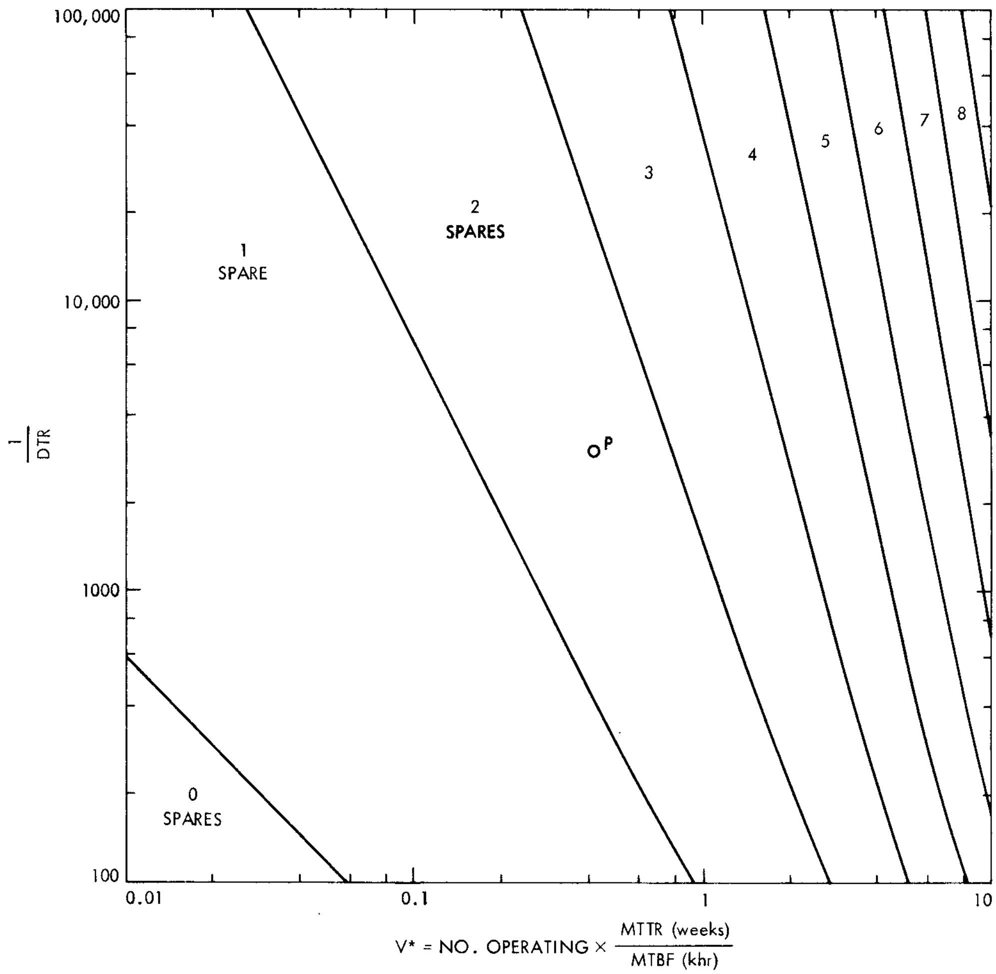

Fig. 1. Down-Time Ratio Chart