A Statistics of Rare Events Method for Transportation Systems¹

Allan L. White
NASA Langley research Center
Hampton, VA 23681-0001
a.l.white@larc.nasa.gov
(757) 864-6199

Abstract—A method is proposed for quantifying the expected number of accidents for a transportation system during some operating period. The operating period is divided into two parts. There is normal operation where everything is working correctly. These intervals can be studied deterministically by arguments-from-design or by tests. There is unsafe operation where equipment has failed, an error has occurred, or traffic perturbations have produced unusual circumstances. Such stochastic phenomena can be studied by experiments or simulation. These two types of operation create a natural partition. This paper proposes a Monte Carlo method based on this partition that appears appropriate for studying scarce events. Estimators for this method are developed. It is shown they are unbiased, and confidence intervals derived. There is also a discussion of integrating random failures with traffic flow in a discrete event simulation.

Table of Contents

7. Introduction
8. Partitioned Monte Carlo
9. Device Failure
10. The Transportation System
11. The Experiment
12. Further Work
13. Summary

1. INTRODUCTION

A number of transportation systems (railroads, rail transit, and air traffic) are planning upgrades to increase performance and safety. The advanced technologies bring both new methods of operating and new failure modes. In all cases, a preliminary quantitative assessment is desirable. It can check that performance and safety requirements are met, compare different proposed architectures, and identify the important features and parameters.

The material below considers estimating a safety parameter: the expected number of accidents during the system's lifetime.

There are a number of reasons that a quantitative assessment of safety is challenging. First, if the

system is well designed with fail safe procedures then accidents are rare events, and it will require numerous simulations to get a statistically meaningful sample. Second, realistic estimates may require the simulation of the entire system or a large portion of the system to account for all the relevant interactions. Third, it may be desirable to study the system over its entire life span in order to take into account the life cycle and aging properties of the equipment. Fourth, there can be large differences between different simulation runs. The system contains a large number of devices, and different simulations can have different sets of devices fail. This can yield a large variance, which implies a large number of trials must be conducted to get an accurate estimate.

Hence, straightforward statistical methods are impractical, while special methods require a demonstration that the estimator is unbiased and a derivation of confidence intervals to ensure the accuracy of the estimate. This paper offers a Monte Carlo method based on partitioning the sample space.

Before proceeding, it is appropriate to discuss the concept of fail safe. In systems that operate for a long period of time, device failures and operator errors are almost certain. Hence, a principle in design and analysis is fail safe: the ability to respond to a failure or error in a manner that prevents accidents. What constitutes fail safe can vary from system to system. In complex systems it can be difficult to determine which responses are adequate. For this reason designers and analysts will tend to be conservative. A certain set of responses will be determined to be adequate; the designers will attempt to design the system to invoke these responses; and the analysts will determine the probability that the selected responses are invoked. In the analysis of the system, all other responses will be declared to be unsafe. The simulation below considers only failures that are unsafe failures. These are the failures that lead to accidents.

Section two describes partitioned Monte Carlo. It contains the probability and statistical analysis for the estimator. The first and second moments of the estimator are derived; it is shown that the estimator

¹ U.S. Government work not protected by U.S. copyright

is unbiased; and the confidence interval is computed. The section also derives the population variance in terms of the stratified parameters. Section three presents a simple failure/repair model, and discusses integrating device failure into traffic simulation. The example application described in section four attempts a simple application with a number of features found in transportation systems. The equipment has fail safe and fail unsafe modes along with repair-on-demand and periodic maintenance. There are protocols for traffic movement. There is a possibility of an accident only in the case of unsafe failure and certain traffic conditions. Section five describes the experiment, which considers the entire operating period of the system. Subsections cover the assumptions of the simulation, the details of the simulation program, and the results of the experiment. For this example, the partitioned Monte Carlo method efficiently yields tight confidence intervals. Section six discusses further work. The number of topics mentioned in this section illustrates the preliminary nature of this paper. Section seven is a summary.

2. PARTITIONED MONTE CARLO

The approach is to divide the problem into two parts: structural and stochastic. The structural part assesses the system under normal operation where all the parts are working correctly. It is done by arguments-from-design and tests that show there will be no accidents if all parts of the system (including human operators) are working correctly. There are traffic perturbations, but the rules governing traffic are deterministic. It is desirable to check that these rules are complete and consistent. On the other hand, reasoning about the possible outcomes in the event of a random failure can be overwhelming. The stochastic part considers equipment failures and operator errors and studies the possible consequences by experiment and simulation. All of this leads to a partitioned Monte Carlo procedure for quantifying the outcomes of rare events. The partition is into normal operation and unsafe operation. This partition requires its own estimators along with a demonstration that they are unbiased and a derivation of their confidence intervals.

There are two categories of partitions. First, if events (unsafe failures) are extremely rare, then there is a significant probability they will not occur during the system's operating life. Second, if an unsafe failure occurs the simulation need only be run for some convenient period that includes the interval from occurrence time to removal time.

The rest of this section conducts a statistical analysis for the first category of partitioning, where the objective is to estimate the expected number of accidents during the systems lifetime. This is equivalent to estimating the population average where the population consists of accidents during the operating period.

Mean and variance of the estimator for the population average

Let X be the subpopulation with no unsafe failures during the operating period, and Y the subpopulation with unsafe failures. Let p_X and p_Y be the probabilities that an operating period belongs to X or Y . Let μ_X and μ_Y be the expected number of accidents for the two subpopulations. By assumption, $\mu_X = 0$. If μ is the population average then $\mu = p_X \mu_X + p_Y \mu_Y = p_Y \mu_Y$.

Let θ be the sample average (of the number of accidents) for the subpopulation Y . Then $E[p_Y \theta] = p_Y E[\theta] = p_Y \mu_Y$ which implies $p_Y \theta$ is an unbiased estimator for the population average.

Since θ is a sample average, it is normally distributed. Suppose the sample size is K ; the sample variance is s_Y ; and N_α is the number of standard deviations for the standard normal that give a $100(1-\alpha)\%$ confidence level. Then

$$\theta \pm N_\alpha \frac{s_Y}{\sqrt{K}} \quad (1)$$

gives a $100(1-\alpha)\%$ confidence interval for μ_Y . Hence,

$$p_Y \left(\theta \pm N_\alpha \frac{s_Y}{\sqrt{K}} \right) \quad (2)$$

gives a $100(1-\alpha)\%$ confidence interval for $\mu = p_Y \mu_Y$, the population average.

The population variance

The subsection above derives an unbiased estimator and confidence intervals for the population average. The population variance indicates the significance of the population average. (If the variance is large, then the population average is not very significant.) We derive an expression for the population variance in terms of the average and variance for the subpopulation Y . Recall that $\mu = p_Y \mu_Y$ and

$E[x]=0$. Let z be the number of accidents during an operating period.

$$\text{Var}[z] = E[(z - \mu)^2]$$

$$= p_X E_X[(x - \mu)^2] + p_Y E_Y[(y - \mu)^2]$$

$$= p_Y (\sigma_Y^2 + \mu_Y^2) - p_Y^2 \mu_Y^2 \quad (3)$$

$$= p_Y (\sigma_Y^2 + \mu_Y^2) - p_Y^2 (\sigma_Y^2 + \mu_Y^2) + p_Y^2 \sigma_Y^2 \quad (4)$$

Expression (3) has been rewritten as (4) as a convenient method of getting an unbiased estimator. (The square of the sample average is not an unbiased estimate for the square of the population average.) Use the usual unbiased estimate for the variance. The term in the parentheses in (4) is the second moment, and an unbiased estimate is the sample second moment.

3. DEVICE FAILURE

This section presents the model for device failure used in the simulation. For models of this type there is a complete analytical solution for unsafe failure occurrence, which is given below.

The Markov mode for failure and repair

The device is assumed to have a single transition from the good state to failure which occurs with constant rate λ . It fails safe if the failure is detected, which happens with probability C (for coverage). It fails unsafe otherwise. Upon safe failure, it is repaired with rate ρ . The Markov model is given in figure 1. The mnemonics are G for the good state, SF for safe failure, and UF for unsafe failure.

In the simulation in this paper, an unsafe failure is detected by periodic maintenance or by occurrence of an accident.

Figure 1: Model for Device Failure

The unsafe failure rate

We will assume that repair proceeds quickly, and use the approximation that the time spent repairing the system can be ignored. With this approximation, failures are a Poisson renewal process with rate r = sum of all failure rates. Suppose the probability of an unsafe failure given a failure is $1-C$. Let T be the operating time. The derivation that unsafe failures are a Poisson process with rate $r(1-C)$ is

$$P\{k \text{ unsafe failures}\}$$

$$= \sum_{n=k}^{\infty} P\{k \text{ unsafe fail} | n \text{ fail}\} P\{n \text{ fail}\}$$

$$= \sum_{n=k}^{\infty} \left[\frac{(rT)^n e^{-rT}}{n!} \right] \left[\binom{n}{k} (1-C)^k C^{n-k} \right]$$

$$= \frac{(rT)^k (1-C)^k e^{-rT}}{k!} \sum_{n=k}^{\infty} \frac{(rTC)^{n-k}}{(n-k)!}$$

$$= \frac{(rT)^k (1-C)^k e^{-rT(1-C)}}{k!} \quad (5)$$

Integrating failures into the traffic simulation

This simulation assumes that failures in traffic control are independent of traffic conditions. (This is not always a valid assumption, especially when considering human operators.) The independence assumption implies that the time and place of failures can be determined separately and then inserted into the list of events for the traffic simulation.

The failure/repair in the first subsection above is simple enough that the time and place of unsafe failures can be determined analytically. In fact, the second subsection above shows that it is a Poisson process. The time and place for failure occurrence for more complicated models can be determined by simulation.

4. THE TRANSPORTATION SYSTEM

The system is a canonical railroad example. There are two yards connected by a single track with a side track midway. Movement is by fixed block protocol. That is, the track layout is divided into disjoint sections called blocks, and there are gates (signals/switches) at the end points of the blocks. Rules based on block occupancy allow trains to leave one block and enter another. There is an accident if two (or more) trains occupy the same block, even if they are going in the same direction.

The diagrams for the track, gates, and blocks are given below. The yards at either end of the track are not indicated in the figures.

Figure 2: The five track gates

Figure 3: The four track blocks

The protocol for moving from left to right is

If the train is at gate #1 it can proceed to gate #2 if blocks 1, 2, and 3 are empty.

If the train is at gate #2 it can proceed to gate #4 if blocks 2 and 4 are empty.

If the train is at gate #2 it can proceed to gate #3 otherwise.

If the train is at gate #3 it can proceed to gate #4 if block 4 is empty.

If the train is at gate #4 it can proceed to gate #5 if block 4 is empty.

The rules for the other direction are similar. For moving from left to right gate #4 is a bit superfluous. The rule for gate #3 would be safer if it were

If the train is at gate #3 it can proceed to gate #4 if blocks 2 and 4 are empty.

The current rule depends on the other rules not letting blocks 2 and 3 be occupied by trains going the same direction. There is no check for trains entering the yards from blocks 1 and 4. It is currently assumed that the train can always safely enter the yard. This can be modified, but it requires information on yard traffic. It is possible that very little information is required, just the percentage of time that the train can safely enter the yard.

For this example, the main track is 60 miles long. The side track is 2 miles long and is positioned midway. There is one kind of train (a slow freight) that moves between any two points at an average speed of 30 miles an hour

A gate can fail in a safe mode or in an unsafe mode. In a safe failure, the gate or system monitors notify the operators and traffic is halted until repairs are made. There are a variety of unsafe failures, but we just consider one kind: the gate lets all traffic pass regardless of block occupancy. For this simulation, an unsafe failure is detected by periodic maintenance or by the occurrence of a train accident.

5. THE EXPERIMENT

The purpose of the experiment is to estimate, and get a confidence interval for, the expected number of accidents during a 25 year operating period of the system.

Assumptions and conditions

The assumptions for this experiment are (with some repetition)

The track, block, and gate layout are given by figures 2 and 3.

The protocol is as described in section four.

Gate failure and repair is given by figure 1.

Fail safe halts all traffic until repair. If a gate fails unsafe it lets all traffic pass regardless of block occupancy.

There is monthly (periodic) maintenance that detects and repairs all faulty gates.

Both repair-on-demand and periodic maintenance occur quickly enough that there is no significant disruption of traffic.

The system operates correctly if fault free or if the gates fail safe. The only hazardous interval lies between the occurrence of an undetected failure and the end of the monthly maintenance check. There is an accident only if during this period two (or more) trains enter the same block.

An accident causes an investigation that finds and repairs the unsafe failure.

Traffic disruptions from an accident dissipate before the next accident.

The traffic consists of 60 trains per month equally divided between the two yards. For each yard, the 30 trains arrive by uniform distribution (during the month) and proceed to the other yard according to protocol.

The yards are 60 miles apart; the side track at midway is 2 miles long; and trains average 30 mph between any two points.

The failure rate for each gate is 1/year. Two coverage values (probability of safe failure given a failure) are used for two different experiments $C = 0.999$ and $C = 0.99999$.

Details of the simulation

The heart of the simulation is an ordered (by time) list of events that occur during a month in which there is a gate failure. An event is either the arrival

of a train at a gate or the (unsafe) failure of a gate. Each event is stored as a vector. A train vector contains its identification as a train, its arrival time at the next gate on its journey, and its direction of travel. A failure vector contains its identification as a failure event, which gate fails, and its time of failure.

Before the simulation begins an unordered list of events for the month is created. The first items on this unordered list are the gate failures. (These are placed first because in the event of a tie the program preserves the original order when ordering a list.) The second items are the arrival times of the trains at the first yard, which is gate 1 in figure 2. The third items are the arrival times of the trains at the second yard, which is gate 5 in figure 2. This entire list is then ordered.

The simulation consists of multiple passes through the list. The program starts at the front of the list and examines each vector until it finds a train that can move or a failure event. When examining a train vector the program first checks its position and direction. The program then checks if the train can move to the next gate according to the block occupancy and protocols. If the train can move, its vector is updated with its new gate and arrival time. The block occupancy record is updated to show that this train now occupies another block. The program checks for several trains in one block, which is an accident. Once the program has done this for one train it reorders the list and begins another pass. When a train reaches its destination yard its vector is removed from the list.

When the program reaches the vector for a failure event it first changes the protocol. The failed gate lets all traffic through regardless of block occupancy. The failure vector is removed from the list and the program begins another pass starting at the beginning of the list. The program begins a new pass through the list upon the failure of a gate because there may be a train waiting at the gate that has just failed. This train will start moving upon failure of the gate.

The program continues until there is an accident or all vectors have been removed from the list. For these experiments it is sufficient to stop after one accident since the random determination of unsafe failures never gave more than one in a single month. If more than one (unsafe) failure had occurred in a single month, the simulation program would have had to be modified.

First experiment

With $C = 0.999$, the unsafe failure rate for five gates is 0.005/year or 0.125/(25 years) which gives a probability of $P_Y = 0.1175$ for an unsafe failure during a 25 year operating period.

The experiment used 1000 trials where each trial is an operating period with at least one unsafe failure. Using the methods in section two gave

927 trials with one unsafe failure
72 trials with two unsafe failures
1 trial with three unsafe failures

The sample average for these 1000 trials is 0.40 and the sample standard deviation is 0.51. Hence, the estimate for the entire population is 0.050 and the 99% confidence interval is [0.045, 0.056].

Extended analysis for first experiment

The estimated standard deviation for the entire population is 0.22 which is rather large compared to the estimated average of 0.047. This indicates that the expected number of accidents is not very meaningful. The variance is large enough that there is a high probability that the actual number of accidents can differ significantly from the expected number.

For this simple case it is possible to derive the probability distribution for accidents. From the simulation, the probability of an accident given an unsafe failure is 0.40. Since unsafe failures occur at the rate 0.125 per operating period, accidents occur (as derived in section two) at the rate $(0.125)(0.40) \cong 0.05$. Hence, accidents have the Poisson distribution with rate 0.05. This explains the large variance since the Poisson distribution has a long tail.

[Note the probability of an accident given an unsafe failure (0.40) is numerically equal (to two decimal places) to the probability of an accident given one or more failures. This is because most (927) of the 1000 trials with one or more accidents have just one accident.]

Second Experiment

A second experiment was conducted to illustrate that this method can handle extremely rare events. The coverage was set equal to $C = 0.99999$. The unsafe failure rate is 0.00125/(25 year period). One thousand trials were performed where there was

at least one unsafe failure during the operating period. There were

998 trials with one unsafe failure
2 trials with two unsafe failures

The estimate for the population average is $4.76e-4$ with a 99% confidence interval of $[4.27 e-4, 5.27 e-4]$.

6. FURTHER WORK

Numerous items come immediately to mind. (1) Apply this method to larger and more complex systems. There are two opposing factors for this. On one hand a larger number of devices decreases the probability of operating periods without unsafe failures. On the other hand the devices can have more safeguards than the simple coverage described in figure 1, which increases the probability of operating periods without unsafe failures. (2) Development and analysis of failure/repair models for the more complicated devices. A guideline is that the safety and its analysis depend only on parameters that can be observed from field data or by experiment. Numerous devices with their different and more complicated models will make it more arduous to determine the number, time, and place of failure events. (3) Include time varying failure rates for equipment that wears out. This extension changes the nature of the simulation since the system is constantly changing during its operating period. (4) Concoct methods for determining the probability distribution of accidents as opposed to just estimating the average per operating period. As seen in section five above, the population variance can be large enough that just the expected number of accidents is not a sufficient description. (5) Include vehicle perturbations. This is especially relevant for air traffic where a plane can deviate from its flight path. One approach is to treat vehicle perturbations in the same manner as unsafe device failure: there is a stochastic model for vehicle perturbation and this is used to determine the frequency, time, and place for unsafe deviations. (6) Study the effect of correlation. Both device failures and vehicle perturbations can have a common cause such as weather or power supply problems. There is also a chain reaction where failures and perturbations place a heavier workload on human operators. It is possible to include correlation in Markov (and other) probability models but at the cost of added complexity. (7) Failures may arise from the system instead of from a particular device. For instance, rules and protocols may be adequate for a certain volume of traffic, but system wide perturbations can produce a

local excess. This type of failure can be handled the same as a part failure, although a quantitative study requires knowing the probability distribution for such events. (8) Develop models and gather statistics for transient and intermittent failures. This opens a large field of study. It includes designing devices that recover from transients, determining how much information is needed about transients and intermittents, and distinguishing between transient faults and permanent but intermittent faults. (9) Include the effect of human operators. This works both ways. On one hand, operators can cause accidents by ignoring signals. On the other hand, operators can act as monitors that prevent accidents in the case of equipment failure.

7. SUMMARY

This paper presents a partitioned Monte Carlo approach for studying rare events in transportation system. The partitioning divides system operation into two types of periods: (1) normal operation where all parts are working correctly or have failed safe and (2) unsafe operation where some failure or perturbation implies there is the possibility of an accident. The objective is to estimate the expected number of accidents during an operating period. Unbiased estimators and confidence intervals are derived. The method is applied to a system that is small but that contains a number of realistic features. The equipment has fail safe and fail unsafe modes along with repair-on-demand and periodic maintenance; there are protocols for traffic movement; failures must be integrated into the traffic flow; and there is a possibility of an accident only in the case of unsafe failure and certain traffic conditions. For this example, the method yields an accurate estimate of the expected number of accidents during the system's lifetime.

ACKNOWLEDGEMENT

The author wishes to thank Ted Giras, Barry Johnson, and Todd Delong at the University of Virginia, Center for Safety Critical Systems, for discussions on rail transportation.

REFERENCES

[1] P. G. Hoel. *Introduction to Mathematical Statistics*, Wiley, New York, 1971.

[2] V. K. Rohatgi, *A Introduction to Probability Theory and Mathematical Statistics*, Wiley, New York, 1976.

Allan L. White is a research engineer at NASA Langley. His interests and publications include probability models, design of experiments, and stochastic control.