Efficient sampling techniques for seismic risk assessment of lifelines

N. Jayaram & J. W. Baker
Dept. of Civil and Environmental Engg., Stanford University, CA.

ABSTRACT: Seismic risk assessment of lifelines is considerably more complicated than that of a single structure on account of the geographical spread of lifelines. Lifeline risk assessment requires knowledge about ground-motion intensities at multiple sites. Further, the link between the ground-motion intensities and lifeline performance is typically not available in closed form. These complications render analytical risk assessment tools insufficient for a probabilistic study of lifeline performance. The current study proposes an efficient simulation-based lifeline risk assessment framework based on importance sampling (IS) . In this framework, 'important' ground-motion fields are preferentially sampled, and their impacts on lifeline performance are studied. Important ground-motion fields are generated by preferentially sampling large magnitude earthquakes and above-average ground motions corresponding to the sampled earthquakes. The study proposes IS density functions that can be used for such preferential sampling. The study also suggests techniques that can be used to estimate the parameters of these sampling densities.

The proposed IS framework is used to evaluate the seismic risk of an aggregated form of the San Francisco bay area transportation network. The accuracy of the proposed method is demonstrated by showing that the risk estimates obtained using the IS framework match those obtained using random Monte Carlo simulation (MCS). It is also shown that the IS approach can produce risk estimates that are comparable in accuracy to results from the MCS approach while using roughly one-hundredth the number of realizations . Finally, the study shows that the uncertainties in the ground-motion intensity and the spatial correlations between ground-motion intensities at various sites must be modeled in order to obtain unbiased estimates of lifeline risk.

1 INTRODUCTION

Lifelines are large, geographically-distributed systems such as transportation networks that are essential support systems for any society. In the past, events such as the 1989 Loma Prieta earthquake and the 1994 Northridge earthquake have exposed the seismic vulnerability of various lifelines. For instance, the Northridge earthquake caused over $1.5 billion in business interruption losses ascribed to transportation network damage (Chang 2003). The city of Los Angeles suffered a power blackout and $75 million of power-outage related losses as a result of the earthquake (Kim et al. 2007). Therefore, it is important to proactively assess and mitigate the seismic risk of lifelines. Traditionally, probabilistic seismic hazard analysis (PSHA) combined with the analytical PEER loss analysis framework has been used to estimate the hazard at a single site and to assess probable losses to the structure at the site using the estimated hazard (McGuire 2007). The site hazard is obtained using ground-motion models, which predict median ground-motion intensities as well as dispersion of the intensities about the median values (e.g., Boore & Atkinson 2008). The PEER methodology cannot, however, be used for lifeline risk assessment, since lifeline risk assessment is based on a large vector of ground-motion intensities (intensities at all lifeline component locations), and because the link between the ground motions at the sites and the performance of the lifeline is usually not available in closed form. Further, while obtaining joint hazard predictions (i.e., prediction of a vector of intensities), it is important to model the spatial correlation between the ground-motion intensities at various sites in order to accurately assess the seismic risk (Park et al. 2007). As a result of these complexities, many past research works use simulation-based approaches instead of analytical approaches for lifeline risk assessment (e.g., Crowley & Bommer 2006, Kiremidjian et al. 2007, Shiraki et al. 2007). One simple simulation-based approach involves studying the performance of lifelines under those scenario earthquakes that are assumed to dominate the hazard in the region of interest (e.g., Adachi & Ellingwood, 2008). While this approach is less computationally demanding, it does not capture the uncertainties in the seismic hazard the way a PSHA-based framework would. A more rigorous approach uses Monte Carlo simulation (MCS) to probabilistically generate ground-motion fields, considering all possible earthquake scenarios that could occur in the region, and then use these for the risk assessment. The methodology used for generating ground-motion fields is based on a form of existing ground-motion models, which is described below. We model ground motion at a site as

$$ln\left(Y_{ij}\right) = ln\left(\bar{Y}_{ij}\right) + \varepsilon_{ij} + \eta_j \quad (1)$$

where Y_{ij} denotes the ground-motion parameter of interest (e.g., $S_a(T)$, the spectral acceleration at period T) at site i during earthquake j; $\bar{Y}_{ij}$ denotes the predicted (by the ground-motion model) median ground-motion intensity (which depends on parameters such as magnitude, distance, period and local-site conditions); ε_{ij} denotes the intra-event residual, which is a random variable with zero mean and standard deviation σ_{ij}; and η_j denotes the inter-event residual, which is a random variable with zero mean and standard deviation τ_j. The standard deviations, σ_{ij} and τ_j, are estimated as part of the ground-motion model.

Crowley & Bommer (2006) describe the MCS approach used to probabilistically sample ground-motion fields. This approach involves simulating earthquakes of different magnitudes on various active faults in the region, followed by simulating the inter-event and the intra-event residuals at the sites of interest for each earthquake. The residuals are then combined with the median ground motions in accordance with Equation 1 in order to obtain the ground motions at all the sites.

Crowley & Bommer (2006) used the above-mentioned approach to generate multiple earthquake scenarios that were then used for the loss assessment of a portfolio of buildings. They found that the results varied significantly from those obtained using other approximate approaches (e.g., using PSHA to obtain individual site hazard and loss recurrence curves, which are then combined to obtain the overall loss recurrence curve). They, however, ignored spatial correlations while simulating ground-motion fields. Further, they used conventional MCS, which is computationally inefficient since it implicitly ascribes equal importance to all ground-motion scenarios irrespective of their potential impact on the lifeline risk. This is inefficient because large magnitude events and above-average ground motions are considerably more important than small magnitude events and small ground motions while modeling lifeline risks, but these events are infrequently sampled in conventional MCS. Kiremidjian et al. (2007) improved the simulation process by preferentially simulating large magnitudes using importance sampling (IS). The residuals, however, were simulated using conventional MCS.

In the literature, there are also non-MCS-based approaches such as that of Kang et al. (2008) (matrix-based system reliability approach) and Dueñas-Osorio et al. (2005) (graph-theory based), which have been used for performance assessment of lifelines. It is, however, difficult to compute exceedence curves for many preferred lifeline performance measures (e.g., the total travel time in a transportation network) using non-MCS-based approaches.

The current research work develops an IS-based framework to efficiently simulate magnitudes and residuals so as to substantially improve the computational efficiency of the simulation procedure. The correlated ground-motion fields generated using this approach are then used for assessing the seismic risk of an aggregated form of the San Francisco bay area transportation network. The resulting risk estimates are shown to be in excellent agreement with those obtained using the conventional MCS approach (the benchmark method). The overall IS framework is shown to be computationally faster than the MCS framework by a factor of 100.

2 SIMULATION OF CORRELATED GROUND-MOTION FIELDS

As mentioned previously, simulating correlated ground-motion fields involves probabilistically sampling earthquake magnitudes and rupture locations (which are required for computing the median ground-motion intensities), the inter-event residuals and the intra-event residuals (Equation 1). This section provides a detailed description of the sampling procedure used in the current work.

2.1 *Generating an earthquake catalog*

Let n denote the number of active faults in the region of interest and ν_i denote the annual recurrence rate of earthquakes on fault i. Let $f_i(m)$ represent the density function for earthquake magnitudes on fault i. Let $f(m)$ denote the density function of the magnitude of any earthquake occurring in the region of interest (i.e., this density function models the distribution of earthquakes resulting from all faults). Typically the distributions $f_i(m)$ are known, but $f(m)$ will be needed later for this procedure. Using the law of total probability, $f(m)$ can be determined as follows:

$$f(m) = \sum_{i=1}^{n} \nu_i f_i(m) \Big/ \sum_{i=1}^{n} \nu_i \quad (2)$$

In the event of an earthquake of magnitude m on a random fault, let p_i denote the probability that the earthquake is due to a rupture on fault i. The p_i's can be calculated using Bayes' theorem as follows:

$$p_i = \nu_i f_i(m) \Big/ \sum_{i=1}^{n} \nu_i f_i(m) \quad (3)$$

Conventional MCS based on Equations 2 and 3 can be used for simulating earthquake magnitudes (using *f(m)*) and rupture locations (using p_i to randomly sample rupture sources).

The drawback of simulating magnitudes directly using the density *f(m)*, however, is that most simulated magnitudes will be small since small magnitude events are considerably more probable than large magnitude events (say, greater than 7) (Fig. 1a). Lifeline losses due to small events are usually negligible. Moreover, lifeline losses are often not sensitive to minor variations in the small-magnitude events. Therefore, it is possible to improve the computational efficiency of the risk assessment process without compromising on the accuracy of the estimates by using a simulation procedure that preferentially samples large events (while still ensuring that the simulated events are seismically representative). This is an application of the importance sampling technique, which is used to obtain the properties of a probability density function *f(m)* using samples from an alternate probability density function *g(m)* (Law, 2007). The function *g(m)* should be chosen to have a high probability of producing samples from the regions of interest (for instance, if *g(m)* is used to simulate magnitudes, *g(m)* at large magnitudes (*m*) should be large (e.g., Kiremidjian et al. 2007)).

In the current work, based on sensitivity analysis, the authors choose the alternate sampling density function, *g(m),* to be a truncated inverse-exponential distribution, as shown below.

$$g(m) = \lambda \exp\{\lambda(m - m_{min})\} \big/ \exp\{\lambda(m_{max} - m_{min})\} \quad (4)$$

where λ is the parameter of the sampling distribution, m_{min} and m_{max} are the minimum and maximum magnitudes of interest respectively. The original and the sampling density functions are shown in Figure 1a. The value of λ that defines *g(m)* needs to be chosen so that the simulated magnitudes match the desirable set of magnitudes. Choosing an optimal λ is discussed in a subsequent section.

2.2 *Simulating the intra-event residuals*

The set of intra-event residuals $\boldsymbol{\varepsilon_j} = \left(\varepsilon_{1j}, \varepsilon_{2j}, ..., \varepsilon_{dj}\right)$ follows a multivariate normal distribution (Jayaram & Baker 2008a). The mean of $\boldsymbol{\varepsilon_j}$ is the zero vector of size *d* (where *d* is the number of sites of interest), while the variances of the residuals can be obtained from the ground-motion model. The correlation between the residuals at two sites is typically a function of the separation between the sites, and can be obtained from a spatial correlation model. In this work, the correlation coefficient between the residuals at two sites separated by *h* km is computed using the following equation, which was calibrated using empirical observations by Jayaram & Baker (2008b, in review).

$$\rho(h) = exp(-3h/30) \quad (5)$$

where the factor 30 controls the rate of decay of spatial correlation and is called the "range" of the correlation model.

Let *f(e)* denote the above-mentioned multivariate normal distribution. A set of correlated intra-event residuals can be simulated from *f(e)* by first generating a set of independent residuals (e.g., using the Box-Muller method), and by incorporating the necessary correlation (using the Choleskey triangle) on to the independent residuals (Law 2007). Direct simulation using *f(e)* (i.e., conventional MCS) will, however, result in a large number of near-zero (i.e., near-mean) residuals and few realizations from the upper and the lower tails. For the purposes of lifeline risk assessment, it is often of interest to study the upper tail (i.e., the ε_{ij} values that produce the largest ground motions). In the current work, this is done using IS, where the alternate sampling density *g(e)* is chosen to be a multivariate normal distribution with the same variance and correlation structure of *f(e)*,

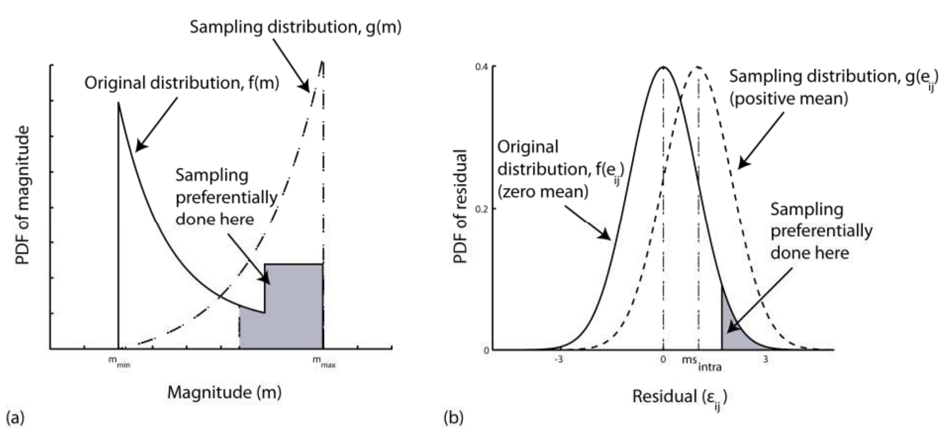

Figure 1: Importance sampling density functions for: (a) magnitude and (b) intra-event residual

but with positive means for the marginal distributions. In other words, the mean vector of *g(e)* is no longer a zero vector of size *d*, but the vector (ms_{intra}, ms_{intra}, ..., ms_{intra}). The standard deviations and correlations of ε_{ij}'s are left unchanged in the sampling distribution. Figure 1b shows the marginal original and sampling distributions for one particular ε_{ij}. It is to be noted that this choice of the sampling distribution results in importance sampling weights (Law 2007) that are extremely simple and computationally efficient to compute. The positive mean of *g* will ensure that the realizations from *g* will be larger than the realizations from *f*. It is, however, important to choose an optimal value of the mean-shift ms_{intra} to ensure adequate preferential sampling of large ε's, while avoiding sets of extremely large intra-event residuals that will make the scenario ground-motion field improbable. The process of selecting an optimal value of ms_{intra} is described in a subsequent section.

2.3 *Simulating the inter-event residual*

The density function of the inter-event residual (η) is as follows:

$$f(\eta) \sim N(0,\tau) \tag{6}$$

where $f(\eta)$ denotes the density function of η, $N(0,\tau)$ is the univariate normal distribution with mean 0 and standard deviation τ (obtained from the ground-motion model).

The IS procedure for η is similar to that for ε, except that the alternate sampling distribution is univariate normal rather than multivariate normal, and has standard deviation τ and a positive mean ms_{inter}.

3 APPLICATION: SAN FRANCISCO BAY AREA TRANSPORTATION NETWORK

In this section, the San Francisco bay area transportation network is used to demonstrate the applicability of the simulation framework developed in section 2 for assessing the seismic risk of lifelines. It is intended to show that the seismic risk estimated using the IS framework matches with a benchmark estimate of seismic risk obtained using the MCS framework.

3.1 *Network data*

The San Francisco bay area transportation network data used in the current study were obtained from Stergiou & Kiremidjian (2006). Figure 2a shows the Metropolitan Transportation Commission (MTC) San Francisco bay area highway network, which consists of 29,804 links and 10,647 nodes. The network also consists of 1,125 bridges from the five counties of the bay area. Stergiou & Kiremidjian (2006) classified these bridges based on their structural properties in accordance with the HAZUS (1999) manual. This classification is useful for estimating the structural damage to bridges due to various simulated scenario earthquakes. The bay area network consists of a total of 1,120 transportation analysis zones (TAZ), which are used to predict the trip demand in specific geographic areas. The origin-destination data provided by Stergiou & Kiremidjian (2006) were obtained from the 1990 MTC household survey.

This application analyzes the network performance under ground-motions generated using both IS and MCS. Analyzing the performance of a network as large and comprehensive as the San Francisco bay area transportation network under scenario ground-motions generated, in particular, using MCS is extremely computationally intensive. Therefore, in the current study, an aggregated representation of the bay area network is used for the illustration. The aggregated network consists predominantly of the region's freeways and expressways, along with the ramps linking the freeways and expressways. The nodes are placed at locations where links intersect or

(a) (b)

Figure 2: (a) San Francisco bay area transportation network (b) Aggregated network (the dots on the links represent bridge locations)

change in characteristics (e.g., change in the number of lanes). The aggregated network comprises of 586 links and 310 nodes (Fig. 2b) (The locations of the 1,125 bridges are also shown in this figure.) Of the 310 nodes, 46 nodes are denoted centroidal nodes that act as origins and destinations for the traffic.

3.2 *Ground-motion hazard*

In the current study, the San Francisco bay area seismicity information is obtained from USGS (2003). Ten active faults/ fault segments are considered in the current work. The characteristic magnitude recurrence relationship of Youngs & Coppersmith (1985) is used to model $f_i(m)$, with upper magnitude thresholds specified by the USGS. For the purposes of our application, 5.0 is considered to be the lower threshold for the magnitudes of interest. The ground-motion model of Boore & Atkinson (2008) is used to obtain the median ground motions and the standard deviations of the residuals needed in Equation 1.

3.3 *Transportation network performance measure*

A popular measure of network performance is the total travel time of passengers in a network (Stergiou & Kiremidjian 2006, Shiraki et al. 2007). The total travel time (T) can be expressed as follows:

$$T = \sum_{i \in links} x_i t_i \left(x_i \right) \quad (7)$$

where x_i denotes the traffic flow on link i and $t_i(x_i)$ denotes the travel time of a passenger on link i.

Travel times on transportation networks are generally computed using the user-equilibrium principle (Beckman et al. 1956), which states that each individual user would follow the route that would minimize his/ her travel time. The travel times under equilibrium are obtained using the commonly-used solution technique provided by Frank & Wolfe (1956).

3.4 *Post-earthquake network performance*

This section describes the process of determining the performance of a transportation network immediately after an earthquake. In the current work, simplified models of transportation systems are adopted for purposes of demonstrating the feasibility of the proposed IS framework. It is to be noted that more realistic network models can be used within the proposed framework, if desired.

Research works such as Kiremidjian et al. (2003) and Cho et al. (2001) focus on a comprehensive treatment of post-earthquake changes in travel demands, which will affect the travel times in the network. The current work, however, assumes for simplicity that the post-earthquake demands equal the pre-earthquake demands. Therefore, the changes in network performance after an earthquake are only due to the structural damage to bridges. The damage states of the bridges are computed considering only the ground-motion hazard (i.e., other possible damage mechanisms such as liquefaction are not considered). The bridge fragility curves provided by HAZUS (1999) are used to estimate the probability of a bridge being in a particular damage state (no damage, minor damage etc.) based on the simulated ground-motion intensity at the bridge site. These damage state probabilities are then used to simulate the damage state of the bridge following the earthquake. Damaged bridges result in reduced capacities and increased free-flow travel times in the links in which the bridges are part of. The capacity reduction and the increase in travel times due to bridge damage were obtained based on Shiraki et al. (2007).

Ideally, each bridge should be considered as part of a separate link, though in practice this could become cumbersome. In particular, aggregated networks can have up to 10 bridges in a single link, and in such cases, the link damages need to be related to the damage to all the bridges present in the link. As a simplification, the current work assumes that the link capacity reduction equals the average of the capacity reductions attributable to each bridge in the link. A similar approach is used to estimate the increase in the link free-flow travel time. The post-earthquake network performance is then computed using the new network characteristics (i.e., the user-equilibrium problem is solved using the new set of free-flow travel times and capacities in order to obtain the link flows), and a new estimate of the total travel time in the network is obtained. It is to be noted that the current work estimates the performance of the network only immediately after an earthquake. The changes in the performance with network component restorations are not considered here for simplicity.

4 RESULTS AND DISCUSSION

This section discusses the San Francisco bay area transportation network risk estimates obtained using the proposed IS framework. The IS framework requires that the parameters of the sampling distribution for the magnitude and the residuals (described in Section 2) be chosen optimally in order to obtain reliable results efficiently. The set of parameters includes λ (Equation 4) for magnitudes, the mean-shift for inter-event residuals (ms_{inter}) and the mean-shift for intra-event residuals (ms_{intra}).

The value of λ was fixed at 1.0 since it produced a desirable earthquake magnitude histogram with a sufficient number of large magnitude events. The value of ms_{inter} can be fixed similarly, by viewing the histogram of the sampled residuals. In this study, ms_{inter} was chosen as 1.0. Choosing an optimal value of ms_{intra} is slightly more difficult since ms_{intra} affects the joint distribution of a set of intra-event residuals, and hence, is not easy to visualize. The first step in fixing the value of ms_{intra} is to note that the optimal value depends predominantly on three factors, namely, the extent of spatial correlations (measured by the range parameter in Equation 5), the average site-to-site separation distance and the number of sites of interest. As mentioned previously, large mean-shifts will result in inefficiency due to the simulation of many improbable ground-motion fields (ground-motion values more extreme than the domain of interest). If sites are close to one another and if the spatial correlations are significant, the correlations between the residuals will permit a larger mean-shift as it is reasonably likely to observe jointly large values of positively-correlated random variables. Similarly, the presence of fewer sites permits larger mean-shifts since it is more likely to observe jointly large values of residuals over a few sites than over a large number of sites. Hence, in this study, it is intended to determine optimal mean-shifts as a function of average site-to-site separation distances normalized by the range and the number of sites. This is done by simulating the intra-event residuals over several fields that contain a varying number of sites with varying average site separation distances, considering several feasible mean-shifts for each field. The feasibility of the resulting residuals (i.e., whether the simulated set of residuals is reasonably probable) is then studied, and the optimal mean-

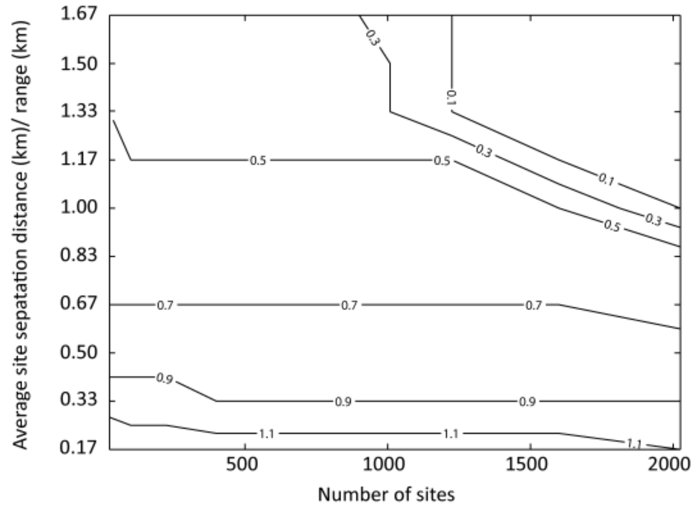

Figure 3: Optimal mean-shift, ms_{intra}, as a function of the average number of sites and the average site to site distance

shifts were determined for each case based on the results. These optimal mean-shifts have been plotted as a function of the number of sites and the average site separation distance/ range in Figure 3. In the current work, for the aggregated bay area network in consideration, the value of ms_{intra} was fixed at 0.3.

In this paper, the lifeline risk estimates are presented in the form of an exceedence curve, which shows the recurrence rates of various possible post-earthquake total travel times. Figure 4a shows the exceedence curve (risk curve) for total travel times obtained using the IS framework. This risk curve is obtained by sampling 25 magnitude-rupture location pairs and 50 sets of inter and intra-event residuals for each magnitude-location pair.

It is important to validate the risk curve shown in Figure 4a in order to ensure the effectiveness of the proposed IS framework. In order to do so, a risk curve is also estimated using the benchmark method, namely, MCS, and is compared to that obtained using IS. Strictly, the benchmark approach should use MCS to sample the catalog of earthquakes (i.e., magnitudes and rupture locations) and the ground-motion residuals. Experience indicates, however, that reliable risk estimates can only be obtained us-

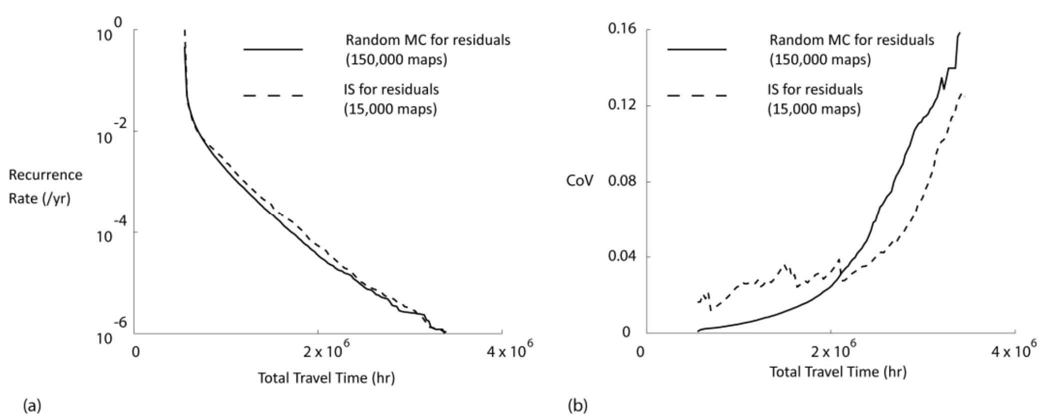

Figure 4: (a) Total travel time exceedence curves (b) Coefficient of variation of the risk estimates

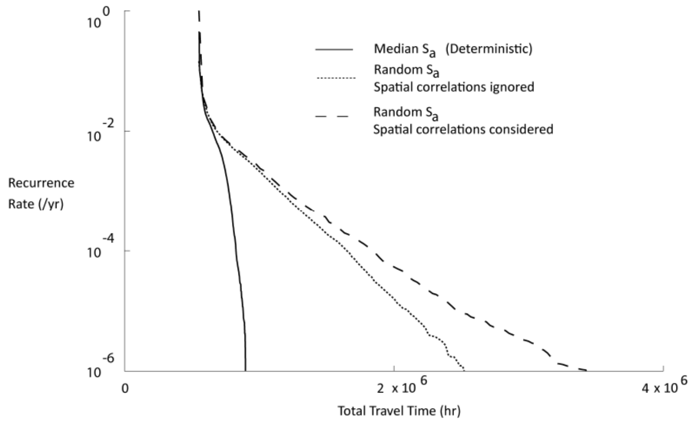

Figure 5: Exceedence curves obtained using simplifying assumptions

ing approximately a million ground-motion maps generated using this approach. This is computationally prohibitive and hence, the benchmark approach used in the current study uses IS for generating the magnitudes but MCS for the residuals. This is reasonable because the IS of a single random variable (magnitude in this case) has been shown to be effective in a wide variety of applications. In fact, Kiremidjian et al. (2003) use IS in place of MCS for generating magnitudes as well. Further applications of IS can be found in areas such as structural reliability (Melchers, 1999). On the other hand, the simulation procedure for intra-event residuals involves IS of a correlated vector of random variables, and hence, is the focus of the validation study described in this section.

Figure 4a shows the exceedence curve obtained using IS for generating magnitudes (25 magnitude-rupture location pairs) and MCS for generating inter and intra-event residuals (500 sets per magnitude-location pair). As seen from the figure, the risk curves obtained using the IS framework closely matches that obtained using the benchmark method, thereby indicating the accuracy of the results obtained using IS. This is further substantiated by Figure 4b, which plots the coefficient of variation (CoV) of the risk estimates obtained using the IS approach and the benchmark approach. It can be seen from the figure that the CoV values corresponding to large travel times are smaller when IS is used, even though the IS uses one-tenth the number of simulations required by the benchmark method. Further, it was also seen that using IS in the place of MCS for simulating magnitudes typically reduces the computational complexity by a factor of 10, and hence, the overall IS framework reduces the CPU time for the risk assessment by a factor of nearly 100.

Finally, in this study, it is also intended to demonstrate the importance of considering spatial correlation while assessing lifeline risk. Hence, the transportation network risk assessment was repeated assuming uncorrelated intra-event residuals, and a new exceedence curve was obtained (Fig. 5). It can be seen that the risk is considerably underestimated when the spatial correlations are ignored. Further, some past risk assessments in the literature have completely ignored the aleatory uncertainty in the residuals (i.e., median ground-motion fields are assumed, and inter- and intra-event residuals are ignored). A risk assessment carried out this way (Fig. 5) shows that the risk is even more substantially underestimated in this case. Such simplifications clearly introduce significant errors into the risk calculations, and should thus be avoided.

5 CONCLUSIONS

An efficient simulation-based framework based on importance sampling (IS) has been proposed in this study that can be used for the seismic risk assessment of lifelines. The risk assessment procedure involves preferentially sampling ‘important’ ground-motion fields, and evaluating the lifeline performance considering the simulated fields. Important ground-motion fields are generated by preferentially sampling large magnitude earthquakes and above-average ground motions corresponding to these earthquakes. The study proposed IS density functions that can be used for such preferential sampling, and also suggested techniques that can be used to estimate the parameters of these sampling densities.

The proposed IS framework was used to evaluate the seismic risk of an aggregated form of the San Francisco bay area transportation network. Simplified transportation network analysis models were used to illustrate the feasibility of the proposed framework. The study showed that the risk estimates obtained using IS match those obtained using the benchmark method, namely, conventional Monte Carlo simulation (MCS). It was also shown that the number of required IS realizations is roughly one-hundredth the number of required MCS realizations for obtaining equally reliable risk estimates. Finally, the study showed that the uncertainties in the ground-motions and the spatial correlations between ground-motion intensities at multiple sites must be modeled in order to avoid introducing significant errors into the lifeline risk calculations.

6 ACKNOWLEDGMENTS

The authors thank Prof. Kiremidjian from Stanford University for useful discussions on this research topic. The authors also thank Evangelos Stergiou and Prof. Kiremidjian for proving data pertaining to the San Francisco bay area transportation network.

7 REFERENCES

Adachi, T. & Ellingwood, B.R. 2008. Serviceability of earthquake-damaged water systems: Effects of electrical power availability and power backup systems on system vulnerability. *Reliability engineering and system safety* 93: 78–88.

M.J. Beckman, C.B.McGuire & C.B.Winsten 1956. Studies in the economics of transportation. *Technical report, Cowles commission monograph, New Haven, Conn.: Yale university press*.

Boore, D.M. & Atkinson, G.M. 2008. Ground-motion prediction equations for the average horizontal component of PGA, PGV and 5% damped SA at spectral periods between 0.01s and 10.0s. *Earthquake Spectra* 24(1): 99–138.

Chang, S. 2003. Evaluating disaster mitigations: Methodology for urban infrastructure systems. *Natural hazards Rev. 4, 186*.

Cho, S., Gordon, P., Moore, J.E., Richardson, H.W., Shinozuka, M. & S. Chang. Integrating transportation network and regional economic models to estimate the costs of a large urban earthquake. *Journal of regional science* 41(1): 39–65.

Crowley, H. & Bommer, J.J. 2006. Modelling seismic hazard in earthquake loss models with spatially distributed exposure. *Bulletin of Earthquake Engineering* 4(3): 249–273.

Dueñas-Osorio, L., Craig, J. I., Goodno, B. J., & Bostrom, A. 2005. Interdependent response of networked systems. *Journal of Infrastructure Systems* 13(3): 185–194.

Frank, M. & Wolfe, P. 1956. An algorithm for quadratic programming. *Naval research logistics quarterly* 3: 95–110.

HAZUS 1999. Earthquake loss estimation technical manual. *Technical report, National institute of building sciences*, *Washington D.C.*

Jayaram, N. & Baker, J. W. 2008a. Statistical tests of the joint distribution of spectral acceleration values. *Bulletin of the Seismological Society of America* 98(5): 2231–2243.

Jayaram, N. & Baker, J. W. 2008b. Correlation model for spatially-distributed ground-motion intensities. *Earthquake engineering and structural dynamics (in review)*.

Kang, W. H., Song, J., & Gardoni, P. 2008. Matrix-based system reliability method and applications to bridge networks. *Reliability engineering and system safety* 93: 1584-1593.

Kim, Y.S., Spencer, B.F., Song, J., Elnashai, A.S., & Strokes, T. 2007. Seismic performance assessment of interdependent lifeline systems. *Mid-American earthquake center, University of Illinois at Urbana-Champaign*.

Kiremidjian, A.S., Moore, J., Fan, Y.Y., Basiz, N., Yazali, O., & Williams, M. 2003. PEER highway demonstration project. *Earthquake engineering*.

Kiremidjian, A.S., Stergiou, E., & Lee, R. 2007. Issues in seismic risk assessment of transportation networks. In Pitilakis, K.D. (eds), *Chapter 19, Earthquake Geotechnical Engineering*, Springer.

Law, A.M. 2007. *Simulation modeling and analysis*. McGraw-Hill.

McGuire, R.K. 2007. Seismic hazard and risk analysis. *Earthquake engineering research institute*.

Melchers, R.E. 1999. *Structural reliability analysis and prediction, 2nd edition*. JohnWiley & Sons.

Park, J., Bazzurro, P., & Baker, J.W. 2007. Modeling spatial correlation of ground motion intensity measures for regional seismic hazard and portfolio loss estimation. *10th International Conference on Application of Statistics and Probability in Civil Engineering (ICASP10), Tokyo, Japan*.

Shiraki, N., Shinozuka, M., Moore II, J. E., Chang, S. E., Kameda, H., & Tanaka, S. 2007. System risk curves: Probabilistic performance scenarios for highway networks subject to earthquake damage. *Journal of Infrastructure Systems* 213(1): 43–54.

Stergiou, E. & Kiremidjian, A.S. 2006. Treatment of uncertainties in seismic risk analysis of transportation systems. *Technical report, Blume earthquake engineering center*, Stanford University.

USGS 2003. Earthquake probabilities in the San Francisco bay region: 2002-2031. *Technical report, Open file report 03-214*, USGS.

Youngs, R.R. & Coppersmith, K.J. 1985. Implications of fault slip rates and earthquake recurrence models to probabilistic seismic hazard estimates. *Bulletin of the seismological society of America* 75(4): 939–964.