COMPARISON OF CONSTRUCTION ALTERNATIVES USING MATCHED SIMULATION EXPERIMENTS

By Photios G. Ioannou,[1] Associate Member, ASCE, and Julio C. Martinez,[2]

ABSTRACT: The comparison of alternative construction methods is one of the principal reasons for using simulation to model construction processes. The efficiency and effectiveness of such comparisons can be greatly improved by the prudent use of "matched pairs," a variance reduction technique based on dedicated and fully synchronized random number streams. The basic methodology is illustrated by using the STROBOSCOPE simulation system to compare two alternative construction methods for rock tunneling [Conventional versus the New Austrian Tunneling Method (NATM)]. For this example the effects are dramatic. The probability of identifying and choosing the cheaper construction method based on a single run increases from 55% to 96%, the variance of the cost difference decreases by two orders of magnitude, and the 95% confidence interval for the true cost difference given by 4,000 independent runs can be obtained by performing only seven replications using matched pairs. Besides this improvement in statistical efficiency, the use of matched pairs is a necessity for this example in order to compare the alternatives on a logical and equitable basis.

INTRODUCTION

A primary objective in using simulation to model construction processes is to evaluate and compare the performance of alternative construction methods to select the best one. A common mode of operation is to construct a simulation model for each method, conduct a limited number of simulation experiments (runs), and then compare the competing alternatives based on the resulting average measure of their performance.

The problem with this approach is that the simulation runs for one construction method are often designed to be statistically independent from those for the competing alternatives. Thus, if the principal causes of uncertainty in the system are not directly due to the methods being compared, but are caused by factors external to the methods and inherent to the problem being analyzed, the alternatives will not be compared on an equal basis.

The tunneling example presented below is an excellent illustration of this problem. In tunneling, most of the risk is due to geologic uncertainty, which is independent of the chosen construction method. As a result, it is very important that competing construction alternatives be compared under the same geologic conditions. Otherwise, the observed cost differences will be primarily due to differences in project geology rather than the construction methods themselves.

To illustrate this point let us consider a tunneling example where all uncertainty about the cost of construction alternatives A and B is due to the project's unknown geologic conditions. Given any particular geologic profile, the costs associated with the two construction methods can be computed deterministically. Let us also assume that these costs are very sensitive to the project geology. The cost of each method is very low if the geology is "favorable" (or "good"), and very high if the geology is "unfavorable" (or "bad"). Furthermore, let us assume that the cost of method A for any geologic profile (good or bad) is always less than (but close to) that of method B for the same geologic profile. Thus, alternative A should always be preferred to alternative B for any geologic conditions that may be encountered during construction.

[1] Assoc. Prof., Civ. and Envir. Engrg. Dept., Univ. of Michigan, Ann Arbor, MI 48109-2125.

[2] Ph.D. Candidate, Civ. and Envir. Engrg. Dept., Univ. of Michigan, Ann Arbor, MI 48109-2125.

Note. Discussion open until February 1, 1997. To extend the closing date one month, a written request must be filed with the ASCE Manager of Journals. The manuscript for this paper was submitted for review and possible publication on July 11, 1995. This paper is part of the ***Journal of Construction Engineering and Management***, Vol. 122, No. 3, September, 1996. ©ASCE, ISSN 0733-9364/96/0003-0231-0241/$4.00 + $.50 per page. Paper No. 11112.

If the costs of using method A or B are obtained by performing independent simulation runs, the results would not always lead to this obvious conclusion. Independent simulation runs imply that the computed cost for method A (in simulation run *k*) is based on a geologic profile that is different from the one used to calculate the cost of method B (in the same simulation run). Given these assumptions, if simulation runs for the two alternatives are compared in pairs, close to (but less than) half of the simulation runs should indicate that method B is indeed less expensive than method A. The reason is that for independent simulation runs, there is a 50% chance that the simulated geologic profile for method A is "worse" than the simulated geologic profile for method B. Thus, a single simulation run for method B may appear to be less expensive than a single run for method A simply because the simulated tunnel geology for B happened to be more favorable.

VARIANCE REDUCTION TECHNIQUES—MATCHED PAIRS

Conceptually, the solution to the foregoing problem is quite simple. The simulation runs for each alternative must be designed so that uncertainty impacts each construction method in a similar manner. Since all uncertainty in a simulation model is determined by random numbers, the key is to ensure that the random numbers used by each alternative follow similar patterns.

A very simple way to do this is to start the corresponding simulation replications for all alternative methods using the same set of random number seeds. Thus, the random number seed that is used to start the *i*th replication (simulation run) for alternative A should also be used to start the *i*th replication for alternative B.

Although this approach is quite effective and should certainly be used whenever possible, it does not always work well by itself and may even backfire. The exact sequence in which random numbers are used during a simulation run for construction method A is not necessarily the same as that for method B.

Let us consider as an example, an earthmoving operation that involves a load activity followed by a haul activity. Alternative A uses one "big loader" and method B uses two "small loaders," each loading separate trucks that when full haul the dirt away. For method A the first random number would be used to determine the duration of the first instance of the load activity while the second random number would be used to determine the duration of the first instance of haul. For method B, the availability of two loaders (each loading a separate truck at the same time) implies that the first and the second random numbers would be used to determine the duration of the first and the second instances of activity load. The third random number for method A is used to determine the second instance of load, whereas for model B it is used to determine the first instance of haul. It is clear that even if the simulation runs for methods A and B are started with the same random number seed, the two runs will use the resulting sequence of random numbers for totally different purposes (load, haul, etc.). For a tunneling project, the random number used to model a transition in geology for construction method A, may end up determining an excavation duration for alternative B.

The sequence of random numbers produced by a given random number generator (RNG) is fixed and periodic. This means that these random numbers are like the links in a circular chain. They follow each other in a fixed order and in the end this sequence is repeated (no matter which starting point was chosen). The seed used to initialize the RNG for a particular simulation run simply determines the entry point to this sequence. Random number streams are successive nonoverlapping segments of this chain (produced by the same RNG).

Modern random number generators provide a mechanism for determining directly the 100,000th consecutive random number from any starting point in this chain of random numbers. Thus, the typical length of a stream is also 100,000. For any given seed, the first 100,000 random numbers are usually called stream 0, the second 100,000 are stream 1, the third 100,000 are stream 2, etc. Changing the seed for stream 0 changes the starting point where stream 0 begins (within the fixed sequence of a specific RNG) and thus changes the random numbers for all other streams as well. A 32-bit RNG with a cycle of over 2 billion, can have more than 20,000 independent random number streams, each having a length of 100,000 numbers. Increasing the spacing between streams (by varying the length of each stream to integer multiples of 100,000) allows for fewer but longer streams of independent random numbers.

Thus, a simple and effective approach to synchronize the use of random numbers across construction methods is to dedicate a stream of random numbers to each uncertain variable that is common to all alternatives. For example, stream 1 may be dedicated to determining the duration of the load activity, while stream 2 is dedicated to the haul activity (for all construction methods).

For the dedicated use of random number streams to be effective, no single variable in any simulation run should use more random numbers than the stream length (typically 100,000). Otherwise, the random numbers used by one variable (e.g., to determine the duration of activity load) will be the same as those for the variable using the next stream (e.g., to determine the duration of activity haul) and the two will be positively correlated. This can be avoided by explicitly increasing the length of each stream to an appropriate multiple of 100,000 (if this capability is provided by the simulation system), or by deliberately skipping streams. For example, using only streams 0, 2, 4, etc., allows for independent streams with a length of 200,000.

Similar care must be taken to select the seeds that initialize successive simulation runs for the same construction method so that these runs are indeed independent. These seeds must be chosen so that all runs for the same construction alternative do not in any way reuse the same random number sequences. This can be accomplished in one of two ways.

The first is to use the last seed value produced by each stream as the seed for the same stream in the next replication for the same construction alternative (i.e., within each stream, the next run picks up where the previous run left off). This requires that the entire set of replications does not use more random numbers than the length of a stream (i.e., it must not exceed the standard stream length of 100,000 numbers, or a specified multiple).

Another option is to use the last random number produced by the highest numbered stream (e.g., stream 3 if a total of four streams are being used) as the seed for stream 0 in the next simulation run. This means that stream 0 in the next run picks up where stream 3 in the previous run left off. This approach requires (roughly) that the total number of replications for each construction alternative, multiplied by the number of streams in each run, does not exceed the total number of independent streams provided by the RNG.

The most appropriate of these two approaches depends on how many random numbers per stream are used in each simulation run, the number of streams in the simulation model, and the total number of replications to be performed. A "back-of-the-envelope" calculation can usually indicate how many random numbers are used per stream per simulation run, and this can be used to select an appropriate stream length, as well as a method to select the seed for the next run.

Another issue that must be considered is how many random numbers are used to produce one sample of a non-uniformly distributed random variable. To keep the random number streams completely synchronized across simulation runs for alternative methods, the set of random numbers used to generate one such sample must be the same for all construction alternatives. For example, the excavation duration for alternative A may be Normal (which may require two random numbers per sample) whereas the excavation duration for alternative B may be exponential (which requires one random number per sample). To ensure the synchronized use of random numbers by the two alternatives, the methods used to generate random samples must adhere to the following guidelines:

1. The number of random numbers used to generate one sample for a nonuniformly distributed random variable must be constant. This means that "acceptance-rejection" techniques for generating samples must be avoided because they use a variable number of random numbers per sample.
2. If possible, competing alternatives should use the same distribution to model the same activity (by using different parameters). For example, the load activity for all alternatives should follow a Normal distribution but with (possibly) a different mean and variance.
3. If the various alternatives cannot use the same distribution to model a particular variable, they should at least use distributions that require the same number of random numbers per sample. For example, if the duration of the load activity for a certain alternative cannot be represented by a Normal distribution (which uses two random numbers per sample), then it should at least be modeled by another distribution (e.g., a Gamma) that also uses two random numbers per sample.
4. The transformation from random numbers into samples must be consistent and monotonic. This means, for example, that "large" random numbers should always produce large samples and "small" random numbers should always produce small samples (or vice versa). This is required in order to induce positive correlation between the simulation runs for the various construction alternatives and thus reduce the variance of their difference in performance (as explained later).
5. If none of the foregoing are possible, then the problem of keeping random number streams synchronized, as well as other related issues, can be avoided by discarding a certain number of random numbers at appropriate points during a simulation run.

Collectively, all the preceding recommendations are variance reduction techniques that rely on "blocking" (i.e., "comparing like with like"). In the following example they are referred to as "matched pairs" to indicate that a simulation run for one construction method is "matched" in terms of random numbers to that for the other alternatives. A comprehensive review of variance reduction techniques can be found in Law and Kelton (1991).

The purpose of using matched pairs as a variance reduction technique for the comparison of alternatives can be explained more clearly by using elementary statistics. Let us define X to be the cost of one construction alternative, Y to be the cost of another, and $Z = Y - X$ to be their difference in cost. The decision criterion of interest is whether the true expected value of Z, E[Z] = E[$Y - X$] = E[Y] – E[X], is positive or negative. Since the actual E[Z] is unknown, it must be estimated by using the results of simulation experiments. The maximum likelihood estimator of E[Z] is $\overline{Z}(n)$, the average of the n samples $Z_i = Y_i - X_i$ produced by n pairs of simulation runs. The samples X_i given by consecutive simulation runs are (by design) independent and identically distributed (IID) with the same mean E[X] and variance Var[X]. The same is true for the samples Y_i. As a result, the samples Z_i are also IID and the variance of the average $\overline{Z}(n)$ is given by

$$\mathrm{Var}[\overline{Z}(n)] = \frac{1}{n}\mathrm{Var}[Z_i] = \frac{1}{n}\left(\mathrm{Var}[X] + \mathrm{Var}[Y] - 2\rho_{X_iY_i}\,\mathrm{SD}[X]\mathrm{SD}[Y]\right) \quad (1)$$

where Var = variance; SD = standard deviation; and ρ_{XiYi} = correlation coefficient between X_i and Y_i in the ith pair of simulation runs. From (1) it is clear that for a given sample size n (i.e., number of simulation runs), the greatest reduction in variance for the average difference in cost can be achieved when the correlation coefficient ρ_{XiYi} approaches +1. Thus, the purpose of matched pairs, from a mathematical perspective, is to make ρ_{XiYi} as close to +1 as possible. This minimizes the variance of the average, $\mathrm{Var}[\overline{Z}(n)]$, and produces the tightest possible confidence intervals for the true E[Z] for any given sample size n. Matched pairs can be quite effective in this respect as illustrated by the following example.

STROBOSCOPE SIMULATION SYSTEM

The effect of using matched pairs in simulation experiments is illustrated by the following example that compares two alternative tunnel construction methods. All simulation runs have been performed using STROBOSCOPE (STate and ResOurce Based Simulation of COnstruction ProcEsses). STROBOSCOPE is a general-purpose simulation programming language based on the activity-scanning paradigm that is particularly suited to modeling construction operations. Simulation models written in this language are processed by a simulation engine that can receive input either from an integrated development environment (edit-compile-run), or from a graphical user interface that automatically generates the code for a model.

A STROBOSCOPE model consists of a network of interconnected links and nodes and a series of programming statements that define its behavior, and statements that control the simulation. The STROBOSCOPE language can set the random number seed for any stream, have any number of random number streams (spaced at any specified multiple of 100,000 random numbers), can perform any number of simulation runs (replications) within each simulation model (within the limits of a 32-bit RNG), can switch simulation models within a replication, can define any number of normal or weighted statistics collectors, and can selectively clear some of the simulation statistics while keeping others across replications. Furthermore, it can produce any output, in any desired format, and save it to a file so that it may be analyzed by another program.

A description of STROBOSCOPE and its applications can be found in Martinez et al. (1994), Martinez and Ioannou (1994), Ioannou and Martinez (1995), Martinez and Ioannou (1995), and Martinez (1996). The STROBOSCOPE program and its documentation are available from grader.engin.umich.edu via anonymous ftp (the file transfer protocol used on the Internet).

EXAMPLE—TUNNEL CONSTRUCTION

We shall use as an example the construction of a two-track tunnel between two existing stations 1.6 km apart. The tunnel is assumed to serve as a connector for two commuter rail systems and thus requires no intermediate stations. All excavation must be performed underground with no intermediate shafts, starting at one of the existing stations. Two construction alternatives will be evaluated, the conventional method and the New Austrian Tunneling Method (NATM). Both options use the same excavation method (drill and blast) and muck method (train muck haulage). Their difference lies in the system used for initial support.

For the conventional method, initial support for the roof and sides of the tunnel is provided by steel sets and lagging. Steel sets are H-beams bent to conform to the rough shape of the excavated tunnel cross section. Wood lagging consists of heavy timber, placed between the steel sets and the rock, to support the walls and roof of the tunnel until the final lining is constructed. The final lining is typically a shell of cast-in-place reinforced concrete.

The NATM is a special case of the observational or adaptable approach, where the initial tunnel support requirements are not determined during the design phase, but rather during the construction phase by monitoring the rock deformation. This deformation is held to within a specified tolerance by varying the quantities of the support materials: rock bolts, wire mesh, and shotcrete. Proponents of observational techniques, such as NATM, point to material and labor cost savings by only using the exact amount of initial support required, rather than using an initial support designed for the suspected worst condition in the tunnel. Opponents of observational techniques cite the time to make support decisions and the imposed time variations (from varying support requirements) during the construction cycle as having major impacts on total cost. This example ignores the institutional barriers and implied liability transfers from adopting the NATM approach.

For both construction alternatives, the "excavation" activity includes drilling holes into the tunnel face and loading them with explosives. This is followed by "shooting" the rock (retracting the jumbo, wiring the "pig-tails", and detonating the explosives). For simplicity, these times have been made part of the duration of the excavation activity and thus shooting is assumed to take zero time. Shooting is followed by the "smoke" activity that ventilates all the smoke out of the tunnel. To bring back the jumbo and resume drilling, all the muck resulting from the last shot must be removed. When mucking is done, excavation (i.e., drilling and loading) and the installation of the initial rock support can start. Thus, excavation and support can occur at the same time. In order to shoot again, the excavation for the next round (drilling and loading of holes) must be complete, and enough initial support must have been installed so that, after the rock is shot, the length of unsupported tunnel is less than the maximum allowed for the current rock class. After the rock is shot, the cycle repeats again. The excavation progress in one cycle (drill, blast, muck) is called a "round."

TUNNEL CONSTRUCTION DATA

The tunnel geology is modeled as a discrete-state, discrete-space Markov process (Ioannou 1987, 1989a). The rock conditions (i.e., system states, from best to worst) are represented by three "ground classes" (1 "good", 2 "medium", 3 "poor"). Space is discretized into "steps" (linear segments along the tunnel's horizontal alignment). Ground classes persist for at least the length of one step. The first step starts out in ground class 1. The ground class transition probabilities (from step to step) are shown in Table 1.

Excavation for both the conventional and the NATM methods is accomplished by drilling and blasting the rock. The pertinent data for both methods are shown in Table 2. The excavation advance rates (linear meters per hour) shown in this table have triangular probability distributions defined by (min/mode/max).

The type of initial tunnel support is the fundamental difference between the conventional method and the NATM. The conventional method uses steel sets and lagging that are designed to withstand the worst possible rock conditions (ground class 3). Thus, the initial support for the conventional method is the same for all ground classes. The NATM uses rock bolts, wire mesh, and shotcrete, the amounts of which vary depending on the current ground class.

Table 3 shows the cost of the support for each type and method, as well as the rate at which it can be installed. The support placement rates (linear meters per hour) in this table have uniform probability distributions in the range (min/max).

Table 4 shows the cost and time delay associated with the decision to switch to a new ground class (and thus to a new support type) at the end of each step when using NATM.

Table 5 shows the rest of the data needed to construct the simulation models.

TABLE 1. State Transition Probability Matrix (from Step to Step)

From ground class (1)	To Ground Class 1 (Good) (2)	2 (Medium) (3)	3 (Poor) (4)
1 (Good)	0.60	0.25	0.15
2 (Medium)	0.10	0.80	0.10
3 (Poor)	0.05	0.20	0.75

TABLE 2. Excavation Data for Conventional Method and NATM

Drill & Blast Excavation Data (1)	Ground Class 1 (Good) (2)	2 (Medium) (3)	3 (Poor) (4)
Advance rate (linear meters/h) (min/mode/max)	0.37/0.38/0.43	0.32/0.33/0.40	0.13/0.17/0.32
Operating cost ($/h), including labor	2,019	1,760	1,750
Overbreak volume (percent of desired excavation)	10%	15%	30%
Maximum distance of support to the tunnel face (m)	16	8	4

TABLE 3. Initial Support Cost ($/m) and Placement Rates (linear m/h)

Ground class (1)	Support Cost: Conventional (2)	Support Cost: NATM (3)	Support Placement Rate (min/max): Conventional (4)	Support Placement Rate (min/max): NATM (5)
1	1,400	940	0.35/0.45	0.55/0.65
2	1,400	1,160	0.35/0.45	0.37/0.47
3	1,400	1,350	0.35/0.45	0.15/0.30

TABLE 4. Time and Cost to Decide/Switch Support Type (NATM Only)

From ground class (1)	TIME (h) To Ground Class 1 (2)	2 (3)	3 (4)	COST ($) To Ground Class 1 (5)	2 (6)	3 (7)
1	0	8	12	0	2,000	5,000
2	8	0	4	1,600	0	3,000
3	12	4	0	3,800	2,300	0

TABLE 5. Miscellaneous Data

Item (1)	Value (2)
Tunnel length	1,600 m
Step length	100 m
Round length	4 m
Finished (inside) tunnel diameter	6 m
"B-line" (excavation) diameter	7 m
Initial capital cost (for either method)	$400,000
Salvage value as a percentage(for either method)	30%
"Smoke" time after every blast	30 min
Muck rate for all options	22 m^3/h
Mucking costs for all options	$10/$m^3$
Concrete unit cost (including labor)	$70/$m^3$
Overhead cost per day (24 h)	$12,000

SIMULATION MODEL

The network model for both the conventional method and NATM is illustrated in Fig. 1. This network, as shown, comes directly from STROBOSCOPE's graphical user interface. At an abstract level, the queues and the combi and normal activities in a STROBOSCOPE network are similar in appearance and function to those in CYCLONE (Halpin and Riggs 1992; Ioannou 1989). In STROBOSCOPE, however, a network defines only the basic model structure. The underlying logic is determined by programming statements that define resources, attributes for the network elements, and other modeling entities (when using the graphical user interface these are entered into appropriate dialog boxes and are not shown in Fig. 1). Some of these are outlined as follows.

Links in STROBOSCOPE networks must be given unique names. These names are used in STROBOSCOPE programming statements to define a variety of link attributes. By convention, link names use two letters (that abbreviate the type of resource that flows through) followed by a number. In this model, RD is the abbreviation for the "generic" resource type Round; SR indicates the "characterized" resource type SupportRound.

The basic cycle of Excavate (drill holes and load explosives), Shoot, Smoke, and Muck is shown by the sequence of links RD1 to RD6. The queue RoundToDo is initialized with one resource of the generic type Round, which cycles through this series of links. Thus, Excavate is the first activity to start. Notice that Fig. 1 shows only one cycle RD1 to RD6, and this cycle is not specific to any particular ground class. In contrast, an equivalent CYCLONE network would require the explicit definition of three such cycles, one for each ground class. Furthermore, Fig. 1 does not include an elaborate mechanism of nodes and links to model transitions in the tunnel geology at the end of each step (i.e., similar to chance nodes in a decision tree). This simplicity in the network representation is made possible by STROBOSCOPE's programmability and the model's parametric design.

The current ground class associated with the basic cycle RD1 to RD6 is kept by a "SaveValue" called CurGClass. A STROBOSCOPE SaveValue is a global storage location similar to a variable in a conventional programming language (such as

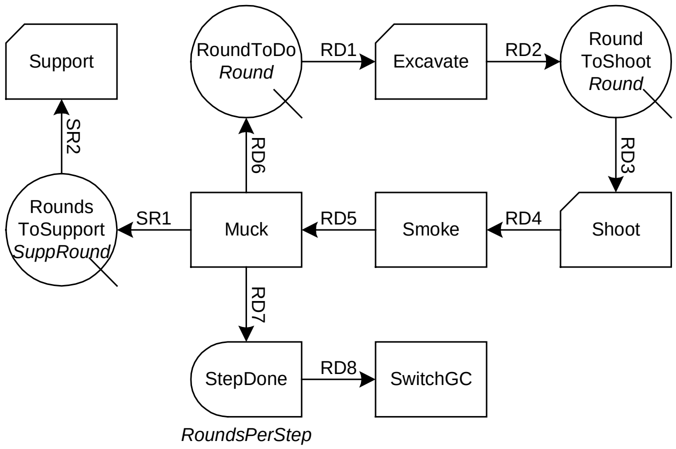

FIG. 1. STROBOSCOPE Network for Tunnel Construction

C). The value of CurGClass is a number that corresponds to the current ground class (1, 2, 3). As explained in the following section, this value changes every time a transition to a new ground class occurs at the end of activity SwitchGC.

The data in Tables 1–4 have a tabular format and can be stored efficiently in one- or two-dimensional arrays. Since these data also depend on the current ground class, they can be easily accessed using the number stored in CurGClass as an array index. For example, to sample the duration of the Excavate activity the model uses CurGClass to access the appropriate excavation advance rate parameters (i.e., min/mode/max for a triangular distribution), depending on the current ground class as shown in Table 2. Thus, the flexibility of STROBOSCOPE's programming language allows modeling of the excavation-blast-smoke-muck sequence, for any ground class, as one generic cycle (RD1 to RD6) that is controlled by one parameter (CurGClass).

The node StepDone is a "consolidator". It receives one Round every time Muck finishes and stores it until a total of 25 Rounds are accumulated. When this happens, StepDone finishes and passes all 25 Rounds to activity SwitchGC, which is then allowed to start. Thus, activity SwitchGC occurs only at the end of each 100-m step.

When SwitchGC starts, the model determines the ground class for the next 100-m step. The entire mechanism for switching ground classes is handled by one STROBOSCOPE programming statement. This statement uses CurGClass as an array index to access the appropriate probabilities in the transition probability matrix (Table 1) and performs Monte Carlo sampling to select the next ground class. The next ground class is then stored in SaveValue NextGClass. At this point, the ground class for the ending step is in CurGClass whereas that for the upcoming step is in NextGClass. For NATM, the duration of the SwitchGC activity and the associated cost for changing the support are given by Table 5. These data are accessed by using CurGClass and NextGClass as indices to a two-dimensional array. For the conventional method the duration and cost of SwitchGC are set to zero. At the end of the SwitchGC activity, CurGClass is set to the value of NextGClass and the ground class transition is complete.

Every time the normal activity Muck finishes, a new round (4 m) of unsupported tunnel becomes available to the Support activity. However, Fig. 1 shows that the queue RoundsToSupport does not hold resources of the generic type Round, but holds "characterized" resources of type SupportRound. There are two reasons for this change in resource type.

The first is that the time it will take to install the initial support depends on the specific ground class to be supported (not the ground class currently being excavated). The ground class for the Support activity is not necessarily the same as CurGClass (which is the current class for Excavate, Shoot, and Muck). This disparity occurs whenever there is a transition to a different ground class at the end of one step and the beginning of the next. Thus, each length of tunnel waiting to be supported must be "tagged" with the ground class it belongs to (so that the duration of the Support activity may be determined accordingly). This requires the use of a characterized resource type, because, in STROBOSCOPE, characterized resources can have dynamically assignable properties called "SaveProps." In this case, one such SaveProp for SupportRound is set equal to its ground class and is used to determine the duration of the corresponding Support activity.

The second reason is to allow the Support activity to lag behind the Excavate, Shoot, and Muck cycle. For this to happen, the Support activity must be slower than the Excavate activity (since the two activities start at the same time, i.e., at the end of the Muck activity), and the Shoot activity must be allowed to start even though the Support activity has not yet finished a full 4-m round (whereas the Excavate activity has). Since it is impossible to blast the rock safely while installation of the support is going on, it is necessary for the Shoot activity to preempt (stop) the Support activity (before it completes a 4-m round). Otherwise, Shoot will have to wait for Support to finish, and blasting will never happen unless the entire excavated tunnel is fully supported up to the tunnel face. Thus, to allow Shoot to preempt the Support activity, each 4-m Round must be subdivided into n "shorter" resources of type SupportRound (so that Support may be stopped at the end of each of the n SupportRounds). In addition, the Shoot activity cannot start (and thus it cannot preempt Support) unless the distance of the installed support to the tunnel face is less than a certain distance. This distance is the smallest of the two distances allowed by the ground class being excavated and the ground class being supported.

Although the issue of subdividing one 4-m Round into n SupportRounds is interesting (from the point of view of modeling accuracy), it is not of any particular consequence for the foregoing example data. The duration of the Excavate activity and the duration of the Support activity happen to be very close. As a result, very rarely does the Shoot activity have to wait for the Support activity to finish before it can start (even for $n = 1$). Experiments using values of n from 1 to 20 produced almost identical results. The results that follow correspond to $n = 1$.

DESIGN AND EXECUTION OF SIMULATION EXPERIMENTS

This tunneling example requires only three independent random number streams. Stream 0 is used to sample from the triangular distribution for the excavation advance rate to determine the duration of the Excavate activity (one sample per 4-m round). Similarly, stream 1 is used to sample from the uniform distribution for the support placement rate to determine the duration of the Support activity (one sample per 4-m round). Stream 2 (the most important for this problem) is used to perform Monte Carlo sampling based on the cumulative form of the transition probability matrix shown in Table 1 (one sample per 100-m step).

A single simulation run for the entire 1,600-m tunnel uses a total of 815 random numbers: 400 each from streams 0 and 1, and 15 from stream 2. The length of each stream is 100,000 (the minimum possible). The seed for stream 0 in the next simulation run is set to the value returned by stream 2 in the previous replication. Thus, the seed for the second simulation run corresponds to the 200,015th random number produced by the seed of the first replication. This strategy ensures that the simulation runs for the same construction alternative are mutually independent.

Each of the two construction alternatives is associated with two sets of simulation runs. Each set consists of 4,000 mutually independent replications for one construction alternative. This exceptionally large number of simulation runs was chosen so that some of the observations that follow would be self-evident (typical simulation experiments rarely exceed 100 replications).

The replications in each set traversed a total of 4,000 × 200,015 = 800,060,000 random numbers (well within the 2 billion limit of a 32-bit RNG). Out of this large set of numbers, only 4,000 × 815 = 3,260,000 were actually used to produce random samples. The rest were skipped over by the mechanism that selects the seed for the next replication (based on the last value returned by stream 2). This ensured that random numbers were used only once and that all 4,000 replications for each construction alternative can be considered mutually independent.

The two sets of 4,000 replications per alternative form two groups of 4,000 paired simulation runs. Together, the first set for the conventional method and the first set for NATM form the 4,000 independent pairs. The second set for the conventional method and the second set for NATM form the 4,000 matched pairs.

The replications for the independent pairs were performed by starting the *i*th replication for NATM where the *i*th replication for the conventional method stopped. For example, the seed for stream 0 for the *i*th replication of NATM is the one that would have produced the 401 random number of stream 0 for the *i*th replication of the conventional method. The same applies to streams 1 and 2 as well. Thus, these 4,000 NATM runs are totally independent from the 4,000 simulation runs for the conventional method.

For the matched pairs, the same seeds that were used for the 4,000 replications of the conventional method were also used for the corresponding 4,000 replications of NATM. The term matched pairs indicates that the *i*th NATM run is matched to the corresponding *i*th run for the conventional method (both runs, for example, are based on the same geology). In this case, the *i*th pair of simulation runs for the two construction alternatives are initialized with the same seed, use the same random number streams, produce the same geologic profile, and have the same excavation durations and costs. The only differences between the two runs in a matched pair are the time to install the support (same random numbers but different distribution parameters), the cost of the support (different deterministic cost), and the cost of overhead (same daily rate but different project durations).

To compare the results produced by independent pairs on an equal footing with those for the matched pairs, the 4,000 simulation runs for the conventional method were designed to be exactly the same in both cases. Only the NATM runs are different. The 4,000 independent pairs and the 4,000 matched pairs form the basis for the statistical analysis that follows.

SIMULATION RESULTS

The results of all simulation runs are shown in Fig. 2 as 8,000 points of the form (conventional method cost, NATM cost). The 4,000 "hollow" points correspond to the independent pairs and the 4,000 "filled" points to the matched pairs.

Fig. 2 illustrates the wide spectrum of construction costs that are possible for this example. These costs range approximately from a low value of $16M to high value of $26M. The most striking aspect of Fig. 2 is that the independent pairs can fall anywhere within a square bounded by these values, whereas the matched pairs are clustered and follow a trendline with a slope slightly higher than 45°.

The main reason for this difference is that the costs in each

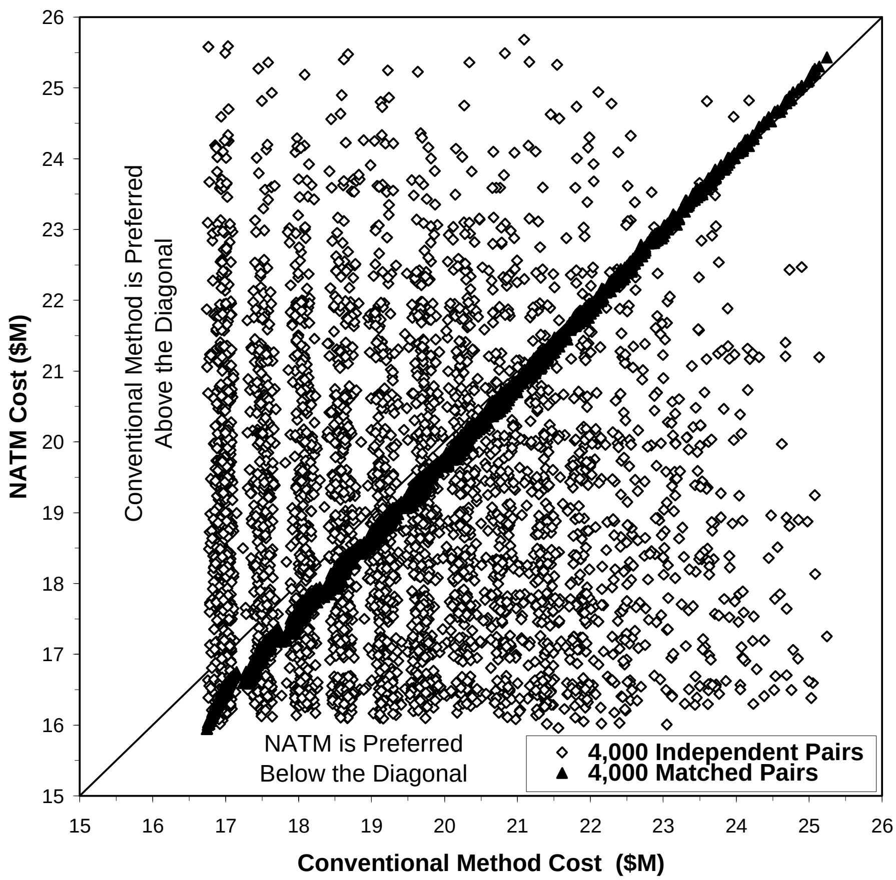

FIG. 2. Total Cost of the Conventional Method and the NATM

independent pair correspond to two different geologic profiles, whereas the geologic profiles for the costs in each matched pair are exactly the same. An independent pair in the northwest (NW) corner, for example, indicates that the conventional method is less expensive because it corresponds to a very favorable geologic profile for the conventional method and a very unfavorable profile for NATM. An independent pair in the southeast (SE) corner indicates that NATM is less expensive because the situation is exactly the opposite.

A total of 2,213 independent pairs, or 55%, are below the SW-to-NE diagonal shown in Fig. 2. This means there is an approximately 45% chance that a single independent pair will indicate that the conventional method costs less than NATM and should be the preferred alternative. Similarly, there is a 55% chance that a single independent pair would indicate exactly the opposite. Overall, however, it is not obvious by looking at Fig. 2 whether one construction alternative should be preferred over the other (with any degree of confidence).

In contrast, the matched pairs in Fig. 2 illustrate that for this example NATM is clearly the preferred alternative. This is most evident for the matched pairs in the SW corner, which correspond to very favorable geology and for which NATM is clearly the less expensive choice. This conclusion makes common sense because NATM is an adaptable method and can thus take advantage of good geologic conditions by minimizing the amount of the required initial support (the conventional method is more expensive because it uses the heavy support required by ground class 3). Matched pairs in the NE corner, however, indicate exactly the opposite. When the geologic conditions are bad, there is no advantage to the adaptability of the NATM, and the conventional method is the preferred alternative because it is less expensive.

Overall, however, 3,826 of the 4,000 matched pairs, or 96%, lie below the SW-to-NE diagonal. Since the exact geologic profile for the tunnel is not known a priori, this is a clear indication that the NATM is less expensive and thus the preferred alternative (from an expected geologic profile point of view).

Another point illustrated in Fig. 2 is that for a given geology the difference in cost between the two methods is actually quite small and also has a small variance (even though the variance of each method's cost is quite large as evidenced by the spread from \$16M to 26M). This is clearly indicated by the narrow width of the band formed by the matched pairs. For a given cost of the conventional method, the spread of the corresponding cost for the NATM matched pairs is less than \$0.5M.

Although not directly related to this discussion, it is interesting that the independent pairs in Fig. 2 tend to be clustered in vertical and horizontal stripes, and that the matched pairs tend to form slanted stripes (at approximately 60°). These stripes are particularly pronounced in the SW corner of the figure. This phenomenon is caused by the assumptions for this example, particularly the way geology impacts project cost, especially for the conventional method.

The conventional method incurs no cost (or time) penalties due to a transition from one ground class to another. Only the total number of steps occupied by each ground class impacts its total cost (the number of transitions is irrelevant). Since the total number of steps is 16 and the first step is given to be in ground class 1, the total number of different geologic profiles (from the point of view of the conventional method) is 16 + 15 + … + 1 = 136 (given that exactly i = 1,…, 16 steps are in ground class 1, each term in this sum represents the ways the number of steps occupied by class 2 plus the number of steps occupied by class 3 can add up to the remaining 16-i steps). Clearly, this number is much smaller than the 3^{15} different possible trajectories for the underlying Markov process. Furthermore, the total costs per step for ground classes 1 and 2 are very close and significantly less than the cost per step for ground class 3. This means that only 16 of the 136 profiles are significantly different from each other (depending on how many of the 16 steps belong to ground class 3, i.e., 0, 1,…, 15). Thus, each of the vertical stripes for the independent pairs corresponds to one of these 16 "significantly different" profiles. The horizontal stripes are caused by similar reasons. However, they are less pronounced because the number of ground class transitions does impact the cost of NATM. This increases the vertical spread of costs within each horizontal stripe causing adjacent stripes to merge together.

COMPARISON OF TUNNEL CONSTRUCTION ALTERNATIVES

If we let ΔC = (total cost of NATM) - (total cost of the conventional method), then the selection of construction method depends on whether the expected value of ΔC is positive or negative. (As shown in the following section, the average difference in project duration, ΔT, for the two construction alternatives is small and will be ignored.)

Table 6 shows the cost and time for the two construction alternatives given by the first 20 as well as the last of the 4,000 replications (for independent and matched pairs). The last four rows give statistics for the entire set of 4,000 runs. These statistics give the average, standard deviation, maximum, and minimum for the corresponding column.

An examination of these statistics shows that the conventional method has an average cost of \$19.337M and a standard deviation of \$1.878M. The NATM has two sets of statistics. The independent pairs produce an average cost of \$19.047M and a standard deviation of \$2.022M, and the matched pairs produce an average cost of \$19.025M and a standard deviation of \$2.037M. As these results illustrate, the two sets of statistics for the total cost of the NATM are very close (i.e., it makes no difference whether they are produced by independent or matched-pair runs). Thus, it is safe to conclude that the two sets of simulation runs are equally effective in estimating the total costs of the two construction alternatives (in absolute terms and without comparing one construction alternative to the other). As explained in the following section, however, independent pairs and matched pairs do not perform equally well when used to estimate the difference in cost between the two construction alternatives.

The (NATM-conventional) columns show ΔC for each replication. The two "running average" columns (one for independent and one for matched pairs) show how the value of the average ΔC changes as more simulation runs are performed. The average, standard deviation, minimum, and maximum of ΔC (as functions of the total runs performed) are also shown in Fig. 3 (for independent pairs) and Fig. 4 (for matched pairs). Notice that the ordinates in these two figures differ in terms of cost scale by an order of magnitude.

The two "running average" columns, as well as the corresponding line graphs in Fig. 3 and 4, are of particular interest because they indicate the best estimate for the expected value of ΔC if no additional simulation runs were to be performed. Thus, they indicate what the preferred construction alternative would have been had we stopped at the corresponding replication. It is clear that for up to 20 replications these two running averages disagree. The one for independent pairs has positive values and favors the conventional method whereas the one for matched pairs has negative values and favors NATM. The average ΔC produced by the entire set of 4,000 independent pairs, however, indicates that for the given data NATM is the better choice with expected cost savings of approximately \$291,000 (which is very close to the expected savings of \$312,000 produced by the matched pairs).

Table 6 shows that the ΔC given by independent pairs varies significantly from one replication to the next. Thus, the fact that the first 20 replications for the independent pairs produce an average ΔC that is positive (leading to the wrong conclusion that the conventional method is less expensive) happens quite by chance. It is due to the large variability in ΔC, which is mainly caused by the fact that, for a given run, the simulated project geology for the two construction methods is not the same.

In contrast, the first 20 replications for the matched pairs (where the two costs correspond to exactly the same geology) produce ΔC values that are much smaller and closer together. Thus, even though the individual costs for the two construction methods vary significantly from one matched pair to the next (different geology), the variability in ΔC is now much smaller, indicating that given the same project geology, the costs for the two methods are quite close.

Table 6 shows that the average ΔC given by independent and matched pairs for 4,000 replications are almost the same (–\$291,000 versus –\$312,000). This is not surprising for the following reasons:

1. Probabilistically, the expected value of ΔC equals the difference of the expected cost values for the two construction methods irrespective of correlation.
2. The average ΔC should be close to its expected value because the sample size of 4,000 is exceptionally large (thus, brute force makes the variance of the average ΔC very small).

In contrast, the two estimates for the standard deviation of ΔC in Table 6 and Figs. 3 and 4 are very different (\$2,762,000 versus \$169,000). This indicates that increasing the number of replications for the independent pairs (even up to 4,000) does not decrease the observed variability in ΔC (even though it decreases

TABLE 6. Total Cost and Time of Conventional Method versus NATM for Independent and Matched Pairs

Replication (1)	Conventional Method: Cost ($M) (2)	Conventional Method: Time (days) (3)	INDEPENDENT PAIRS — NATM: Cost $M (4)	NATM: Time (days) (5)	NATM–Conventional: ΔC ($M) (6)	NATM–Conventional: ΔT (days) (7)	Running average ΔC ($M) (8)	MATCHED PAIRS — NATM: Cost ($M) (9)	NATM: Time (days) (10)	NATM–Conventional: ΔC ($M) (11)	NATM–Conventional: ΔT (days) (12)	Running average ΔC ($M) (13)
1	18.480	352	20.072	401	1.592	49	1.592	18.049	357	–0.431	5	–0.431
2	19.594	370	21.811	434	2.218	64	1.905	19.171	377	–0.423	8	–0.427
3	17.477	336	18.162	360	0.685	24	1.498	16.990	336	–0.487	0	–0.447
4	16.890	322	17.848	355	0.958	32	1.363	16.309	320	–0.581	–2	–0.480
5	19.760	377	17.004	337	–2.757	–39	0.539	19.418	384	–0.342	7	–0.453
6	16.989	326	16.385	323	–0.605	–3	0.348	16.472	326	–0.518	–1	–0.464
7	16.952	325	22.553	450	5.601	125	1.099	16.389	322	–0.563	–2	–0.478
8	18.585	360	16.444	325	–2.140	–35	0.694	18.297	366	–0.288	6	–0.454
9	19.630	378	18.190	361	–1.439	–18	0.457	19.400	387	–0.229	9	–0.429
10	17.026	331	16.168	315	–0.857	–16	0.326	16.588	330	–0.437	–1	–0.430
11	20.389	394	20.777	413	0.387	19	0.331	20.231	404	–0.159	10	–0.405
12	17.519	338	22.099	439	4.580	101	0.685	17.099	340	–0.420	2	–0.406
13	19.534	378	21.356	426	1.822	48	0.773	19.325	387	–0.209	9	–0.391
14	23.668	452	22.911	458	–0.756	6	0.663	23.637	468	–0.031	16	–0.366
15	18.923	359	16.524	328	–2.399	–31	0.459	18.479	364	–0.444	4	–0.371
16	18.590	355	17.749	353	–0.841	–2	0.378	18.158	358	–0.433	3	–0.375
17	21.958	423	23.385	466	1.427	44	0.440	21.904	439	–0.055	16	–0.356
18	17.673	344	19.980	397	2.307	53	0.543	17.312	346	–0.361	2	–0.356
19	18.623	359	20.424	405	1.800	47	0.610	18.270	363	–0.353	5	–0.356
20	18.435	348	18.347	364	–0.088	16	0.575	17.906	351	–0.529	3	–0.365
—	—	—	—	—	—	—	—	—	—	—	—	—
4000	18.694	361	18.028	353	–0.666	–8	–0.291	18.360	365	–0.335	4	–0.312
Average	19.337	371	19.047	378	–0.291	6.9	—	19.025	378	–0.312	6.5	—
Standard deviation	1.878	36	2.022	41	2.762	55.0	—	2.037	42	0.169	5.9	—
Maximum	25.244	485	25.680	513	8.819	196.6	—	25.433	510	0.192	25.6	—
Minimum	16.738	312	15.958	307	–8.647	–158.2	—	15.942	307	–0.797	–5.1	—

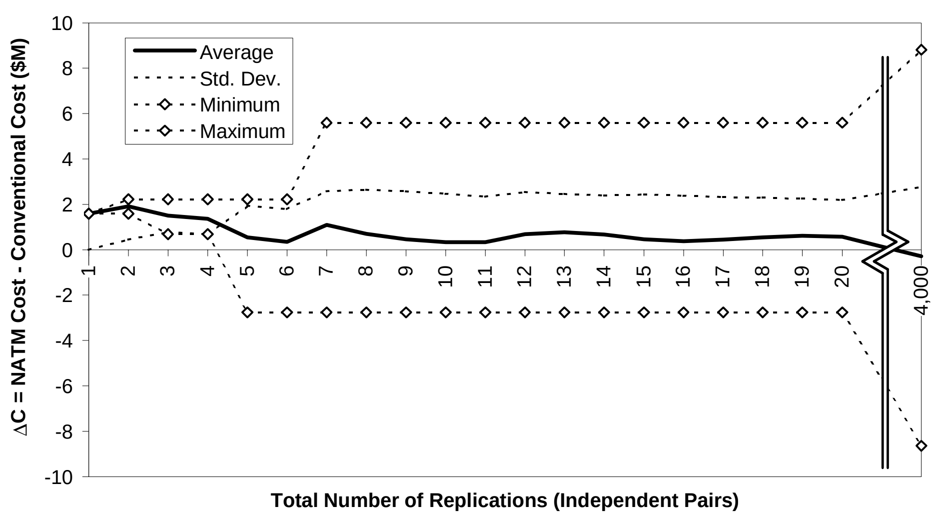

FIG. 3. Effect of Number of Replications on Difference in Cost between NATM and Conventional Method (Independent Pairs)

the variance of the average ΔC by brute force). Thus, expending computing resources is not an acceptable substitute for good experimental design.

The use of matched pairs reduces the standard deviation SD[ΔC] by an order of magnitude (from $2,762,000 to $169,000), and thus reduces the variance Var[ΔC] by two orders of magnitude. This large reduction in variance is due to the high correlation induced by matched pairs between the costs of the two methods. Visually, this correlation is evident in the two sets of points plotted in Fig. 2 and their difference in scatter. The correlation coefficient for the 4,000 independent pairs is –0.2% (which, as expected, is very close to 0). In contrast, the correlation coefficient for the 4,000 matched pairs is +99.95% (very close to +1). As a result, the variance of ΔC for independent pairs is very large because it is approximately equal to the sum of the variances of the costs of the two methods (the covariance term is very small). For matched pairs, the standard deviation of ΔC is approximately equal to the difference in the standard deviations of the two costs and thus is much smaller.

This reduction in variability also increases the probability that any single pair would produce a negative ΔC (and thus indicate correctly that NATM is the preferred alternative) from 55% (for independent pairs) to 96% (for matched pairs). It also reduces the number of simulation runs required to achieve the desired confidence interval width for the expected value of ΔC. Given the same number of replications, the variance of the average ΔC equals the variance of ΔC divided by the sample size n. This means that in order to achieve the same confidence interval for the true expected value of ΔC we can either perform a few replications using matched pairs or rely on a significantly larger number of independent replications.

Fig. 5 shows how the 95% confidence intervals for the expected value of ΔC change as more replications are performed. For independent pairs these intervals are very wide, reflecting the aforementioned high standard deviation of ΔC. Moreover, they

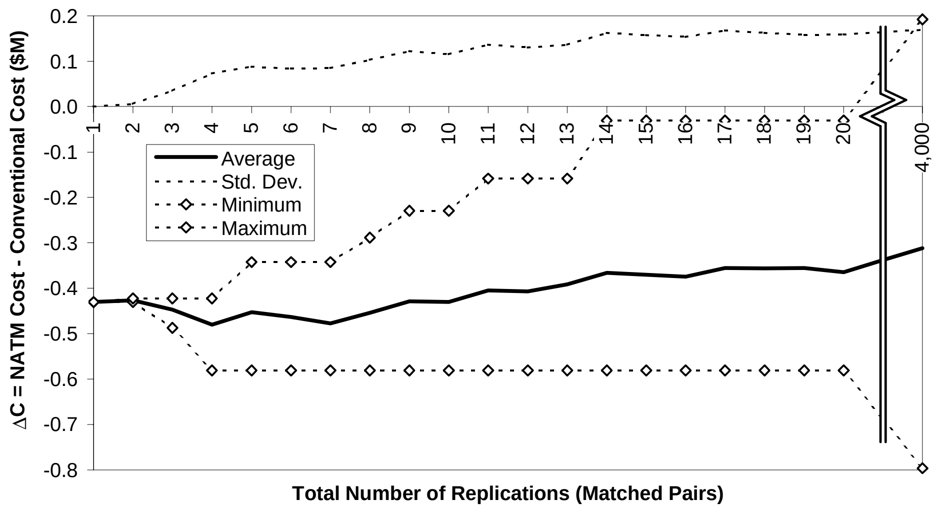

FIG. 4. Effect of Number of Replications on Difference in Cost between NATM and Conventional Method (Matched Pairs)

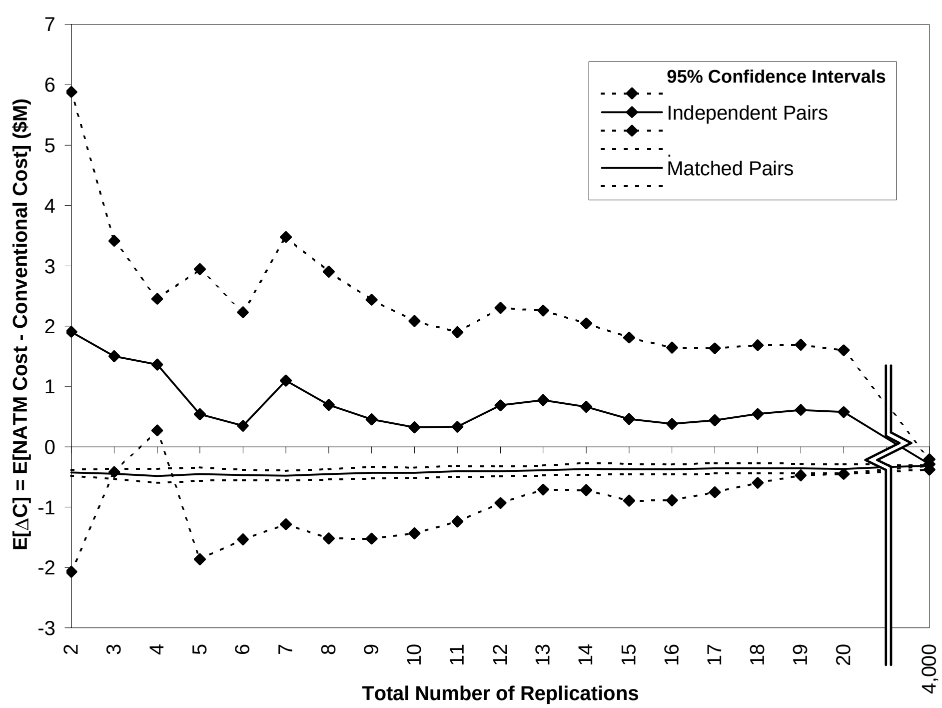

FIG. 5. Ninety-Five Percent Confidence Intervals for Expected Difference in Cost (E[ΔC]) between NATM and Conventional Method as a Function of Number of Replications

are centered over positive values of ΔC (which in this example leads to the wrong choice of method) and include both positive and negative regions for the value of the true mean of ΔC (indicating ambiguity as to which method is less expensive). In contrast, the intervals for the matched pairs are very tight and stable, even for as few as three or four replications. In fact, the width of the confidence interval given by 4,000 independent pairs (brute force) can be achieved with as little as seven matched pairs.

WHEN TO USE MATCHED PAIRS

It is important to point out that there is a statistical as well as a logical reason for using matched pairs to compare construction alternatives. From a statistical standpoint, the use of matched pairs ensures that the way chance affects the two alternatives is similar and equitable. In other words, it ensures that alternative A does not appear good because of a stroke of good luck whereas alternative B appears bad because of a stroke of bad luck. Furthermore, the effectiveness of matched pairs (as well as the need for using the technique) increases as more of the variance in the system is due to factors that are external to the construction alternatives being compared (i.e., when the performance attribute used for comparison depends on these outside factors). The use of matched pairs in the foregoing tunneling project, for example, results in a spectacular reduction in variance (by two orders of magnitude) because most of the variance in the cost of both tunneling alternatives is caused by geologic uncertainty.

It is not difficult to show that the expected reduction in variance should be smaller (but still existent) if one of the tunneling methods is not sensitive to the geology. To illustrate this, let us assume that the excavation and initial support data for the conventional tunneling method do not depend on the geology in any way (i.e., they follow probability distributions that are independent of the current ground class). In this case, producing one sample cost for the conventional method, and then using the same geologic profile to produce one sample cost for NATM (to compare the two), does not make any difference because the former does not depend on the geology anyway. Even in this hypothetical case, however, the statistical reason for using matched pairs remains valid from a variance reduction standpoint. Having the excavation advance rate for both methods face the same string of luck still makes sense because it reduces the variance of their cost difference (albeit, not as much).

In addition to the preceding statistical reasons, certain construction settings require the use of matched pairs for purely logical reasons. In these cases, there is no choice: use of the technique is mandatory, otherwise the comparison of alternative methods does not make logical sense.

An example of such a case is again provided by the tunneling project. The estimate of \$0.169M for the standard deviation of ΔC produced by matched pairs has a clear meaning. It refers to the difference in cost for the two construction alternatives in the conditional space of “other things being equal” (i.e., same project and same geology). In contrast, the standard deviation estimate of \$2.762M produced by independent pairs is a mathematical construct that is not as easy to describe. Independent pairs cannot logically be used to decide between the conventional method versus NATM for a particular project because they provide no simple and defensible answer to the question, “under what conditions exactly are the two alternatives being compared?” No matter which construction method is selected, the geology of the project will be the same (even though it is unknown prior to construction). Thus, it makes no sense to compare the two methods without first assuming the exact same geologic conditions. The use of matched pairs to compare the conventional method and NATM in this example is not simply a matter of choice or efficiency (i.e., to decrease the variance of the average ΔC), it is a necessity for the results to make engineering common sense.

It is very easy for the uninitiated in simulation experiments to assume that the large standard deviation produced by independent pairs is indeed a meaningful measure of the variability of ΔC (i.e., to confuse it, in terms of meaning, with the estimate produced by matched pairs). Such an oversight is easy to make and can lead to erroneous conclusions. For example, one can easily argue that if the costs of the two methods are so variable (from \$16M to \$26M) it “makes sense” for their difference in cost to be even more variable (i.e., to have a true mean of –\$0.291M and a standard deviation of \$2.762M). If one views the actual construction of the project as the next experiment, this perceived (but wrong) large variance would make it appear that the chance that one method is less expensive than the other is close to 50% (actually 55% versus 45%). As a result, it would make little sense for any company to switch from the conventional method, that has been used for years, to this new method that the company has little experience with. The fallacy in this argument is made obvious by the use of matched pairs. The probability that NATM will be less expensive, if used for this hypothetical example, is 96% (and not 55%). In other words, the conventional method is less expensive for only 4% (and not 45%) of the possible geologic profiles for this project.

CONCLUSIONS

Simulation is often used to evaluate and compare the merits of different construction methods according to some decision criteria in an attempt to select the best alternative. For such comparisons to be valid, it is very important to compare like with like. That is, to assume the same conditions for all alternatives so that none is favored more than the others due to external factors that have little to do with the method itself.

The selection of the least costly tunneling alternative for a given project is an excellent example of this problem. As illustrated, it makes little sense to compare the cost of one tunneling alternative against the cost of another if the costs of the two methods are based on different geologic profiles. The alternative that happens to correspond to the more favorable geologic conditions is also the most likely to appear as the least expensive.

To compare like with like, it is necessary to structure simulation experiments so that chance impacts all alternatives in a similar manner. This is most easily accomplished by using matched pairs, a variance reduction technique that applies specifically to the comparison of alternative system configurations. Despite its simplicity, matched pairs is a very effective and popular variance reduction technique. The basic idea behind this technique is to design the simulation experiments for the various system alternatives so that they all use random numbers in an identical or a very similar manner (i.e., so that chance does indeed impact all alternatives the same way).

To perform a meaningful comparison of construction alternatives it is imperative to know when to continue and when to stop performing additional simulation runs. The typical methodology for making this decision is based on the width of the confidence interval for the true expected value of the difference in cost between the alternatives. In general, simulation runs should be performed until the confidence interval for the true expected cost difference does not include the zero point. Otherwise, the confidence interval is ambiguous because the true expected difference in cost can be either positive or negative. In this respect, the use of matched pairs to reduce the variance of the cost difference can decrease the number of simulation runs significantly.

In addition to increasing statistical efficiency, matched pairs must also be used whenever construction alternatives need to be compared in the same conditional space (i.e., other things being equal). In these cases, the use of the technique is mandatory in order for the comparison of construction alternatives to make common sense. Examples of such cases are construction projects that are strongly affected by geologic uncertainty, such as tunnels, subways, and dams.

ACKNOWLEDGMENTS

The authors would like to acknowledge the University of Michigan Horace H. Rackham School of Graduate Studies and the National Science Foundation (Grant No CMS-9415105) for their support.

APPENDIX. REFERENCES

Halpin, D.W., and Riggs, L.S. (1992). *Planning and analysis of construction operations.* John Wiley & Sons, Inc., New York, N.Y.

Ioannou P.G. (1987). "Geologic prediction model for tunneling." *J. Constr. Engrg. and Mgmt.*, ASCE, (113)4, 569–590.

Ioannou P.G. (1989a). "Evaluation of subsurface exploration programs." *J. Constr. Engrg. and Mgmt.*, ASCE, (115)3, 339–356.

Ioannou P.G. (1989b). "UM-CYCLONE, discrete-event simulation system." *Tech. Rep. UMCE 89–11, 89–12*, Dept. of Civ. Engrg., Univ. of Michigan, Ann Arbor, Mich.

Ioannou, P. G., and Martinez, J.C. (1995). "Evaluation of alternative construction processes using simulation." *Proc., 1995 Constr. Congr.*, ASCE, New York, N.Y.

Law A.M., and Kelton W.D. (1991). *Simulation modeling and analysis*, 2nd Ed., McGraw-Hill Book Co., Inc., New York, N.Y.

Martinez J.C. (1996). "STROBOSCOPE — state- and resource-based simulation of construction processes." PhD thesis, Dept. of Civ. and Envir. Engrg., Univ. of Michigan, Ann Arbor, Mich.

Martinez J.C., Ioannou P.G. (1994). "General purpose simulation with STROBOSCOPE." *Proc., 1994 Winter Simulation Conf.*, IEEE, Piscataway, N.J., 1159–1166.

Martinez J.C., Ioannou P.G. (1995). "Advantages of the activity scanning approach in modeling complex construction processes." *Proc., 1995 Winter Simulation Conf.*, IEEE, Piscataway, N.J., 1024–1031.

Martinez J.C., Ioannou P.G., and Carr, R.I. (1994). "State- and resource-based construction process simulation." *Proc., ASCE 1st Congr. on Computing in Civ. Engrg.*, New York, N.Y.