This document was prepared in conjunction with work accomplished under Contract No. DE-AC09-76SR00001 with the U.S. Department of Energy.

DISCLAIMER

TECHNICAL DIVISION
SAVANNAH RIVER LABORATORY

DPST-83-792

ACC. NO. 106041

DISTRIBUTION

D. A. WARD, 773-A, A-231
H. F. RING, M-6600, WILM.
G. F. MERZ, 703-A, A-239, SRP
H. F. ALLEN, 706-C, 4-A
W. B. DASPIT, 706-C, 2-A
J. L. JONES, 706-C, 8-E
J. JIUNNIES, 703-A, A-123
M. N. NELSON, 777-10-A
C. W. SWINDELL, 777-10A
C. G. HARDIN, 777-10A
M. E. BACINOGLU, 777-10A
M. R. BUCKNER, 773-A, A-230, SRL
D. B. RANKIN, 773-22A
A. A. TUDOR, 773-22A
J. P. CHURCH, 773-22A
R. L. POSTLES, 773-16A
J. A. SMITH, 773-22A
SRL RECORDS (4)

August 22, 1983

TO: D. A. WARD

FROM: J. A. SMITH JAS

LOSS OF 115 kV POWER

INTRODUCTION

Continuous electric power to SRP reactors is necessary to maintain water flow for heat removal, and essential monitoring and control. Should power supplied to the plant 115 kV system from offsite be lost, on-site generation is sufficient to maintain all reactors in a safe shutdown mode for an indefinite period. Should on-site generators for the 115 kV grid also be lost, diesel-electric generators within each reactor building are also sufficient to maintain safe shutdown but only for a finite period. Primary cooling flow would be provided by dedicated, online, diesel-electric DC pump drives. Secondary flow would be assured for a few hours by gravity flow from a storage basin. Startup of emergency diesel-electric AC generators would permit operation of basin recirculation pumps and extend the finite period for secondary flow to a few days.

In all cases DC power for necessary monitoring and control would be available from battery systems with AC converter backup.

Loss of all power to the 115 kV system could be an initiating event leading to melting a reactor core (for lack of secondary cooling water). Such an accident might begin with loss of all commercial power, which then triggers loss of all on-site generators because of overloading. Some means of reducing the likelihood of losing on-site generators have been proposed recently[1,2]. To aid in assessing some of these proposals, the postulated accident is reviewed here.

SUMMARY

Data in hand provide means of deriving the probability of losing onsite generators (OSG) under overload conditions. But perception of several relevant variables prompts offering three different estimates of the "expected" frequency of core melting due to loss of 115 kV power.

(a) Based on the recent history of SRP power reliability (past 10 years, or so) the probability of core melting due to loss of 115 kV power has been estimated to be on the order of 10^{-7} per year.

(b) More recently, there is evidence that the OSG have entered a "wearout" phase, characterized by an accelerating probability of failure under overload conditions. An assumption that this condition applies to the immediate future (prior to operation of automatic load-shedding) increases the estimated frequency of core melting to the order of 10^{-4} per year.

(c) The automatic load-shedding system now being installed should reduce the probability of core melting back below the estimate of 10^{-7} per year.

These estimates of the probability that a loss of offsite power might lead to a core melting accident should be considered with careful regard to the assumptions made. First, the estimates of component and system failure rates or failure probabilities are not the product of a comprehensive analysis. They are the best estimates that can be made at the present time with existing data and resources. They are judged to be reasonable. Second, the estimated rates are based upon extrapolations of experience. They do not include the probability of initiating events which could result in common failures of several safety systems, and which can be postulated, but for which there is no experience base from which to estimate probabilities. For example, a very large earthquake, well beyond the design basis earthquake for the reactor, might render inoperative all or several of the power systems. The frequency of occurrence of such an earthquake is not known —— it

might be truly zero; it is certainly less than once in 10,000 years. However, when the results of probability calculations yield values as low as 10^{-7} per year (as above), it is appropriate to recognize that there may very well be exceedingly rare events which, if they could be quantified, would be seen as contributing more to risk than the experience-based events whose risk contributions have been quantified. The important conclusion is that an event so rare as to occur only once in 10^7 years, as in cases (a) and (c) above, should be regarded as having effectively, zero probability. There is no incentive to further reduce its probability or its consequence. An event having a probability of once in 10^4 years, as case (b) above, might be considered as a significant contributer to risk if the consequence of the event is known or judged to be very large. Thus, there is incentive to reduce its probability or consequence. In this case, the program to provide automatic load shedding is appropriate.

DISCUSSION

A. Nominal Accident Sequences

The pre-accident situation is assumed to include three reactors operating at full power. The plant electrical load is assumed to be 160 MW (16 times the average capacity of an OSG). We assume only 6 OSG on line, although it is more nearly typical to have 7 or 8 OSG on line. Commercial power (CP) supplies the balance. We assume the pre-1983 procedures for load shedding to be in effect. Thus, the analysis is aimed at the recent past rather than the future. And it is intended to provide a conservative estimate of chances of encountering difficulties.

The postulated accident begins with loss of all CP. The succeeding events are illustrated in Figure 1 and discussed briefly below. More detail is given in Section B. We have not considered every imaginable sequence of events. Specifically, we assume early shutdown of the reactors and that the DC motors circulate the primary coolant as required.

If CP were lost, we would rely on OSG to supply electrical power for essential functions. The most pivotal of these functions is pumping cooling water from the river (or PAR pond) to the reactor areas.

Loss of CP overloads the OSG by about 170%. This degrades the power output (lower voltage and frequency) and subjects machinery to stresses capable of causing damage. If one OSG were rendered inoperable by overloading, this would increase the degree of overload on the remaining OSG. We therefore entertain the possibility of a cascading effect that eliminates all OSG. This would leave no source of power for pumping water to the reactor areas.

If load shedding is called for, abnormal condition control procedures (DPSOL 105-MC-15) prescribe recirculating secondary cooling water to preserve the supply on hand. This requires a supply of 480V power to realign valves and to operate a circulating pump. But the normal supply of 480V power was lost with the OSG. So the emergency power supply (diesel generators) is called upon to enable recirculation.

If emergency power were not available, the reactor could still be cooled as long as the supply of secondary coolant lasts. But the supply could be exhausted within 8 hours if no steps were taken to reduce flow in the secondary system. In this event, core melting could result unless power to river water pumps were restored within 8 hours (to provide additional secondary cooling water).

If emergency power were available, recirculation of secondary coolant would begin. Recirculation must be continued until one of two "saving" events occurs: restoration of power to river water pumps or discharge of the fuel from the reactor.

We assume that it would require a minimum of 4 days to make preparations for fuel discharge, plus one more day to complete discharge. If recirculation should fail any sooner than 4 days, discharge would not be possible. The only way to avoid melting in this case would be to restore pumping power.

Explicit calculations(3) indicate that 5 days of recirculation is about the best we can hope to accomplish because some of the "auxiliary" applications of secondary cooling water cannot be recirculated. (About 2000 gpm is spent on cooling diesel engines, electric motors, pumps and compressors.) Improvements have been recommended(7) but are assumed not available here. Thus, either discharge must be completed or power must be restored before the assumed 5 day deadline passes. And, in the meantime, recirculation must be maintained.

There are other possible failures that could prevent recirculation of cooling water for the full 5 days (human error, valve or pump failure, failure of emergency power). Chances of recovering power to river water pumps during this interval should be estimated by considering the full continuum of possible times between 8 hours and 5 days. For purposes of this study, we make the conservative assumption that one day is a "representative" time for estimating the mean time over the continuum. This assumption is indicated in Figure 1.

B. Individual Events

1. Loss of Commercial Power

Total loss of commercial power (CP) has occurred 4 times in the 29 year history of SRP. One of these events may be dismissed as avoidable in the future (Reference 4, Section C.1.a), indicating a frequency of about 10^{-1} per year.

2. Loss of On-Site Generators

a. Experience with Failure in Service

If we postulate that the probability of failure per unit time for a given OSG unit is a constant, we can deduce values for that constant from experience. Detailed arguments are presented in Appendix A. The value deduced for normal operating conditions is about 1.3 failures per year per unit. The value deduced for the overload conditions experienced immediately following loss of commercial power is about 1000 times larger.

The exact value deduced for overload conditions varies with the assumptions made about how the probability of failure depends on the degree of overload. Arguments presented in the Appendix consider independence of overload, proportionality to overload and proportionality to the square of the degree of overload. The last assumption yields the highest probability for losing all OSG from cascading failures due to overload, and was adopted to derive the probabilities cited below.

For the case at hand, we assume that the total plant load is 160 MW (16 times the output of an average OSG), while we have only 6 OSG on line. (It would be customary to have 7-8 OSG on line.) Upon loss of commercial power, then, the immediate overload is 167%. We assume the absolute value of the extra load to remain constant for 20 minutes (the average duration of such incidents in the past). Thus, as individual OSG are lost, the degree of overload on the survivors would increase. And their probability of failure per unit time increases in proportion to the square of the overload. The calculated probability of a cascading loss of all six OSG (see Appendix) is about 10^{-3}.

If the number of OSG on line were only four, with the same total plant power load, the calculated probability of losing all four would be 0.15. We might apply such data (see Appendix) to arguing how many OSG should be kept on line. But the issue at hand is to estimate the probability of losing all OSG, given loss of commercial power, under typical operating conditions. And the answer seems to be something on the order of 10^{-3} or less.

b. Experience with Failure in Offline Tests

Power Technology has reported the history of failures of OSG insulation observed in offline testing.[5] A review of this experience, described in Appendix B, shows clear evidence of "wearout" of the OSG. The specific evidence is that the cumulative average rate of failure has doubled or tripled in recent years. This implies that the instantaneous failure probability is higher than it was just a few years ago, by a factor bigger than 2 or 3.

Power technology judges that some such acceleration in probability must also apply to failure in service, especially under overload conditions. If we assume a factor-of-3 increase probability for the case chosen for illustration in the preceding section (6 OSG, probability proportional to square of overload), the probability of losing all 6 OSG increases from about 10^{-3} to about 1/2.

3. Failure of Emergency Power

The emergency power system (GM diesels) would be called upon early to effect the valve changes required to recirculate the secondary coolant. It would then be called upon to operate the secondary recirculating pumps for (up to) 5 days. For the accident sequence diagram (Event Tree) of Figure 1, we assume that the only mode of failure of the emergency power system is failure to start on demand. The probability that the system starts well enough, but fails later, will be included in the next section (as one of the mechanisms by which we could fail to sustain recirculation of secondary coolant).

Experience cited in Reference 4 (Section E.1) puts the probability of failure of a single diesel to start on demand at 0.02 per demand. Since we have a recirculating pump powered by each of the two GM diesels (plus options for supplying any transformer room from either diesel), both diesels must fail to prevent recirculation. The probability of both diesels failing should be $(0.02)^2 = 0.4 \times 10^{-3}$.

4. Recirculation of Secondary Coolant

Recirculation of secondary coolant must be sustained for at least 5 days. The basic events contributing to failure are of two general kinds: equipment failures and human errors. The equipment failures are calculable (e.g., using data from WASH-1400) and are demonstrably small.

Calculation of the human errors is beyond the scope of this study. We simply assume that the probability of doing something that stops flow abruptly (e.g., close a valve or turn off a pump) and failing to take corrective action is on the order of 1/10 or less. But the probability of failing to correct subtler effects (e.g., wasted flow to auxiliaries) is assumed to be significantly larger, say 1/3.

The overall probability derived from these assumptions is about 1/2 that secondary cooling will be sustained.

5. Restoration of River Water Pumps

It is not necessary to restore full power to the 115 kV system to operate the few river water pumps required to start supplying cooling water to the reactor areas again. But, given the loss of all OSG, the most likely mechanism for restoration of power to river water pumps would seem to be recovery of the commercial power (CP) system, which would enable operation of all pumps.

We have no detailed statistics on the durations of CP outages at SRP. But some such information is presented in Appendix III of WASH-1400[6] (p. III-72). The data was fitted to a log-normal distribution. This means that the probability of having power restored by time t is given by

$$F(t) = \frac{1}{\alpha\sqrt{2\pi}} \int_0^t \exp\left\{-\left[\ln (y/\beta)\right]^2 / (2\alpha^2)\right\} dy$$

$$= \frac{1}{2}\left[1 \pm \operatorname{erf}(|z|)\right],$$

where $z = [\ln (t/\beta)]/(\alpha\sqrt{2})$. In the last expression, erf represents the error function tabulated in handbooks. Also, the + and - signs apply when $t>\beta$ and $t<\beta$, respectively. The parameter β was set equal to 0.12 hour to force agreement with the data of WASH-1400 at the point of 50% restoration. The parameter α was then set equal to

2.15 to force agreement with the data for greater than 50% recovery.

The fitted distribution yields probabilities of 0.03, 0.007 and 0.7×10^{-3} that power would not be restored within 8 hours, 1 day and 5 days, respectively. For the sake of consistency, we will use the more conservative values (derived and quoted earlier, but unpublished) derived by fitting a log-normal distribution to the more limited SRP experience. These values, shown in Figure 1, are 0.02, 0.005 and 0.5×10^{-3}, respectively.

6. Discharge of Fuel

There has been no explicit assessment of the probability of failure of the discharge process. The 5 day allowance for discharge is based on normal operating conditions, with provision for complete checkout of all equipment on both charge and discharge machines. Thus, the probability of completing discharge within 5 days under these (accident) conditions might be expected to be better than for completing a normal charge-discharge operation in 5 days.

On the other hand, the present case is subject to failure of the emergency power system and human error under stressful conditions. And one strong contributor to that stress is the challenge of discharging three reactors at once.

Another factor that might interfere with rapid discharge of a reactor is availability of space (hangers) in the spent fuel pool. This problem is aggravated by the need to discharge all reactors at once.

We must assign this task a fairly high probability of failure. We assume a value of 0.5.

C. Discussion of Results

The probability of suffering fuel melting because of a loss of 115 kV power varies directly with the probability of losing all OSG. This factor seems to be subject to the greatest variability of all of the factors involved.

There are 3 or 4 major reasons for variability in this factor. One of these is the number of OSG operating at the time the postulated loss of commercial power occurs. A typical number in practice is 7 or 8. We picked a value of 6 OSG for this illustration to err on the pessimistic side. The probability of losing all of 6 OSG is more than 3 decades higher than that of losing 8 OSG. The probability increases more slowly as we reduce the number of OSG working (by about 2 decades from 6 OSG to 4 OSG).

Another major source of variability of the probability of losing all OSG is the question of how (or whether) that probability depends on the degree of overload. There is a difference of about 3 decades between assuming independence and proportionality to the degree of overload. Roughly one more decade is incurred by assuming quadratic rather than linear dependence on degree of overload.

The third source of variability in probability of losing all OSG arises from the question of whether the evidence of "wearout" of the OSG, so apparent in offline testing of insulation, applies to the probability of failure in service, as well. The effect on probabilities is on the order of 3^N, where N is the number of OSG in service.

There is a fourth factor capable of exerting strong influence that was not considered in this study. This is the automatic load-shedding system now being installed. This system should reduce the period of severe overload by a factor on the order of 10^2. The corresponding effect on probability of losing all of N OSG should be roughly 10^{2N}.

To extract some kind of stable conclusions from these wide variations we make some assumptions and retain some distinctions among past, present and future. Thus, we assume only 6 OSG online and assume that the probability of failure varies as the square of the degree of overload.

The temporal distinctions are to declare separate estimates of what the probabilities have been in the past few years (denying evidence of wearout), what they may be in the immediate future (accepting the evidence of wearout but denying the benefits of automatic load-shedding), and what they should be in the near future (taking credit for automatic load-shedding).

JAS:cor

REFERENCES

1. M. E. Bacinoglu, "Automatic Load Shedding System", TA 5-16, March 1983.

2. R. A. Allen, "100 Areas Automatic Load Shedding", TA 1-2144 (DPSOX 9622), June 1983.

3. D. C. Wood, "186 Basin Recirculation", DPST-82-918, October 1982.

4. J. A. Smith, "SRP Electric Power Systems", DPST-83-249, January 1983.

5. C. G. Hardin, Memorandum, May 24, 1983.

6. U. S. NRC, "Reactor Safety Study", WASH-1400 (NUREG-75/014), October 1975.

7. W. H. Baker, "Cooling Water Conservation - 100 Areas", RTM-4542, June 1983.

OSR 25-10 (Rev 6-78) ENGINEERING DU PONT COMPUTATION SHEET SHEET No. Pg-11

TITLE OF PROJ. [illegible] PROJ. NO. Fig. 1

SUBJECT WORKS

COMPUTER DATE 19

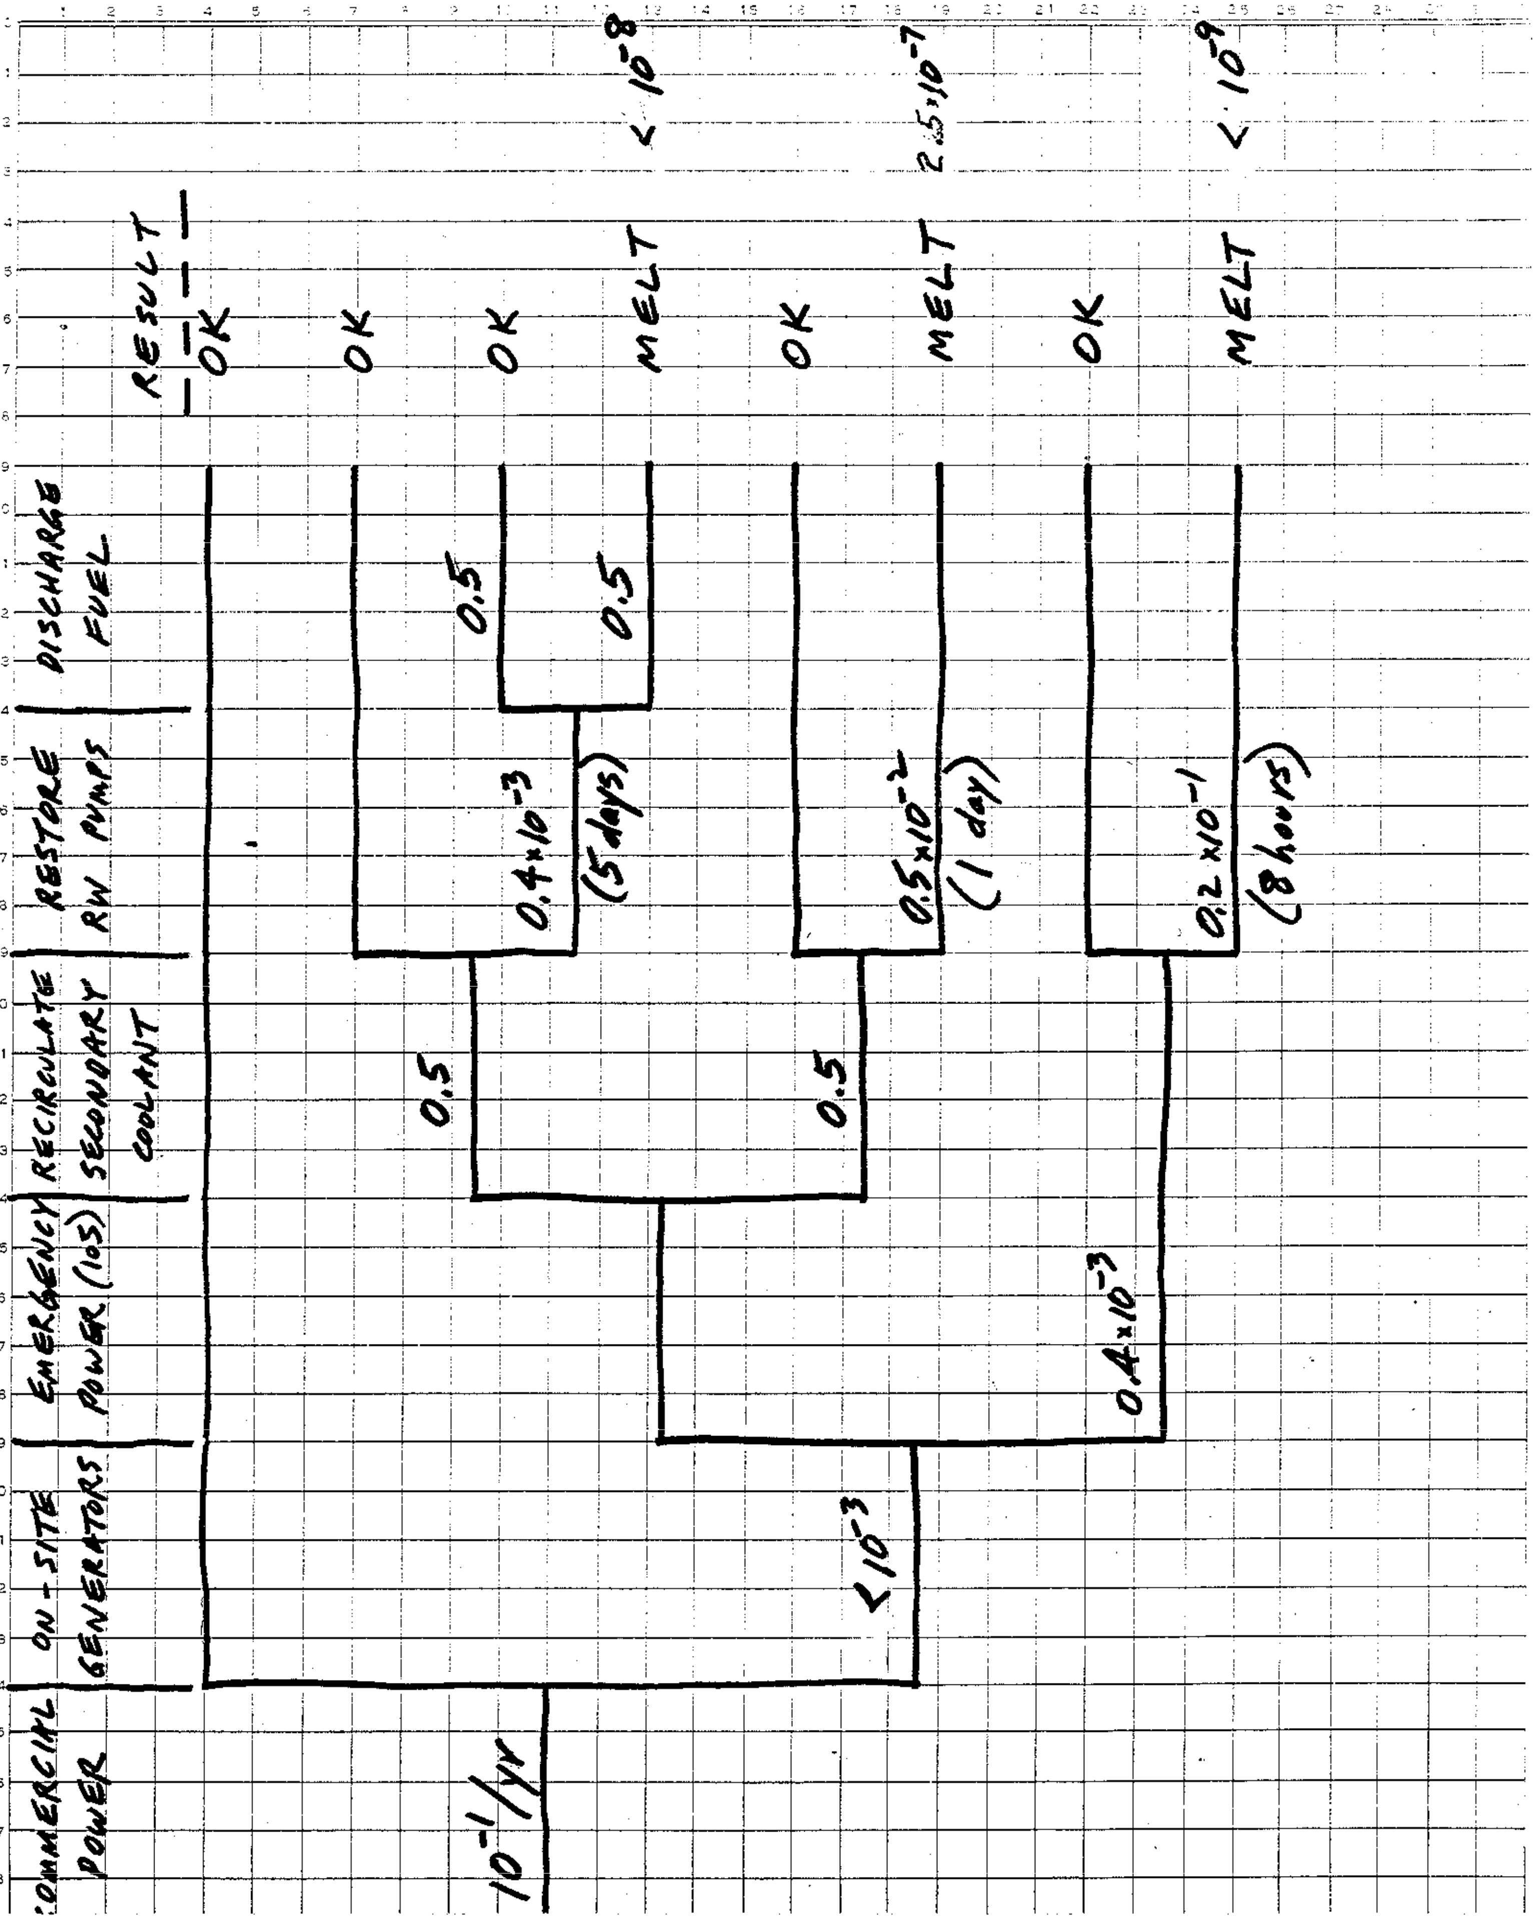

APPENDIX A: PROBABILITY OF LOSING ON-SITE GENERATORS

On-site generators (OSG) have failed both during normal operating conditions and during the stressful conditions of overload. This experience provides a basis for predicting the frequency of OSG losses under future conditions, both real and hypothetical.

For normal operating conditions we can postulate that the probability of failure per unit time for a given OSG unit is some constant value, f. Once lost, the unit could be restored to operational status after an interval, R, on the average. A given unit should therefore be available for operation for a fraction of the time (roughly) equal to (1-fR). If we had N such units in operation, we should expect to observe failures at a rate,

$$N(1-fR)f$$

per unit time. Power Tech observes[5] that we experience about 10 failures per year with an average of about 8 units on line. Although the mean time to repair has not been established, a reasonable estimate is about 2 weeks (and the results here are not sensitive to that number). From these observations, we can deduce a value for f of about 1.3 failures per year per unit during normal operation.

During conditions of significant overload, we should postulate a different value of f (call this one F) on grounds that new forces come into play that could increase the probability of failure. But overload conditions arise infrequently and persist for only a short time. So repairability and unavailability (due to failure earlier in the interval of interest) are not significant factors in this case. If we had N OSG on line during overload conditions, we should expect to observe failures at a rate, NF, per unit time. At the end of an interval, Δt, we should expect to observe a number of failures equal to $NF\Delta t$.

Our experience with severe overload conditions includes the four cases of losing all commercial power. The relevant characteristics of these cases are itemized in Table 1. After these four incidents, we should expect to observe a number of failures equal to

$$F(N_1\Delta t_1+N_2\Delta t_2+N_3\Delta t_3+N_4\Delta t_4),$$

assuming the same failure probability applies to all four incidents. We did observe two failures after the four incidents. Using this value, and values for N_i and Δt_i from Table 1, we can deduce a value for F of about 1500 failures per year per unit.

This result indicates a probability of failure under overload conditions very much larger than the value under normal conditions. If the probability is so sensitive to overload, then it is apt to vary with the degree of overload. It might be reasonable to assume direct proportionality to degree of overload. If so, we could express the probability of failure per year per unit as FL, where F is the value under 100% overload conditions, and L is the ratio of extra load actually experienced to normal load.

To quibble for a minute, this probability (FL) would decrease to zero under normal load conditions. So the "true" probability should be expressed as (f+FL). As long as we are considering severe overload conditions, though, we are practically correct in using just (FL).

If we assume direct proportionality (probability=FL) and apply the data of Table 1, we can deduce a value for F (under 100% overload) of about 1200 failures per year per unit.

It might also be reasonable to assume proportionality to some higher power of L. For example, if we assumed a probability of FL^2, we deduce a value for F of about 860 failures per year per unit under 100% overload.

Armed with an explicit (although presumed) value for probability of failure of OSG under overload, we can make quantitative estimates of the probability of cascading failures that might result from loss of commercial power.

I needed some help with the probability formulas. R. L. Postles of CSD provided that help by pointing out that our case (characterized by discrete states and continuous time) fits into an established system called the "Markov chain model". The technique is described in standard texts on reliability analysis (e.g., McCormick's "Reliability and Risk Analysis", Chapter 7). The results are described below. The derivation for the case at hand is outlined in Appendix C.

If we start with a total of N units, and the probability of losing a unit changes with the number left on line, then the probability of losing all N units after time t is

$$P_N(t) = \sum_{i=0}^{N-1} \frac{PK}{PI} \frac{F_{N-1}}{F_i(N-i)} \left(1 - e^{-(N-i)F_i t}\right),$$

where F_i represents the probability of failure per year per unit, given that i units have failed already, and

$$PK = \prod_{K=0}^{N-2} (N-k)F_K ,$$

$$PI = \sum_{\substack{K=0 \\ K \neq i}}^{N-1} \left[(N-k)F_k - (N-i)F_i\right] .$$

The expression proves difficult to evaluate because it involves small differences of large terms. It helps to expand the exponential into a series. Terms involving $F_i t$ to a power less than (N-1) drop out. The expression then reduces to

$$\sum_{i=0}^{N-1} \frac{PK}{PI} \frac{F_{N-1}}{F_i(N-i)} \frac{\left[F_i(N-i)t\right]^N}{N!} \left[1 - \frac{F_i(N-i)t}{N+1} + \frac{(F_i(N-i)t)^2}{(N+1)(N+2)} - \cdots\right] .$$

The preceding equations do not apply to the special case in which the probability of failure per year per unit changes in direct proportion to the degree of overload. (In this case, $(N-i)F_i$ is the same for all values of i.) The probability of losing all N OSG for this case becomes

$$P_N(t) = e^{-x} \frac{x^N}{N!} \left[1 + \frac{x}{N+1} + \frac{x^2}{(N+1)(N+2)} + \cdots\right],$$

where $x = NF_0 t$.

The "bottom line" answer we want from such formulas is the probability of losing all N of OSG, given that the commercial power source has already failed. The numerical answer we calculate depends on what degree of sensitivity to overload we assume and to the general boundary conditions we assume (how many OSG on line, what degree of overload immediately after losing CP and how long the overload persists).

Consider recently typical conditions of eight OSG on line, an immediate overload of about 100%, and a 20 minute duration for the overload interval. Then the calculated chance of losing all generators would be too small to be considered seriously. If we assumed the probability of failure to be proportional to the square of the overload factor, the probability of losing eight OSG would be on the order of 10^{-7} (according to our formulas). Probabilities this small are automatically suspected of errors of oversight. But it does seem safe to conclude that the phenomenon of cascading losses, due to increasing overload, is not apt to cause loss of all OSG if we have a full complement of OSG on line.

Another possibility that would aggravate the overload problem would be to operate with fewer-than-normal OSG, so the immediate overload factor would be larger and the total number of OSG available to fail is smaller. For example, suppose the CP load were equivalent to that of 8 OSG, so the immediate overload factor would be 100% if we had 8 OSG on line. If, in fact, we have only 6 OSG on line, the immediate overload factor would be 10/6=167%. Again assuming probability of failure to be proportional to the square of overload, the probability of losing all six OSG would be about 10^{-3}.

If the number of OSG on line just before the loss of CP were reduced to four, while the total plant load remained the same, (causing an immediate overload factor of 300%) the calculated probability of losing all four OSG would be about 1/7. This value was calculated assuming probability proportional to overload squared. If the dependence on overload were only linear, the calculated probability would drop to about 1/200.

Playing further with these formulas, conditions and numbers might lead us to some defensible bases for arguing how many OSG must be on line for a given plant electrical load. But they already suggest that the place to start worrying is somewhere around half the normal number of OSG. We cannot justify calling for special actions because we have one or two fewer OSG than would be normal for a given plant load.

All of the preceding numerical results were calculated for the heretofore typical duration of 20 minutes. If we had automatic load-shedding to reduce that duration to a few seconds, the calculated probabilities would drop to negligible values. Reducing the duration by a factor on the order of 10^2 reduces the probability of losing N OSG by a factor on the order of 10^{2N}. In other words, cascading failures due to overload should not be a problem worth worrying about.

Table 1: Data from Periods of OSG Overload

Date	No. of OSG on line	Duration, Minutes	Overload Factor[a]
2/61	10[b]	10	1.33[b]
10/70	10	19	0.79
12/79	7	38	1.43
6/81	7	18	1.75

Notes: (a) Ratio of load carried by commercial power to load carried by OSG, before losing commercial power.

(b) These data were not recorded. But there were 5 reactors on line, implying a plant electrical load of about 210MW. We assume 10 OSG on line, putting out 90 MW, as in 10/70.

Table 2: Cascading Failure Probabilities

No. of OSG on line	Probability of Losing All OSG in 20 Minutes[a] independent[b]	linear[b]	square[b]
8	.89 E-10	.31 E-7	.14 E-6
7	.16 E-8	.12 E-5	.20 E-4
6	.29 E-7	.29 E-4	.80 E-3
5	.52 E-6	.48 E-3	.17 E-1
4	.94 E-5	.53 E-2	.15
3	.17 E-3	.39 E-1	.54
2	.31 E-2	.19	.92
1	.55 E-1	.57	1.00

Notes: (a) Assume plant load constant at equivalent of 16 OSG. If 8 OSG on line, overload upon loss of commercial power is 100%. If N generators on line, overload is (16-N)/N. Assumes no load shedding for 20 minutes.

(b) Probability of failure assumed to vary with overload in 3 ways: independent, proportional to OL or proportional to $(OL)^2$.

APPENDIX B: OSG INSULATION FAILURES

Power Technology issued a memorandum(5) describing the history of failures of insulation tests performed on OSG that were out of service. The memo also cited the history of overloads of the OSG (when commercial power is lost).

Stator failures were detected in the process of conducting offline tests in which a high potential is imposed, looking for a breakdown of the insulation. If a breakdown is observed, the stator must be rewound. Five such breakdowns were reported. Rotors are not tested in the same way. The two reported failures were detected by inspection. Only the stator data were considered for quantitative analysis.

The raw data is summarized in Tables B-1 and B-2. Startup dates for the units were not given, so I made the assumptions shown in Table B-1.

The Power Tech memo cited 7 cases in which the OSG were severely overloaded, presumably, by loss of all offsite power. These incidents were dated 10/54, 2/56, 7/56, 2/61, 10/70, 10/79 and 6/81. I found no evidence of the first two in Reactor Tech monthly reports at those times. So I entertained the notion that those overloads were not as severe as the last 5. Thus, Table B-2 lists two sets of data for the number of overloads experienced by each unit.

In their memo(5) Power Tech made note of their perception of a correlation between failures and overloads. I tried to quantify this proposition by postulating an exponential distribution of failures as a function of number of overloads. This means that, for a small number of overloads, N, the probability of failure should be expressible as N. And, in general, the probability of failure (i.e., the fraction of available units that can be expected to fail) should be given by

$$P = 1 - e^{-\lambda N}$$

The data of Table B-2 give us two milestones at which to test this hypothesis. By the Power Tech count, there were 3 failures by the time we reached 4 overloads, while 9 other units operated that long without failure. The probability of failure is therefore $3/12 = 0.25$, and $\lambda = 0.072$ per overload. There were 5 failures by the time we reached 6 overloads, while 4 other units operated that long without failure. The probability of failure is therefore $5/9 = 0.56$, and $\lambda = 0.14$ per overload.

By my count of the number of overloads, the milestones occur at 2 and 4 overloads. At 2 overloads, P =2/12 = 0.17 and $\lambda = 0.091$. At 4 overloads, P=5/10 =0.50 and $\lambda = 0.17$.

While the two ways of counting overloads yield different absolute values for λ, they yield one common result: the postulated exponential distribution does not fit very well. There are any number of ways to interpret this misfit. But all such rationales have one point in common: if the probability of failure should, in fact, be correlated with the number of overloads, then that probability is not simply proportional to the number of overloads. It seems to accelerate as the number of overloads increases. The cumulative average failure rate roughly doubled between 4 and 6 overloads (or between 2 and 4 overloads). This implies that the instantaneous failure rate accelerated by a factor more than two.

The phenomenon just cited, accelerating failure rate, may be rationalized by asserting that the equipment is simply "wearing out" after repeated overloads.

The same effect might be expected if we try to correlate the failure rate with service age. The data of Table B-2 provide 4 milestones to try fitting an exponential distribution based on age, $P = 1-e^{-\lambda t}$. The results, summarized in Table B-3, show a similar result: the exponential distribution does not fit well. In this case, the cumulative average failure rate triples from the first to the last milestone, even stronger evidence of wearout.

If engineering judgement leads us to expect correlation between failure rate and either or both of service age and overload history, we must perceive clear evidence of wearout of the OSG stators. The only tempering consideration that occurs to me is that the statistical sample might be too small to warrant this conclusion. I will pursue this question with professional statisticians.

Engineering judgement by Power Tech leads them to draw another conclusion: if the probability of failure of offline insulation tests is accelerating, then so is the probability of failure in service. I do not feel competent to make this judgement. Neither did Pickard, Lowe & Garrick (PLG), the consulting firm now analyzing reliability of the electric power system. But PLG promised to pursue the question with their industry contacts.

TABLE B-1: OSG Insulation Failure History

Unit	Assumed Start Date	Date of Failure Rotor	Date of Failure Stator	Date of Retirement	Age of Failures R	Age of Failures S	Age of Survivors (a) R	Age of Survivors (a) S	Age of Replacements (a) R	Age of Replacements (a) S
R1	1/54			6/64	-	-	10.5	10.5	-	-
R2					-	-	10.5	- (c)	-	-
P1	1/54				-	-	29.5	29.5	-	-
P2		2/82 (b)	6/66		28.1	12.4	-	-	1.3	27.5 (c)
L1	7/54			6/68	-	-	14.0	14.0	-	-
L2					-	-	14.0	14.0	-	-
K1	10/54		6/67		-	12.7	28.7	-	-	16.0
K2		2/83	4/80 (b)		28.3	25.5	-	-	0.3	3.2
HP1	1/53				-	-	30.5	30.5	-	-
HP2					-	-	30.5	30.5	-	-
HP3			5/82 (b)		-	29.4	30.5	-	-	1.1
LP1	1/55				-	-	28.5	28.5	-	-
LP2			2/83		-	28.1	28.5	-	-	0.3
LP3					-	-	28.5	28.5	-	-
LP4					-	-	28.5	28.5	-	-

Notes: (a) as of 6/83
(b) failed within one year of last overload
(c) the P2 stator that failed in 1966 was replaced with one from R

TABLE B-2: Stator Failure Histories

Unit	Failers	Survivors	No. of Severe OLs[a] Failers	No. of Severe OLs[a] Survivors	No. of Severe OLs[b] Failers	No. of Severe OLs[b] Survivors
R1	-	10.5	-	0	-	1
R2	-	-[c]	-	-	-	-
P1	-	29.5	-	3	-	5
P2	12.4	27.5[c]	2	4	4	6
L1	-	14.0	-	1	-	3
L2	-	14.0	-	1	-	3
K1	12.7	16.0	2	1	4	1
K2	25.5	-	4	-	6	-
HP1	-	30.5	-	4	-	6
HP2	-	30.5	-	4	-	6
HP3	29.4	-	4	-	6	-
LP1	-	28.5	-	4	-	5
LP2	28.1	-	4	-	4	-
LP3	-	28.5	-	3	-	4
LP4	-	28.5	-	5	-	6

Notes: (a) by JAS count
(b) by Power Tech count
(c) the P2 Stator that failed in 1966 was replaced with one from R.

TABLE B-3: Probability of Failure with Service Age

Service Age, Years	No. of Failures At Or Before That Age	No. Of Survivors At That Age	P(fail)	λ, y^{-1}
12.5	2	13	2/15	0.011
25.5	3	9	3/12	0.011
28.1	4	7	4/11	0.016
29.4	5	3	5/8	0.033

APPENDIX C: Derivation of OSG Failure Probabilities

We have a system of N onsite generators, some of which may fail. Let P_n denote the probability that precisely n of them have failed (not more, not less). If n units have already failed, we have (N-n) units left. In this situation, the chance of one more unit failing is $(N-n)F_nP_n$, where F_n is the failure probability per year per unit, discussed in Appendix A.

For n<N, we can write a general equation for the rate of change of P_n,

$$\dot{P}_n = -\lambda_n P_n + \lambda_{n-1} P_{n-1} ,$$

$$\lambda_n = (N-n) F_n .$$

For n = N,

$$\dot{P}_N = \lambda_{N-1} P_{N-1} .$$

The last equation is different because no more OSG are available to fail after all N have failed.

These equations are the same as those for a radioactive decay chain, with the familiar solutions,

$$P_0(t) = e^{-\lambda_0 t}$$

$$P_1(t) = \frac{\lambda_0}{\lambda_1 - \lambda_0}\left(e^{-\lambda_0 t} - e^{-\lambda_1 t}\right)$$

etc.

As long as no two λ_i are the same, the general solution is

$$P_n(t) = \sum_{i=0}^{n} \frac{\left(\prod_{k=0}^{n-1} \lambda_k\right)}{\prod_{\substack{k=0 \\ k \neq i}}^{n} (\lambda_k - \lambda_i)} e^{-\lambda_i t} , \quad n < N .$$

But we have one special case in which all λ_i are equal. When F_n increases with n in direct proportion to the degree of overload, as discussed in Appendix A, F_n is inversely proportional to (N-n). In this case, the general solution is

$$P_n(t) = \frac{(\lambda t)^n}{n!} e^{-\lambda t}, \quad n < N.$$

Integrating the last equation (n = N) yields a slightly different result. And applying the boundary condition, $P_n(o) = 0$, yields

$$P_N(t) = \sum_{i=0}^{N-1} \frac{\left(\prod_{k=0}^{N-2} \lambda_k\right)}{\prod_{\substack{k=0 \\ k \neq i}}^{N-1} (\lambda_k - \lambda_i)} \left(\frac{\lambda_{N-1}}{\lambda_i}\right)\left(1 - e^{-\lambda_i t}\right), \quad \lambda_i \neq \lambda_k .$$

For $\lambda_k = \lambda_i = \lambda$, it is easier to calculate the probability of losing all OSG from

$$P_N(t) = \sum_{k=N}^{\infty} P_k(t) = e^{-\lambda t} \frac{(\lambda t)^N}{N!}\left[1 + \frac{\lambda t}{N+1} + \frac{(\lambda t)^2}{(N+1)(N+2)} + \cdots\right].$$

The last equation, for equal λ, can be evaluated without difficulty by substituting the appropriate λ t and N values. But the preceding equation, for $\lambda_i \neq \lambda_k$, is difficult to evaluate because it involves small differences of many terms. This problem can be avoided by expanding the $e^{-\lambda_i t}$ terms into series, and taking advantage of the fact that all terms of order less than $(\lambda t)^N$ drop out.

The result is

$$P_N(t) = \sum_{i=0}^{N-1} \frac{\left(\prod_{k=0}^{N-2} \lambda_k\right)}{\prod_{\substack{k=0 \\ k \neq i}}^{N-1} (\lambda_k - \lambda_i)} \left(\frac{\lambda_{N-1}}{\lambda_i}\right) \frac{(\lambda_i t)^N}{N!}\left[1 - \frac{\lambda_i t}{N+1} + \frac{(\lambda_i t)^2}{(N+1)(N+2)} - \cdots\right].$$