MEMORANDUM
RM-3367-PR
FEBRUARY 1963

USES OF MONTE CARLO IN PERT

Richard Van Slyke

PREPARED FOR:
UNITED STATES AIR FORCE PROJECT RAND

The RAND Corporation
SANTA MONICA • CALIFORNIA

28p

MEMORANDUM
RM-3367-PR
FEBRUARY 1963

USES OF MONTE CARLO IN PERT

Richard Van Slyke

This research is sponsored by the United States Air Force under Project RAND—contract No. AF 49(638)-700 monitored by the Directorate of Development Planning, Deputy Chief of Staff, Research and Development, Hq USAF. Views or conclusions contained in this Memorandum should not be interpreted as representing the official opinion or policy of the United States Air Force. Permission to quote from or reproduce portions of this Memorandum must be obtained from The RAND Corporation.

The RAND Corporation
1700 MAIN ST. • SANTA MONICA • CALIFORNIA

PREFACE

This Memorandum is concerned with the application of Monte Carlo, or random-variable, techniques to the PERT management system. It should be of particular interest to operations and systems analysts working with PERT or other complicated management systems and of general interest to statisticians and mathematicians concerned with numerical methods of analysis.

The author, a graduate student at the University of California, Berkeley, participated in RAND's 1962 summer program for graduate students.

SUMMARY

The mathematical assumptions underlying the PERT (Program Evaluation and Review Technique) management system are of doubtful validity, and even granting these assumptions there are still serious computational difficulties involved in getting the desired answers. This report outlines some of the weaknesses of the present system, and suggests how the use of Monte Carlo methods can lead to improvement, first in allowing less restrictive mathematical assumptions to be made, and second in extending the kinds of computational results that can be obtained. Moreover, the method can be used to check the validity of the commonly used approximations.

A "criticality index" for an activity is defined; it is the probability that an activity will lie on a critical path. This is an example of a quantity that can be calculated by Monte Carlo methods, but not by presently used methods.

A experimental computer program for Monte Carlo treatment of PERT networks has been coded for the IBM 7090 computer. The program is discussed, and some shortcuts are suggested for reducing the time of computation.

CONTENTS

USE OF MONTE CARLO IN PERT

1. INTRODUCTION

By a management system such as the PERT system* we mean a mathematical model of the project to be managed, methods of collecting data for the model, solution techniques, and a set of outputs. The usefulness of such a system depends on the extent to which it can answer the following questions:

(a) How well does the model represent the object of interest? Does it include all the relevant features of the project in question?

(b) What are the data required? Are they easily available, and if so how accurately?

(c) How easy is it to do analysis on the model? Are the available solution techniques adequate? What approximations have to be made?

(d) How much information do you get? Is it in a form that is relevant to the manager's problems and to the decisions he must make? Is there so much output that he cannot make sense of it, or so little that he does not have the information necessary for action?

These four groups of questions clearly are not independent; e.g., the solution method certainly is intimately connected with what output is desired and what input is required. Nevertheless, these areas form a convenient basis for evaluation. Similarly, the interfaces between the assumptions of the model itself, and the assumptions

---

*We distinguish between a management _model_ and a management _system_. The latter includes the solution method, data collecting procedures, and outputs in addition to the model itself.

used in solution and in data-collecting methods, are not always well defined. Therefore what is given here as "the PERT model," for instance, is highly subjective and must be considered as one man's opinion.

In Sec. 2 a somewhat abstract version of the PERT model is given. See References [1, 3, 5] for a more extended discussion. In Sec. 3 we see how the current PERT system answers our questions, and in Secs. 4 and 5 we see what the Monte Carlo approach is and how it answers the questions posed. Some of the appropriate statistical analysis appears in Sec. 5, and in Sec. 6 some computational short cuts are introduced. In Sec. 7 some results using an experimental code are given, as well as some suggestions for further study. Finally in the appendix the results of a Monte Carlo simulation are compared with PERT output for a sample network.

2. THE PERT MODEL

We are given a finite acyclic directed network of arcs, which we call <u>activities</u>, and their nodes, which we call <u>events</u>. Associated with each activity is a nonnegative random variable, which we call its <u>duration</u>. There are two special events, the <u>initial</u> event and the <u>terminal</u> event, such that no activity leads into the initial event or out of the terminal event and such that each activity is contained in a directed path leading from the initial event to the terminal event. We assume that the durations have independent distributions, each with a finite range. A <u>realization</u> of the network is the network with a fixed value for each of its durations. For a particular realization of the network, we call a longest path from the initial to the terminal event a <u>critical path</u> and its length the <u>project duration.</u> We often speak also of the random variable corresponding to the project duration and of its distribution.

3. THE PERT SYSTEM

The PERT model given in Sec. 2 has some fairly obvious deficiencies. For instance, it cannot handle:

(a) Resource allocation. Activities often do not have to occur in any set sequence, but use the same resources such as manpower and facilities so that they cannot all be carried on simultaneously. This situation is not taken into account by the present model.

(b) Correlations between activity durations. As an example, suppose that one approach to a design problem did not work and another had to be used. Many of the following activities would have entirely different distributions depending on which alternative was used. This would introduce correlation among these activity durations.

(c) Feedback. The data and estimates are obtained from engineers and other technical people in the absence of schedule dates and deadlines. Then this data is used to generate schedule dates, which in turn affect the duration times.

(d) "Beating the system." Finally, there are the confounding effects of the technical people trying to "beat the system" when they give estimates. This is not difficult to do and must be taken account of in any realistic use of the system.

The basic data required for the PERT system is the distribution of the activity durations. It is here that the approximations are particularly gross, if to some extent unavoidable. The data for these distributions are obtained from technical people who have had some experience with the type of activities involved, although in research and development projects, for example, the activity in question may never before have been attempted. The most one can expect are estimates for a few parameters of the distribution, commonly the range and the mode. Since, as will be seen later, only the mean and variance of the distributions are used in current calculation methods, the character of the distribution is treated somewhat cavalierly. The distribution is assumed to be a beta distribution with standard deviation equal to 1/6 the range. These are of course highly arbitrary assumptions and should not be taken too seriously.

We shall be concerned primarily with the computational aspects of the system. The procedure in current use is as follows:

(a) Reformulate the problem as a deterministic problem, using the expected value of an activity duration as its deterministic length.

(b) Find the corresponding critical path P_E.

(c) Use the associated project duration ℓ_E as the mean of the project-duration distribution; and use the sum of the activity variances along the path as the variance of the true project duration, which is then assumed to be normal.

(d) Calculate the probability that schedule dates will be met.

The assumptions used to justify this approach are mainly the following:

(A) A critical path P_E of the derived problem is assumed to be "enough longer"* than any other path so that the probability of a realization having a different critical path is negligible.

(B) The critical path P_E has enough activities so that a central limit theorem applies.

We note, however, that assumption (B) is not needed for calculating the expected length of the critical path or to calculate its variance, but is used only for making inferences about probabilities of meeting schedule times.

---

*A very loose measure of this for individual alternative paths can be obtained using Kolmogoroff's inequality [9], p. 107.

To illustrate what can happen when these assumptions are not satisfied, we consider the following simple network:

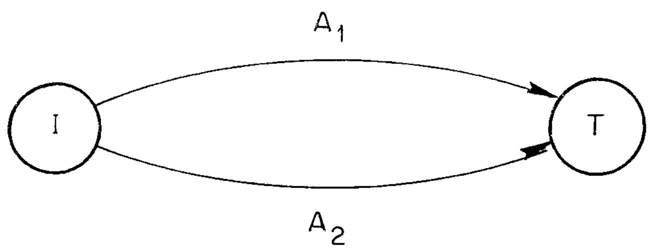

Here the arc durations A_1 and A_2 are normal,* with means μ_1, μ_2, and variances σ_1^2, σ_2^2, respectively. Clark [2] shows that for the random variable which is the maximum of A_1 and A_2, the mean is given by

$$\mu_M = \mu_1 \, \Phi(\alpha) + \mu_2 \, \Phi(-\alpha) + a \, \varphi(\alpha),$$

and the variance by

$$\sigma_M^2 = (\mu_1^2 + \sigma_1^2) \, \Phi(\alpha) + (\mu_2^2 + \sigma_2^2) \, \Phi(-\alpha) + (\mu_1 + \mu_2) \, a \, \varphi(\alpha) - \nu_1^2,$$

where

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-t^2/2} \, dt, \quad \varphi(t) = \frac{1}{\sqrt{2\pi}} \, e^{-t^2/2},$$

$$a = \sqrt{\sigma_1^2 + \sigma_2^2}, \qquad \alpha = \frac{\mu_1 - \mu_2}{a},$$

---

*We should realize, of course, that strictly speaking we should not use normal distributions because they do not have finite ranges.

and μ_M and σ_M are the mean and variance of the distribution of the maximum. Assume for simplicity that $\mu_1 = \mu_2$; then $\alpha = 0$ and $\mu_M = \mu_1 + \sqrt{\sigma_1^2 + \sigma_2^2}$. The method outlined above gives μ_M "=" μ_1. So as σ_1^2 or σ_2^2 increases, we can make the difference, $\Delta = \sqrt{\sigma_1^2 + \sigma_2^2}$, between the correct mean and the mean calculated using the above procedure arbitrarily large. Similarly, we have

$$\sigma_M^2 = \frac{1}{2} (\sigma_1^2 + \sigma_2^2),$$

where the above procedure would consider σ_M^2 "=" σ_1^2 or σ_2^2, so that the estimate could be 100 per cent too high or arbitrarily too small. Similar exercises could be indulged in for other distributions than the normal. The point is simply that it is trivial to construct examples violating (A) for which the above technique will yield almost arbitrary errors.*

More interesting questions are: What does "enough longer" mean? and Given a network, how can we discover if (A) is satisfied? Two factors enter here; one is the "closeness" of two paths, and the other is the number of paths that may become critical. Even if each alternative path has a small probability of being critical, if there are enough of these alternative paths the error can accumulate significantly. In comparing two paths for criticality, the following considerations are important:

---

*This is not entirely true; for example, the estimate of expected project—duration time will always be optimistic under the above procedure. This can be shown in several ways. One informative way is to write the problem as a linear program with a random constant vector, the components of which are simply the activity durations. Then the general results of Madansky [4] for stochastic linear programming problems can be applied. In particular, for any linear programming problem with a random right—hand side the answer obtained by using expected values is always too optimistic. Another connection between linear programming and PERT is the relation between sensitivity analysis and the short cut of Sec. 2.

(a) The means of the path lengths.

(b) The variances.

(c) The correlation of paths, i.e., the activities the paths have in common.

Probably a large fraction of networks encountered in practice will satisfy assumption (A) at least to the extent needed for the desired accuracy, but general analytic methods of determining this are hard to come by because of the strong dependence on the particular network involved. The problems connected with assumption (B) are classical and will not be gone into here except to note that the probabilities assigned to meeting schedule dates depend on all the assumptions and are therefore especially unreliable. Nevertheless, the current procedure is extremely simple and rapid, and any other solution method yielding significantly improved results must be expected to take considerably more computation time.

The purpose of the PERT system is to help management spot bottlenecks and overruns before they occur, so that correction action can be taken before it is too late. Emphasis is on providing a tool for "management by exception," that is, to point up the few activities that are particularly critical and require close managerial supervision.

To do this, the current PERT system essentially gives us output "early" and "late" times for each activity, and the probability of meeting given bench-marks or schedule dates. Since these are expected values, a variance for these numbers is also given to indicate the reliability of the estimates. Generally in this Memorandum we will restrict our attention to the expectation of the project duration, its variance, and the distribution itself. All the other commonly given output is obtained using essentially the same methods.

One of the more misleading aspects of current PERT solution methods is the implication that there is a unique critical path. In general any of a number of paths could be critical, depending on the particular realization of the random activity durations that actually occurs. Thus it makes sense to talk about a "criticality index," which is simply the probability that an arc will be on the critical path. Current solution procedures by their very nature cannot give any information on this.

In the next sections we shall see how applying Monte Carlo techniques can alleviate some but by no means all of these problems.

4. SOLVING THE PERT MODEL USING MONTE CARLO TECHNIQUES

When we say that we "Monte Carlo" a PERT network, we simply mean that we apply the longest path algorithm to a long series of realizations each one obtained by assigning a sample value to every activity drawn from its proper distribution. Given this information, we use standard statistical methods to estimate the distribution and parameters of interest. This technique can be used in two basically different ways. The first involves a relatively small sample size and is used to test, for example, the classical type of hypothesis as to whether the true mean equals ℓ_E, to give some indication as to whether assumption (A) is violated. These statistical tests could be sequential and involve the application of standard statistical techniques. This approach has the disadvantage that it gives only incidental information about the distribution of the project duration.

The second application of the technique offers more fascinating opportunities, and that is what will be considered here. That is to take a large number of samples and Monte Carlo the network actually to obtain the answers rather than to run a check on traditional methods. Let us see how this approach can improve our answers to the four sets of questions in Sec. 1.

In the Monte Carlo framework, the model can be generalized to some extent. Including resource allocation would be complicated but could be accomplished, especially for certain cases. For example, if two activities use the same scarce resource, things could be set up in such a way that the first activity to start would have to finish before the second could begin, or, for three activities, the last to start might have to wait for the first two to finish. Alternatively, the three could be constrained to go one at a time, etc. Correlation between activities could very easily be handled with Monte Carlo techniques. For obvious reasons, the other shortcomings mentioned in the preceding section, namely feedback and "beating the system," could not be handled by this approach.

The data-collection problems are just as grevious when treated by Monte Carlo techniques as when treated by the general methods, if not more so. In current solution methods, much of the output does not depend on the structure of the activity duration distributions but only on their means and variances. The Monte Carlo approach, in order to gain extra accuracy, does depend on the shape of the distribution. On the other hand, the Monte Carlo approach has greater flexibility in that any distribution can be used for activity durations—beta, normal, triangular, uniform, or discrete in any sort of mix. This flexibility allows one, in particular, to try different distributions and observe the effect of neglecting or making highly arbitrary assumptions on the shape of these distributions.

It is in the area of computational accuracy, however, that the Monte Carlo techniques yield the greatest benefits. For example, the estimate for the mean using the procedure of Sec. 3 is always low, whereas the Monte Carlo procedure gives an unbiased estimate.

Again, while the earlier procedure is the least accurate in making estimates of the probability of meeting schedule dates, we shall see in the next section that the Monte Carlo technique yields very good answers indeed and in particular makes no recourse to central limit theorems.

Finally, with respect to output, this approach yields all the output of previous methods plus information on the "criticality indexes," or probabilities of activities being on the critical path.

5. ACCURACY OF MONTE CARLO TECHNIQUES

The only type of statistical analysis we shall investigate here is predicated on a fixed number of trials determined before sampling. We then ask what is the probability that our estimators for the various parameters and distributions will be "close" to the correct ones.

The first parameter we shall consider is the mean of the project duration. Each complete realization is considered as one sample observation, and the estimator we consider is the sample mean. More formally, let X be the random variable (r.v.) corresponding to the project duration. Let X_i, $i = 1, \ldots, N$, be N independent r.v.'s each with the same distribution as X. Then $(X_1, \ldots, X_N)$ is our sample. The r.v.

$$\hat{\mu} = \frac{\Sigma X_i}{N}$$

is the estimator of

$$\mathrm{Exp}\ \{X\} \equiv \mu_X .$$

Let

$$\sigma_x^2 = \mathrm{Var}\ \{x\};$$

then we have

$$\mathrm{Var}\ \{\hat{\mu}\} = \sigma_x^2/N.$$

Since we shall be talking in terms of thousands of samples, the central limit theorem tells us that $\hat{\mu}$ is to a very good approximation normal with mean μ_x and variance σ_x^2/N, denoted Nor $(\mu_x, \sigma_x^2/N)$. Clearly as N increases, the uncertainty in our estimate of Exp $\{x\}$ decreases. Since the standard deviation, i.e., the square root of the variance, is in the measure of spread with the same units as X, we see that roughly speaking the accuracy increases inversely with the square root of N.

Example: Suppose with probability 0.95 we want our estimate of the mean to be within $\sigma_x/50$ of the true mean. That is, we want to choose N so that

$$\Pr\ \{\mu_x - \frac{\sigma_x}{50} < \mu < \mu_x + \frac{\sigma_x}{50}\} = 0.95.$$

Now we have

$$\Pr\ \{\mu_x - \frac{\sigma_x}{50} < \hat{\mu} < \mu_x + \frac{\sigma_x}{50}\}$$

$$= \Pr\ \{\frac{-\sqrt{N}}{50} < \frac{\hat{\mu} - \mu_x}{\sigma_x/\sqrt{N}} \leq \frac{\sqrt{N}}{50}\},$$

and

$$\frac{\hat{\mu} - \mu_x}{\sigma_x/\sqrt{N}}$$

is distributed Nor (0, 1), which implies that one should choose t such that

$$\Phi(t) - \Phi(-t) = 0.95,$$

where $\sqrt{N}/50 = 1.96$ from tables of the normal distribution. This gives $N \doteq 10,000$.

In general, however, we do not know the variance of the distribution, so this also must be estimated. To do this, we use as our estimator

$$S^2 = \frac{1}{N} \Sigma (X_i - \hat{\mu})^2.$$

Assuming the distribution of X is approximately normal, NS^2/σ^2 is χ^2_{N-1}, so we can get rough estimates of the error involved [9].

Example: Suppose we wish to know with probability 0.95 that our estimate of σ^2 is correct within 5 per cent; that is, we wish to chose N such that

$$\Pr \{0.95\sigma^2 < S^2 \leq 1.05\sigma^2\} = 0.95.$$

But

$$\Pr \{0.95\sigma^2 < S^2 \leq 1.05\sigma^2\} = \Pr \{0.95N < \frac{NS^2}{\sigma^2} \leq 1.05N\}.$$

For large N, χ^2_{N-1} is approximately Nor (N − 1, 2(N − 1)), whence we have, upon some more manipulations,

$$\Pr \{0.95N < \frac{NS^2}{\sigma^2} \leq 1.05N\}$$

$$= \Pr\left\{\frac{1-0.05N}{\sqrt{2N-2}} < \frac{(NS^2/\sigma^2)-N+1}{\sqrt{2N-2}} \leq \frac{1+0.05N}{\sqrt{2N-2}}\right\} = 0.95,$$

where $\dfrac{(NS^2/\sigma^2)-N+1}{\sqrt{2N-1}}$ is approximately Nor (0, 1). Proceeding as before, we find

$$\frac{1+0.05N}{\sqrt{2N-2}} = 1.96,$$

which yields $N \doteq 3000$.

For the "criticality index," one is simply sampling from a binomial distribution. The estimator is the ratio of the number of sample realizations for which the arc is critical to the total sample size. The mean is NP and the variance $NP(1-P)$, where P is the probability of being critical. Since the number of samples is fixed, the estimator is simply a sum of independent random variables and therefore asymptotically normal.

Example: Suppose with probability 0.95 we want our estimator of **P** to be correct within 0.01. Now $(\hat{P}-P)/\sqrt{P(1-P)/N}$ is asymptotically Nor (0, 1), where P is our estimator. We want

$$\Pr\{|P-P| > 0.01\} = 0.95,$$

or

$$\Pr\left\{\frac{-0.01}{\sqrt{P(1-P)/N}} < \frac{P-P}{\sqrt{P(1-P)/N}} \leq \frac{+0.01}{\sqrt{P(1-P)/N}}\right\} = 0.95.$$

Since $P(1-P) \leq \frac{1}{4}$, if

$$\Pr\left\{\frac{0.02}{1\,N} < \frac{P-P}{\sqrt{P(1-P)/N}} \leq \frac{0.02}{N}\right\} = 0.95$$

the result will be even better. This implies

$$.02\sqrt{N} = 1.96, \quad \text{or} \quad N \doteq 10000.$$

Finally, we look at estimates of meeting specific schedule dates. For any given fixed schedule date, the problem is again sampling from a binomial distribution, and the previous analysis goes through. On the other hand, we may not want to limit our attention to any fixed date, and may want to know the accuracy for any date we choose; i.e., we may want to know, in some sense, how well our sample project duration distribution "fits" the real one. If we are willing to assume that the project duration cumulative distribution function is continuous, we can make probability statements about the greatest absolute differences between the sample cumulative distribution function (c.d.f.) and the true one <u>independent</u> of the distribution itself [9] . Kolmogorov in 1933 gave the asymptotic results that are tabulated, for example, in [6] .

<u>Example:</u> Let $D_n = \sup|F_n - F|$, where F is the c.d.f. of the parent population and F_n is the sample cumulative function. We find tabulated in tables Z, $L(Z)$, where

$$\lim_{n\to\infty} \Pr\{\sqrt{n}\,D_n \leq Z\}.$$

Suppose with probability 0.95 we want the maximum deviation to be less than 0.01. We find for $L(Z) = 0.9505$ that $Z = 1.36$. Solving for N in $N(.01) = 1.36$, we find $\sqrt{N} \doteq 18{,}500$.

The statistical analysis for other outputs of interest is carried through using the same methods found in this section.

6. COMPUTATIONAL SHORT CUTS

As in most Monte Carlo applications, there are techniques available for reducing computational expense in solving PERT problems. In solving PERT problems using Monte Carlo methods, most of the time is consumed in generating random numbers. Our approach will be to avoid sampling activity duration distributions that are rarely if ever on the critical path. These activities do not significantly affect the solution. The trick of course is to find these activities. Two methods will be outlined here. The first is an analytic approach, based on a suggestion of Ken MacCrimmon, and the second is a statistical approach.

Because the distributions of the activity durations have finite ranges, it may turn out that many paths can never be critical. One way to discover this is the following: Consider the realization obtained by setting all the activity durations at the lowest point of their ranges, and find the corresponding critical path P and path duration ℓ. Now set all activity durations <u>not</u> on the critical path at their highest value. Then only the activities used in paths with length strictly greater than ℓ can ever be critical. The remaining activities may be eliminated. References [7, 8] give algorithms for finding these paths. This procedure does not in general remove all the arcs that can never be critical; moreover, if the number of paths longer than ℓ is very large, then the algorithm for finding them can become quite unwieldy.

The statistical approach is quite a bit easier and can be applied to activity duration with infinite ranges. The approach here is to sample a relatively few times using all the activities and then stop and eliminate all activities that were never critical. For a given network the size of the initial sample determines the probability of making a mistake, that is of eliminating an arc that "should have" remained. For instance, if the initial sample were 1000, then the probability that a given activity with probability 0.01 of being critical would not be critical is $(0.99)^{1000} \doteq 7.3 \times 10^{-5}$. Another way to prevent the elimination of the wrong activities is not to remove the activities at all but to use each random sample for the activities with low criticality for K realizations before drawing a new one. In other words, draw new activity durations for these arcs only every K–th realization. The same estimators as before are respectively the unbiased estimators of the project duration distribution, its mean, and the criticality probabilities for the arcs. The estimate for the variance of the project duration is no longer unbiased but is consistent. The variance is not decreased but as a matter of fact is very slightly increased. The savings are obtained by having to sample the "unimportant" activities only $(1/k)$–th the number of times that the more important activities are sampled. Under this approach it is not necessary to separate only those activities that were never critical in the initial sample, but more generally all those activities that were on the critical path less than some given number of realizations could be distinguished.

Finally, consider a situation in which the exact distribution of the project–duration distribution is not desired, but only its mean and variance. In this case an activity which is <u>always</u> critical can be replaced by a deterministic activity with duration equal to the mean of the original distribution. The same procedure is gone through as before, except that the original variances of all the changed activities must be added to the estimated project–duration variance.

7. FEASIBILITY OF METHOD AND PROPOSALS FOR FUTURE INVESTIGATION

An experimental computer code has been coded by R. J. Clasen for the IBM 7090 computer to test the feasibility of Monte Carlo evaluation of PERT networks. Runs of 10,000 samples and up to 45 activities were run. The code can handle 1000 activities and can generate uniform, triangular, and beta duration-distribution functions. The output is the project-duration density function, and its mean and variance, and the criticality index for each of the activities. This code is not suitable for general use because it does not include many of the accounting and input-output functions of production systems; also much output of more practical than theoretical interest is not computed. To use this method for production runs, the next step is to incorporate the solution method into a current PERT system, using its accounting and input-output routines. Then "authentic" networks should be evaluated and the results compared with those obtained using the old program to see if the hoped-for advantages materialize.

Experience with the program has indicated that most of the time is consumed generating random numbers, so that the running time is essentially a linear function of the number of random numbers generated—or, for fixed sample size, a linear function of the number of activities in the network. This is especially true for larger-sized networks. For a 200-activity network and 10,000 samples, the running time would be about 20 minutes for triangular distributions and about 5 minutes for uniform distributions.

APPENDIX

Figure 1 shows an example given in [1] p. V–2.

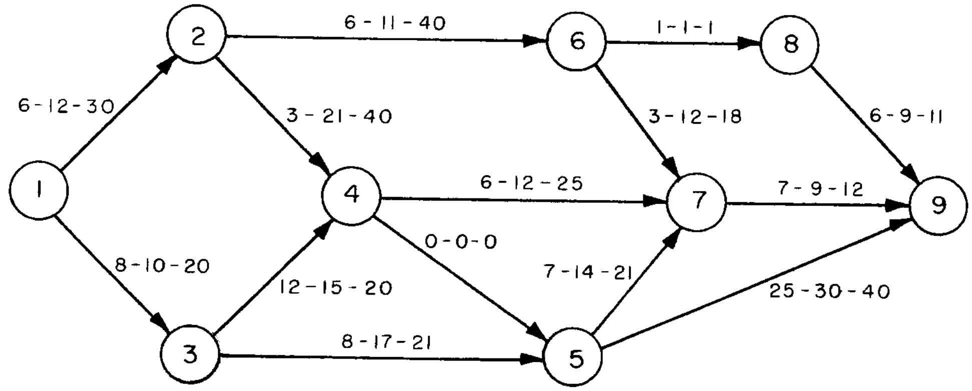

Fig. 1

The three numbers assigned to each activity are respectively the optimistic, most likely, and pessimistic estimates for each activity.

The expected project duration time using PERT is 66.0 with a variance of 60.27. The Monte Carlo method using 10,000 realizations yields a mean of 67.0 and a variance of 42.39.

The distributions used for the activity durations were the beta distributions with end points and modes given by the three parameters indicated in Fig. 1. The standard deviation was taken to be 1/6 the range.

In Fig. 2 a normal distribution with mean 66.0 and variance 60.27 is compared with the distribution obtained by Monte Carlo. Figure 3 displays the network with the criticality index for each activity. The heavy line is "the critical path" calculated using expected values for the activity durations.

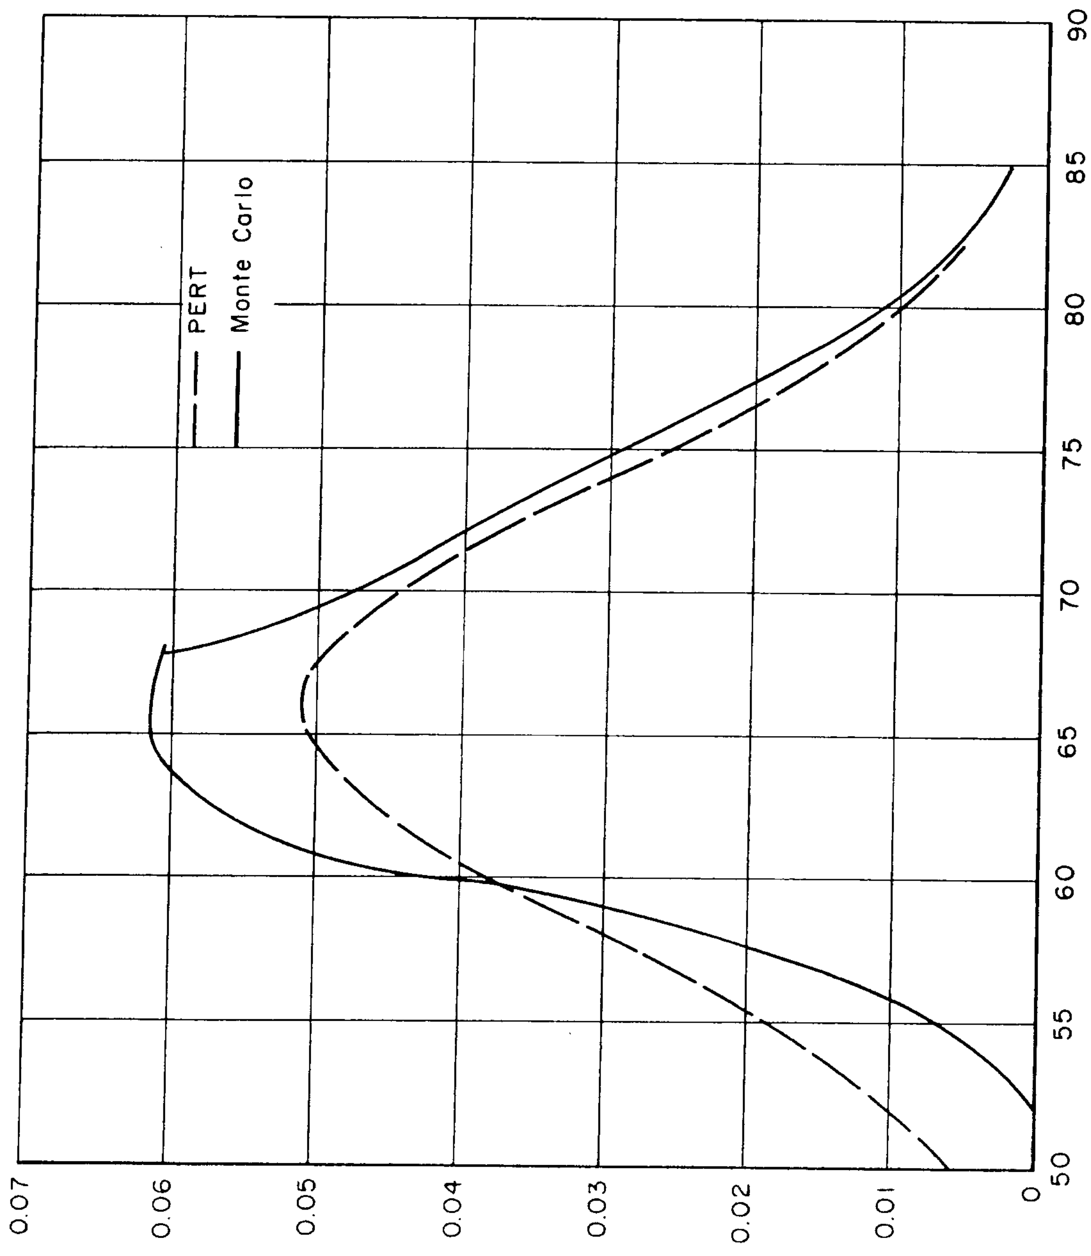

Fig. 2 — Probability density curves

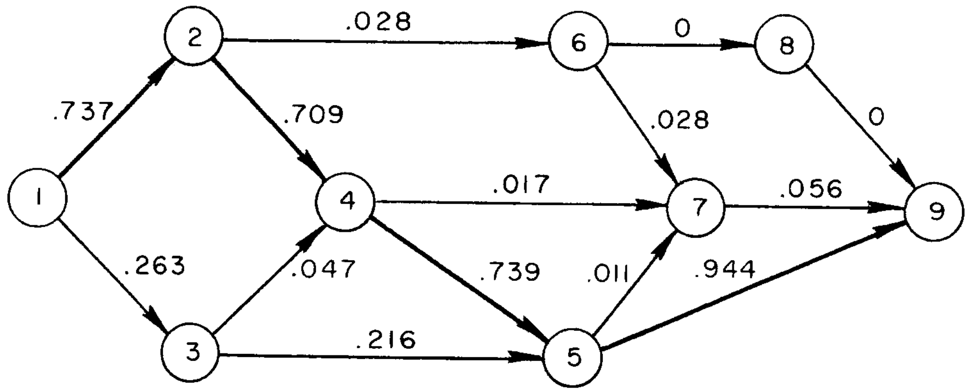

Fig. 3

Carrying through the analysis of Sec. 5 we see that with $N = 10,000$ we can expect with probability 0.95 that our estimated mean is within $\sigma_{x/50}$ of the true mean. Using as an approximation of σ_x the square root of the estimated variance, we obtain

$$\sigma_{x/50} \doteq \sqrt{42}/50 = .13 .$$

Since the Monte Carlo estimate differs by 1.0 from the PERT estimate the difference is certainly statistically significant at the 95% level. Finding the shortest and longest possible project durations by using for activity durations respectively the optimistic and pessimistic times we find that the discrepancy, 1, is on the order of 1 1/2 per cent of the range. On the other hand for the variance of the distribution the variance estimated by PERT methods is

1.42 times as large as the Monte Carlo estimate, but as we saw in Sec. 5 the probability is better than 0.95 that the Monte Carlo estimate is within 5 per cent of the correct value. Finally, in Fig. 4 the probability that the project will be completed by X is charted as estimated by the PERT technique and as estimated by the Monte Carlo approach. With $N = 10,000$ in $N(D_n) = 1.36$, we see that with probability 0.95 the sample cumulative distribution is within 0.0136 of the true distribution at every point.

We also note that activities 6–8 and 8–9 were never critical. In fact they cannot be critical, but the deterministic algorithm given in Sec. 6 will not reveal this. So in this case the statistical approach is clearly preferable.

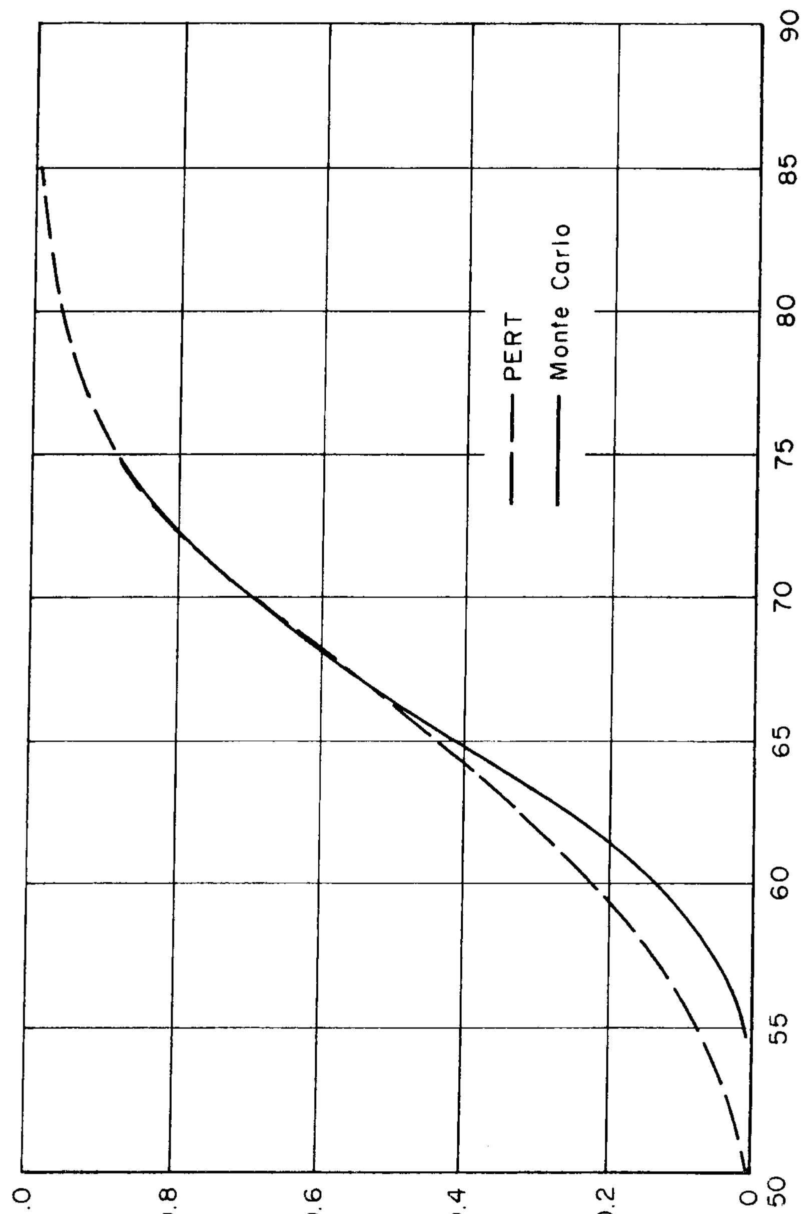

Fig. 4 — Cumulative distributions

REFERENCES

1. Air Force Systems Command, "PERT–AFSC Policies and Procedures Handbook," ASD Exhibit AS00 61–1.

2. Clark, C. E., "The Greatest of a Finite Set of Random Variables," ORSA, Vol. 9, No. 2, March–April, 1961, p. 145.

3. Kelley, J. E. and M. R. Walker, "Critical Path Planning and Scheduling," '959 Proceedings E. J. C. C.

4. Madansky, A., "Inequalities for Stochastic Linear Programming Problems," Management Sc., Vol. 6, No. 2., January 1960, p. 197.

5. Malcom, D. G., J. H. Roseboom, C. E. Clark, and W. Fazar, "Application of a Technique for Research and Development Program Evaluation," ORSA, Vol. 7, No. 5, Sept.–Oct., 1957, p. 646.

6. Owen, D. B., Handbook of Statistical Tables, Addison Wesley, 1962.

7. Pollock, M., "Solutions of the k–th Best Route Through a Network – A Review," J. Math. Analysis and Appl., Vol. 3, 1961, p. 547.

8. Suurballe, J. W., "Network Algorithms for Combinatorial and Discrete Variable Optimization Problems," to appear in Proceedings of a Symposium on Mathematical Programming, Chicago, 1962.

9. Wilks, S. S., Mathematical Statistics, John Wiley and Sons, 1962.