

OPTIMIZATION OF SCENARIO CONSTRUCTION FOR LOSS ESTIMATION IN LIFELINE NETWORKS

Nathanael J.K. Brown
Jared L. Gearhart
Dean A. Jones

Sandia National Laboratories
PO Box 5800
Albuquerque, NM 87185-1138, USA

Linda K. Nozick
Natalia Romero
Ningxiong Xu

Cornell University
School of Civil and Environmental Engineering
220 Hollister Hall
Ithaca, NY 14853, USA

ABSTRACT

Natural disasters have become a pressing national and international problem. Population growth, aging infrastructure, and climate change suggest that mounting losses will continue into the foreseeable future, hence mitigation and response planning is of increasing importance. The conduct of studies to support this type of regional planning often requires an estimation of the impacts of a single earthquake scenario on a region. This paper describes a method to identify a set of consequence scenarios that can be used in regional loss estimation for lifeline systems when computational demands are of concern, and the spatial coherence of individual consequence scenarios is important. This method is compared with Monte Carlo simulation.

1 INTRODUCTION

Natural disasters have become a pressing national and international problem. Population growth, aging infrastructure, and climate change suggest that mounting losses will continue into the foreseeable future. Estimates for losses from landmark events in the United States are the Northridge Earthquake (1994), Hurricane Andrew (1992) and Hurricane Katrina (2005) which reached \$40B (State of California Department of Conservation 2011), \$32B (Sun-Sentinel 2011) and \$125B (USA Today 2005), respectively. Figure 1 illustrates worldwide, economic losses from natural disasters (Kunreuther et al. 2009). Notice that these losses have increased 15-fold since the 1950s. Considering just losses in the U.S. or just insured losses, the trends are similar. For simplicity, for the duration of this paper we focus on earthquakes with the understanding that the methods described in this paper can be applied to hurricanes.

There are many questions that arise in mitigation planning and response that require an estimate of the impact of a single earthquake scenario on a region. This event might be the repetition of a historical event of significance like the 1812 New Madrid earthquake or an event developed from a seismic study, similar to the ShakeOut scenario developed for Southern California in 2008 (USC 2008).

HAZUS is a loss estimation methodology developed by FEMA (2010) that can be used to provide estimates of damage to buildings and individual elements of transportation and utility systems based on a single event. HAZUS accomplishes this task by bringing together two key pieces of data for each element in the built infrastructure: fragility curves that are dependent on the type of structure and a probability distribution for the ground shaking at that location based on the event. A fragility curve specifies the probability that a structure of a specific type experiences a certain level of damage or greater for a given amount of ground shaking. Figure 2 gives an example of a set of fragility curves for a single type of highway bridge. HAZUS categorizes damage to highway bridges into five classes: no damage, slight,

moderate, extensive and complete. The curves are interpreted as follows. If the peak spectral acceleration at 1 second is 0.4g, the probability that the bridge experiences at least slight damage is 50%. This fact implies that the probability that no damage occurs is also 50%. Furthermore, the probability it experiences at least moderate, extensive or complete damage is 37%, 18% and 10%, respectively.




Figure 1: Worldwide economic loss by decade




Figure 2: Example fragility curves for a class of highway bridge (non-California, built prior to 1990 of conventional design with length >150 m)

The second element that is needed is a probability distribution for the ground shaking at the location of interest. This ground shaking is often estimated using empirical relationships called attenuation equations, which predict the distribution of ground motion amplitudes using a logarithmic mean that represents the distance, magnitude, type of rupture, soil classification, etc., and a standard deviation. It is commonly assumed that the residuals, as estimated by the standard deviation, are normally distributed about the mean.

The result of bringing together these two types of information is an estimate of the probability, for each component in the study area, that it will be in each of the several mutually exclusive and exhaustive damage states. For example, for a specific highway bridge there might be a 50%, 30%, 10%, 6% and 4% chance that it will be in no damage, slight damage, moderate damage, extensive damage, and complete

damage, respectively. Notice, this does not say which damage state the bridge will be in, only the probability that it will be in each state. Further, the performance of a highway system is driven by the damage to all the bridges because each bridge affects the connectivity of the network. Therefore, in order to assess the functionality of a lifeline system it is not only important to understand what might occur to each of the components but it is also important to understand how these components will interact when subsets of the components are in degraded conditions. To achieve this, it is necessary to develop scenarios that are consistent at the component level with this probabilistic information derived from the loss methodology.

The common strategy used to address this issue is Monte Carlo sampling. Chang, Shinozuka, and Moore (2000) randomly generate 10 consequence scenarios for a single event scenario using the probabilistic information derived for each bridge to analyze the impacts of earthquakes on the Los Angeles highway system. Similarly, Jayaram and Baker (2010) apply Monte Carlo sampling to the San Francisco highway system. Çağnan, Davidson, and Guikema (2006) use a mixture of Monte Carlo simulation and expert opinion to model the Los Angeles electrical system.

This paper describes an alternative method to identify a set of consequence scenarios that can be used in regional loss estimation in lifeline systems when computational demands are of concern, and the spatial coherence of individual consequence scenarios is important. Computation may be of particular concern: (1) when there are computationally intensive analyses to be conducted for each consequence scenario, as when one estimates damage to a lifeline network for each consequence scenario, then estimates service restoration times for each consequence scenario (e.g., Xu et al. 2007), or (2) when loss estimation must be repeated many times, as when evaluating the relative benefits of many mitigation alternatives (e.g., Dodo et al. 2005).

The following section describes an optimization model and a solution procedure for that model that can be used to identify a set of consequence scenarios and their probabilities of occurrence. The third section describes a realistic case study. The fourth section presents conclusions and opportunities for future research.

2 OPTIMIZATION MODEL

The input to the model is the vulnerability of each component of an infrastructure to an earthquake event of interest, which is the probability that the component is in each of a set of mutually exclusive and exhaustive damage states (as defined by a HAZUS analysis or other similar procedure). The optimization model is then used to identify a suite of consequence scenarios that match this vulnerability as “closely” as possible, given a user-specified limit on the number of consequence scenarios that may be used. Figure 3 illustrates this idea. The “true” vulnerability of a specific bridge is given by the diamonds. Based on the consequence scenarios generated, the probability that the bridge falls into each of the damage states is then as given by the squares. The goal is to select a set of scenarios so that across all the scenarios based on their likelihoods of occurrence, the difference from the probabilities generated by the loss estimation is as small as possible. This difference is illustrated by the value of the variable e_{kd} ; which is defined as the difference in the probability that bridge k is in damage state d as given by the loss methodology in contrast to that implied by the identified consequence scenarios.

The objective of the optimization is to minimize the sum of the errors over all bridges and damage states:

$$\min \sum_{kd} (e_{kd}^+ + e_{kd}^-), \quad (1)$$

where e_{kd}^+ and e_{kd}^- are, respectively, the errors resulting from overestimating and underestimating the probability that bridge k falls into damage state d . The error e_{kd}^+ is positive if the “true” values are overestimated and zero otherwise; e_{kd}^- is positive if the “true” values are underestimated and zero otherwise.

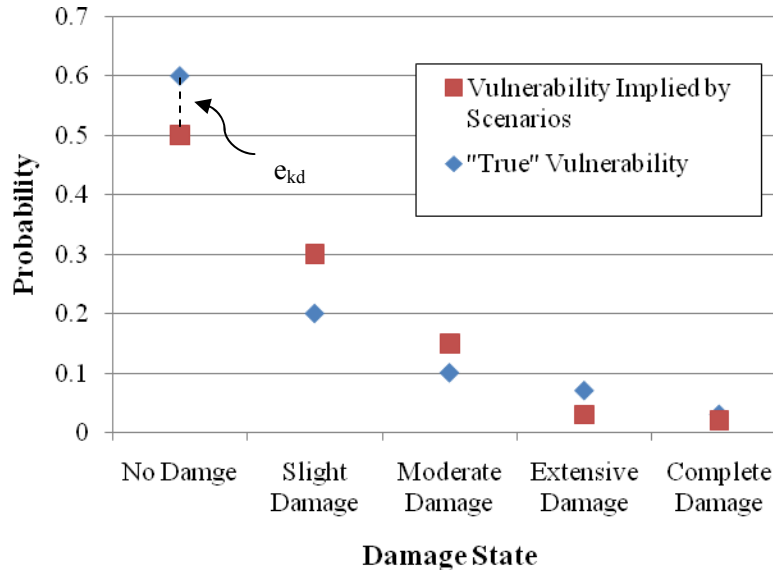


Figure 3: Schematic defining the error terms between (a) the “true” vulnerability of the bridge under the event and (b) the vulnerability as represented by the set of scenarios.

Suppose s_j is the probability of consequence scenario $j \in (1, \dots, J)$, and J is the number of consequence scenarios to be generated. The following constraint defines the error terms for each bridge and damage state, as the difference between the “true” vulnerability and that estimated using the consequence scenarios:

$$\sum_j s_j b_{jkd} - e^+_{kd} + e^-_{kd} = m_{kd} \quad \forall k, d, \quad (2)$$

where m_{kd} is the probability that bridge k falls into damage state $d \in (1, \dots, D)$ and b_{jkd} is a one if bridge k falls into damage state d under consequence scenario j and zero otherwise.

Under each consequence scenario, j , each bridge must be in exactly one damage state d . The following constraint ensures this:

$$\sum_d b_{jkd} = 1 \quad \forall k, j. \quad (3)$$

The probability of each consequence scenario j , must be between user-specified values s_{min} and s_{max} , where $0 \leq s_{min} < s_{max} \leq 1$, as given in equation (4):

$$s_{min} < s_j < s_{max} \quad \forall j. \quad (4)$$

One would typically assume that s_{min} and s_{max} are zero and one, respectively but if desired, the user can specify $s_{min} > 0$ and $s_{max} < 1$. It may be awkward to explain very small as well as very large values for s_j and this provides a mechanism to prevent that situation. Also, the sum of the probabilities across all consequence scenarios must equal one. The constraint which enforces this requirement is given in equation (5).

$$\sum_j s_j = 1 \quad (5)$$

Finally, the error terms must be nonnegative and each bridge under each consequence scenario must be either in a particular damage state or not.

$$e^+_{kd}, e^-_{kd} \geq 0 \quad \forall k, d \quad (6)$$

$$b_{jkd} \in \{0, 1\} \quad \forall j, k, d \quad (7)$$

The optimization model is a nonlinear integer program where the objective is given in equation (1) and the constraints in equations (2)-(7). The model determines the probability of each consequence scenario j , s_j , as well as the damage state for each bridge under that scenario, b_{jkd} . It also provides the errors, e^+_{kd} and e^-_{kd} for each bridge and damage state, so the user can see the magnitude and the distribution of the errors.

This model is challenging to solve because it is both nonlinear and integer. It is worth noticing, however, that if the value of the variables s_j are known, then the resultant optimization is a mixed-integer linear program (which decomposes by bridge where each mixed-integer program is on the order of $J*D$). Similarly, if the values of the variables b_{jkd} are known, the resultant optimization is a linear program. This observation leads to the following heuristic solution procedure.

1. Assume the values of all the variables s_j are $1/J$ and solve the resultant mixed-integer program.
2. Fix the values for b_{jkd} using the solution identified in step 1 and solve the resultant linear program for the variables s_j .
3. Fix the variables s_j and again solve the mixed-integer program for the variables b_{jkd} .
4. Iterate between steps 2 and 3 (where step 2 is executed using the values for the variables b_{jkd} identified in step 3) until there is no improvement in the objective.

It is worth noting that at each iteration of this algorithm the objective given in equation (1) will either improve or remain the same. It is this observation that leads to the simple stopping criteria given in Step 4 of the algorithm.

3 CASE STUDY

This section presents a case study application of the model to assess the vulnerability of the highway network in Memphis, Tennessee. Memphis is located in the southwest corner of the New Madrid Seismic Zone (NMSZ) which is illustrated in Figure 4. There are three key characteristics associated with the NMSZ which make understanding the consequences of earthquake events in this zone important. First, studies of historic records of large earthquakes in the area suggest that the time between large events is somewhere between 200 to 800 years with an average of about 500 years. The last recorded high magnitude earthquake associated with a fault rupture in this area was in 1811-1812. Based on the minimum recurrence rate, we might be close to an event (Tuttle et al. 2002). Second, due to soil conditions, ground shaking in this area is expected to affect a greater area than would be expected in California for a similar magnitude earthquake (Gomberg and Schweig 2007). Third, since earthquakes in the NMSZ are not as frequent as those on other faults, like San Andreas, there is an insufficient understanding of the earthquake risk in the area and therefore mitigation strategies in use may not be adequate.

3.1 Input Data

This case study focuses on an event with a magnitude of 7.7 at the Mississippi River halfway between Mississippi and Kentucky. This location is on the Mid-East branch of the synthetic faults created by the United States Geologic Survey (USGS) to represent the hazard in New Madrid (Peterson et al. 2008).

For this study we focus on developing scenarios to estimate the performance of the Memphis road network where the key vulnerabilities are in the 335 highway bridges built prior to 2000. Figure 5 illustrates the Memphis highway network. The spectral accelerations produced by this event at each of the bridges needed to compute the damage probabilities were estimated by code provided by USGS (2008).

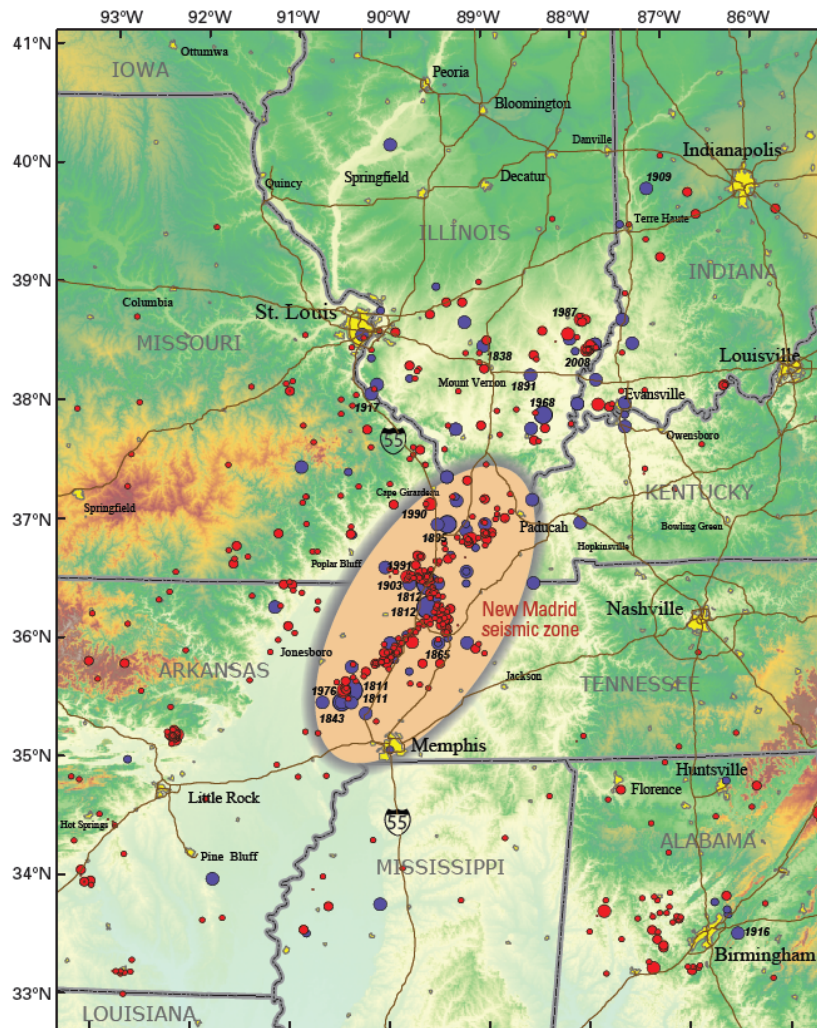


Figure 4: New Madrid Seismic Zone. Red and blue dots are historical earthquakes which occurred prior to and after 1973, respectively. Yellow polygons are metropolitan areas (Frankel et al. 2009).

3.2 Results of the Optimization and Comparison with Monte Carlo Simulation

Figure 6 gives the errors (objective value given by Equation 1) based on the number of consequence scenarios desired. There are 335 bridges so the sum over all m_{kd} is 335. This implies that a solution with no scenarios would have an error of 335. With only 5 scenarios the error is about 10%. With 13 scenarios the error is on the order of 2%.

Generally, as the number of scenarios needed rises, so does the number of iterations (which serves as a measure of computational burden). However, it is useful to notice that when 20 scenarios are needed, the number of iterations drops to only 3. While not illustrated in Figure 6, the solutions from step 1 (the first iteration) are quite good. The sum of errors are 66, 52, 48, 48, 35, 32 and 20 for 5, 7, 8, 9, 11, 13 and 20 scenarios, respectively. For 20 scenarios, the optimization of the values for b_{jkd} when all scenarios are assumed equally likely (result of the step 1 in the algorithm) produces a sum of errors of about 20 which is about 30% lower than the first iteration when 13 scenarios are needed (32). Furthermore, the coefficient of variation in the optimal set of scenario probabilities is about 11% for 20 and about 25% for 13. This indicates that the solution from step 1 is significantly closer to the optimal solution for 20 scenarios than the initial solution for 13 or fewer scenarios.




Figure 5: Memphis highway network, traffic analysis zones and bridges (squares)




Figure 6: Error and iterations needed as a function of the number of scenarios desired

It is useful to compare the magnitude of these errors to those that would result from a Monte Carlo simulation strategy. It is common to assume that there is no correlation in the damage across bridges given the ground shaking (FEMA 2010). Therefore, a Monte Carlo sample is obtained by simply drawing an observation from the probability distribution for damage for each bridge k generated by the loss estimation methodology (m_{kd}). Under this structure, each scenario generated is treated as equally likely.

A hybrid of Monte Carlo simulation and the optimization described above can also be developed with the goal of producing a better match with the probabilities m_{kd} , permitting the use of fewer consequence scenarios. This is done by using the optimization described by equations (1)-(7). The b_{jkd} are fixed based on the generated Monte Carlo samples and the optimization determines the probability of each scenario, s_j . Since the binary variables b_{jkd} are fixed, the resultant optimization is a linear program.

Figure 7 illustrates how well both Monte Carlo methods match the probabilities m_{kd} . One hundred replicates were performed for various numbers of scenarios, the average sum of errors is reported in Figure 7. In both cases, it is useful to notice that with a sample size of 500 there are larger discrepancies from the probabilities m_{kd} than the 9 scenarios identified by the optimization (sum of errors of about 10 for the optimization and an average of about 18 and 12 for the equal weight and optimized weight Monte Carlo methods, respectively).

Figure 7 also illustrates the performance of hybrid Monte Carlo sampling and optimization. The ability to optimize the probability of each scenario generated by Monte Carlo simulation results in a larger percentage improvements in the errors as the sample size rises. For example, when there are only 10 scenarios, the reduction in the errors via optimization of the probabilities is about 7% (125 in comparison to 117). However, when the number of scenarios is 500, that reduction is about 32% (18 in comparison to 12).

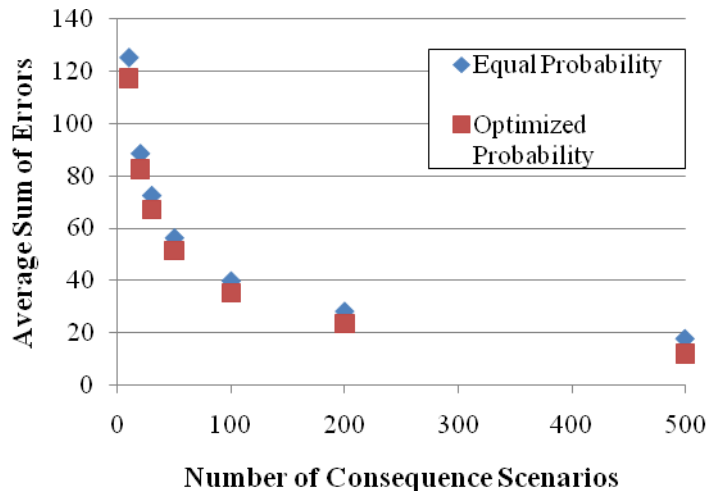


Figure 7: Average Monte Carlo sampling errors as a function of the number of scenarios for the equal probability and optimized probability methods

These gains in performance achieved by the optimization do come at a “price” as measured by the variance and covariance in the damage between pairs of bridges. As mentioned previously, it is commonly assumed that the correlation in the damage of two different bridges is zero (once the ground shaking has been considered). This assumption is commonly made, however, there is little understanding of how accurate this assumption is in practice. Since it is the common assumption, we explore the degree of departure from this assumption in this example using these methods. None of the methods described in this paper include this as an explicit constraint, though with Monte Carlo simulation, as long as the sample size is large enough, this will be true.

In order to analyze the correlations in the damage across bridges given the ground shaking from a given procedure, we define a random vector $\zeta = (\zeta_k; k = 1, \dots, 335)$, where $\zeta_k = 0,1,2,3,4$. We let $\zeta_k = 0,1,2,3,4$ indicate bridge k 's damage state to be no damage, slight damage, moderate damage, extensive damage, and complete damage, respectively. Then for the given set of scenarios and associated probabilities, we calculate the covariance in damage state between two bridges as follows:

$$cov(\zeta_k, \zeta_{k'}) = E(\zeta_k - E\zeta_k)(\zeta_{k'} - E\zeta_{k'}) \quad \forall k, k'. \tag{8}$$

Analytic calculations show that the sum of the variances in the damage across all bridges is about 444. Figure 8 gives the sum of the variance across all bridges as a function of the number of scenarios generated by the optimization and Monte Carlo methods. For the Monte Carlo case, 100 replicates were used to estimate this quantity, hence the mean and the 5th and 95th percentiles are reported. Our estimates of the confidence limits have some additional variability because we did not use common random numbers in the sample generation. Except for the case with 7 consequence scenarios, the sample means are not statistically different from the analytic value ($\alpha = 0.05$). As the number of scenarios increases the variability around each sample decreases. For the optimized method, the sum of the variances is initially below the calculated value but approaches the correct value with an increasing number of consequence scenarios.

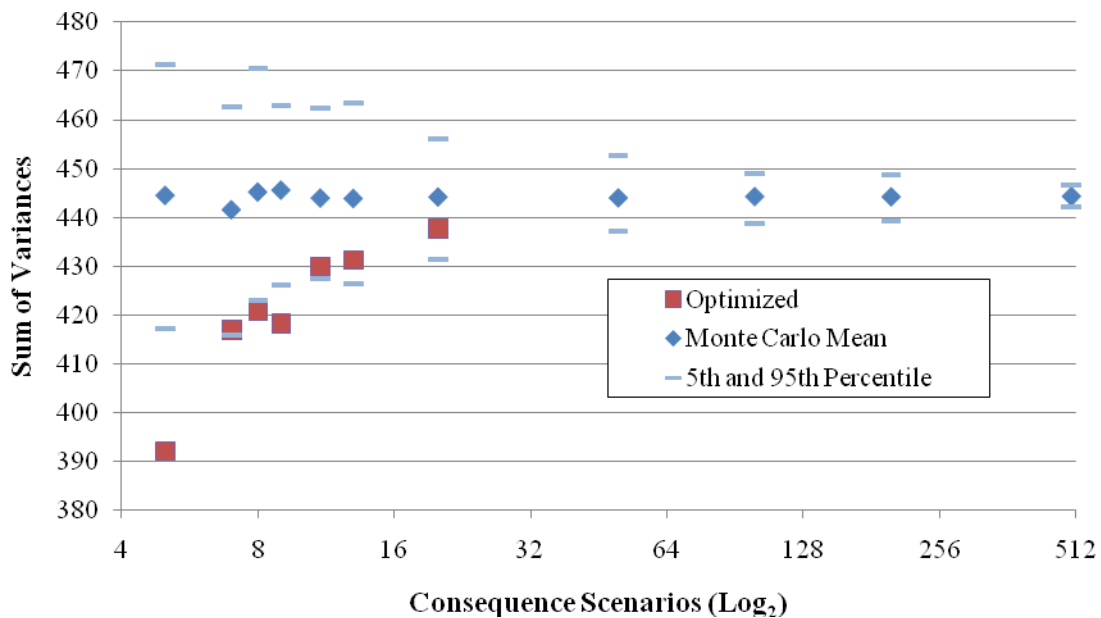


Figure 8: Total variance in the optimized and Monte Carlo methods as a function of the number of scenarios desired

Figure 9 is similar to Figure 8 but gives the sum of the absolute value of the covariance in place of the total variance. The Monte Carlo case is consistent with theory, as more scenarios are added the total deviations in covariance reduce to zero. It is important to view these data in the context of the over 111,000 pairs of bridges which are considered in this sum. In the case where the total covariance is about 33,000 the covariance for each pair of bridges is on the order of 0.3. For the optimized method the covariance follows the same trend but the total covariance is greater. When 13 and 20 scenarios are desired, the covariance is about 18% and 22% greater using optimization than the Monte Carlo approach, respectively.

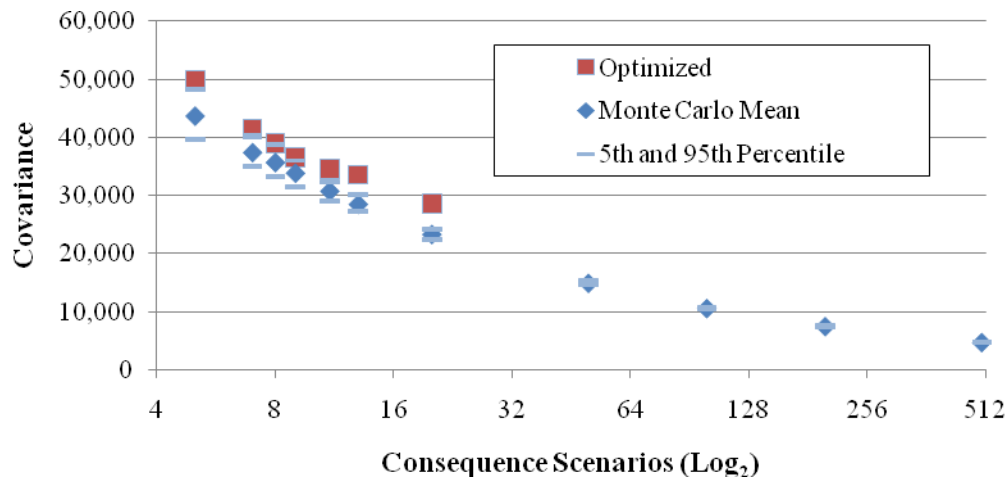


Figure 9: Total covariance in the optimized and Monte Carlo methods as a function of the number of scenarios desired

It is important to recall that the discrepancies with damage state probabilities are drastically reduced when optimization is used. The deviations for optimization compared to the Monte Carlo methods are 6.8 versus 125 for 13 scenarios and 6.4 versus 83 for 20 scenarios. While the optimization approach does introduce more covariance than Monte Carlo simulation, the increase is not drastic and it decreases as the number of scenarios increases. Furthermore, the importance of increased covariance should be balanced by the fact that the actual correlation between bridges is not well understood and independence is a simplifying assumption.

4 CONCLUSIONS

This paper presents a new optimization approach to identify consequence scenarios to estimate the impact of an earthquake event on lifeline systems. A novel nonlinear integer program is developed and used to identify a small set of consequence events and the probability of occurrence for each, such that together they match the damage probabilities for each component as given in the loss estimation modeling (e.g., HAZUS 2010). This set of events will allow system performance to be estimated with a high degree of accuracy while keeping the computational burden to a minimum.

There are opportunities for further research in at least a couple of areas. This case study focused on developing consequence scenarios for a single event, which is of significant value. However, if the goal is to use these scenarios to understand the seismic vulnerability of an area, it is important to extend the optimization to consider simultaneously identifying consequence scenarios for the range of events (and the hazard-consistent probability of occurrence for each) that represent the seismic vulnerability of the area. Clearly this optimization could be run for each of the events separately, but that may make it hard to identify how many consequence scenarios should be used for each event when there is a maximum total number of scenarios which can be included in further analyses. It would be better if that determination could be done via optimization.

Second, if the ultimate goal of the modeling is to support decision-making for seismic retrofit of structures, it is important that the scenarios identified specify how those structures would change under each consequence scenario, based on the mitigation performed. Since a loss analysis can be performed assuming each type of mitigation is performed on each structure, it is possible to augment the constraint set in the optimization described above to include this information so that each consequence scenario reports the condition of each bridge with no mitigation, as well as with each feasible package of retrofit options.

ACKNOWLEDGEMENTS

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND 2011-2150C.

REFERENCES

- Çağnan, Z., R. Davidson, and S. Guikema. 2006. "Post-Earthquake Restoration Planning for Los Angeles Electric Power." *Earthquake Spectra* 22(3):1-20.
- Chang, S., M. Shinozuka, and J. Moore. 2000. "Probabilistic Earthquake Scenarios: Extending Risk Analysis Methodologies to Spatially Distributed Systems." *Earthquake Spectra* 16(3):557-572.
- Dodo, A., N. Xu, R. Davidson, and L. Nozick. 2005. "Optimizing Regional Earthquake Mitigation Investment Strategies." *Earthquake Spectra* 21(2):305-327.
- FEMA (Federal Emergency Management Agency). 2010. *HAZUS User and Technical Manual*, Department of Homeland Security. Accessed June 14, 2011. http://www.fema.gov/plan/prevent/hazus/hz_manuals.shtm.
- Frankel, A. D., D. Applegate, M. P. Tuttle, and R. A. Williams. 2009. *Earthquake Hazard in the New Madrid Seismic Zone Remains a Concern*. U.S. Geological Survey Fact Sheet 2009-3071.
- Gomberg, J., and E. Schweig. 2007. *Understanding Earthquake Hazards in the Central United States. Earthquake Hazard in the Heart of the Homeland*. U. S. Geological Survey Fact Sheet 2006-3125.
- Jayaram, N., and J. Baker. 2010. "Efficient Sampling and Data Reduction Techniques for Probabilistic Seismic Lifeline Risk Assessment." *Earthquake Engineering and Structural Dynamics* 39:1109-1131.
- Kunreuther, H., E. Michel-Kerjan, N. Doherty, M. Grace, R. Klein, and M. Pauly. 2009. *At War with the Weather: Managing Large-Scale Risks in a New Era of Catastrophes*. Cambridge: The MIT Press, Cambridge.
- Petersen, M. D., A. Frankel, C. Harmsen, C. Mueller, K. Haller, R. Wheeler, R. Wesson, Y. Zeng, O. Boyd, H. Perkins, N. Luco, E. Field, C. Wills, and K. Rukstales. 2008. *Documentation for the 2008 Update of the United States National Seismic Hazard Maps*. U. S. Geological Survey Open-File Report 2008-1128.
- State of California Department of Conservation. 2011. "Northridge Earthquake 10 Year Anniversary." Accessed June 14, 2011. http://www.consrv.ca.gov/cgs/geologic_hazards/earthquakes/Pages/northridge.aspx.
- Sun-Sentinel. 2011. "Ten Years After Hurricane Andrew, Effects are Still Felt." Accessed June 14, 2011. <http://www.sun-sentinel.com/sfl-1992-ap-mainstory,0,7290462.story>.
- Tuttle, M. P., E. S. Schweig, J. D. Sims, R. H. Lafferty, L. W. Wolf, and M. L. Hayes. 2002. "The Earthquake Potential of the New Madrid Seismic Zone." *Bulletin of the Seismological Society of America* 92(6):2080 – 2089.
- USA Today. 2005. "Katrina Damage Estimate Hits \$125B." Accessed June 14, 2011. http://www.usatoday.com/money/economy/2005-09-09-katrina-damage_x.htm.
- USC (University of Southern California). 2011. "The Great California ShakeOut." Accessed June 14, 2011. <http://www.shakeout.org/>.
- USGS (United States Geological Survey). 2008. *2008 NSHM Software*. U.S. Geological Survey Website. Accessed June 14, 2011. <http://earthquake.usgs.gov/hazards/products/conterminous/2008/software/>.
- Xu, N., S. Guikema, R. Davidson, L. Nozick, Z. Çağnan, and P. Vaziri. 2007. "Optimizing Scheduling of Post-Earthquake Electric Power Restoration Tasks." *Earthquake Engineering and Structural Dynamics* 36:265-284.

AUTHOR BIOGRAPHIES

NATHANAEL J.K. BROWN is a Software/Electrical Engineer at Sandia National Laboratories in Albuquerque, New Mexico. He has a Masters in Electrical Engineering from Purdue University and has worked at various companies including Intel Corporation and Compaq Computer Corporation. The main focus of his career has been on the design and development of advanced software algorithms ranging from handwriting recognition to digital filtering. His e-mail address is njbrown@sandia.gov.

JARED L. GEARHART is an Operations Research Analyst at Sandia National Laboratories in Albuquerque, New Mexico. He has a Masters of Engineering in Applied Operations Research from Cornell University. His work interests are in developing simulation and optimization models, and solution procedures for complex optimization problems. His e-mail address is jlgearh@sandia.gov.

DEAN A. JONES is a Distinguished Member of Technical Staff at Sandia National Laboratories. He is the principal technical and programmatic lead of the Operations Research & Computational Analysis Group (ORCA). His emphasis is on the research, design and application of mathematical models for use in the analysis of complex systems with particular emphasis on systems that can be represented mathematically as graph-based networks. His e-mail address is dajones@sandia.gov.

LINDA K. NOZICK is a Professor of Civil and Environmental Engineering at Cornell University. Her Ph.D. is in Systems Engineering from the University of Pennsylvania. Her research interests are in the modeling of complex systems with a particular focus on infrastructure systems and resiliency to natural and man-made hazards. Her email address is lkn3@cornell.edu.

NATALIA ROMERO is a Ph.D. student in the School of Civil and Environmental Engineering at Cornell University. Her research interests are in the modeling of civil infrastructure systems under natural and man-made extreme events. Her e-mail address is nr229@cornell.edu.

NINGXIONG XU is a Research Associate in the School of Civil and Environmental Engineering at Cornell University. He earned a Ph.D. in Operations Research from Stanford University. His research interests are in large scale optimization under uncertainty and the application of those tools to transportation systems. His e-mail address is nx22@cornell.edu.