A Framework For Selecting A Suite Of Ground-Motion Intensity Maps Consistent With Both Ground-Motion Intensity And Network Performance Hazards For Infrastructure Networks

M. Miller & J. Baker
Stanford University, Stanford, CA, USA

ABSTRACT: To assess the risk of natural disasters to complex urban infrastructure networks, researchers commonly use variations on Monte Carlo simulation because the the link between natural disasters and the network performance and community impact is difficult to express analytically. Particularly if evaluating the network performance is computationally intensive, a challenge then is to choose an appropriate suite of scenarios to capture the risk posed by natural disasters that allows for a computationally tractable evaluation. In seismic risk assessment, for example, efforts to choose a representative suite of scenarios, as expressed by weighted ground motion intensity maps, have traditionally focused on consistency with the seismic hazard only, whereas in some applications consistency with the exceedance curves of a performance measure may be more relevant. Here we show a method for selecting and evaluating a reduced suite of ground motion intensity maps consistent with both ground motion intensity hazard and performance measure exceedance curves for a regional network of bridges, highways and minor roads. We first use optimization to select a reduced suite of ground motion intensity maps. We then evaluate the consistency with the ground motion hazard. In the second stage, we choose a computationally efficient performance measure that is representative of a metric of greater interest and evaluate how well the reduced suite of ground motion intensity maps matches the performance measure exceedance curves obtained using an extensively-sampled catalog of ground motion intensity maps. To illustrate our approach, we present an example application of the San Francisco Bay Area road network for which we evaluate the hazard consistency region-wide at all highway bridge locations and the consistency with selected performance measures of social or technical resiliency. Our results indicate that with a reduced suite of re-weighted ground motion intensity maps, we can reasonably estimate the exceedance rates of potential ground motion intensity and performance measures, including percentage change in average morning travel time 2-3 post-earthquake. While we have focused on the seismic risk to urban road networks, this framework is generalizable to assessing network risk from a wide range of perils.

1 INTRODUCTION

The post-disaster functionality of infrastructure networks is crucial for the overall community resiliency, such as indicated by access to hospitals or economic impacts to supply chains. Because the link between natural disasters and community resiliency is difficult to express analytically, researchers commonly use variations on Monte Carlo simulation (e.g., (Shiraki et al. 2007, Rokneddin et al. 2012)). One strategy popular in current practice is to evaluate the damage to a portfolio of components of an infrastructure system given one single scenario earthquake and assuming deterministic exposure and vulnerability. However, a key limitation with this approach is that the range of possible events and damages is not captured. An alternative framework in the research literature involves Monte Carlo simulation to generate ground motion intensity maps as described by Crowley and Bommer (2006) among others, and then the following 4-step framework to compute the network performance measures as described by Jayaram and Baker (2010), Han and Davidson (2012) and others, as illustrated in Figure 1. The network performance measures capture aspects of the social or technical resiliency of the communities and infrastructure networks (Bruneau et al. 2003).

Step 1: Generate ground motion scenarios from appropriate seismic source models using Monte Carlo simulation and its variants.

Step 2: For each ground motion scenario, generate ground motion intensity maps from ground motion scenarios, ground motion prediction equations, and models of the inter-event and intra-event residuals using Monte Carlo Simulation and its variants.

Step 3: For each ground motion intensity map,

generate damage maps by sampling a damage state for each structure of interest given the ground motion map and appropriate fragility functions.

Step 4: Compute the network performance measures based on the damage map and the performance measure algorithm.

The number of Monte Carlo simulations depends on a number of factors including computational resources, computational efficiency of the network performance measure algorithm, and sampling procedure. Since the size of a catalogue of ground motion intensity maps could be arbitrarily large given the various sources of uncertainty and since some performance measure evaluations are particularly computationally intensive, researchers including Jayaram and Baker (2010) and Han and Davidson (2012) have proposed approaches to limit the number of ground motion intensity maps generated at the end of Step 2. In both cases, the evaluation of the resulting suite of ground motion intensity maps has been limited to the ground motion intensity (Step 2) exceedance curves and has not considered possible consistency with the performance measure exceedance curves (Step 4). Note that in this paper, exceedance curves refer to the annual exceedance rates of different values of intensity or performance measures.

Here we show a framework for selecting and evaluating a reduced suite of ground motion intensity maps consistent with both ground motion intensity hazard and performance measure exceedance curves for a full network of bridges, highways and minor roads. The resulting reduced suite of maps can then be used as an input for yet more computationally expensive performance measure calculations, such as the exceedance rates of different levels of post-earthquake accessibility of various communities. To illustrate our approach, we present an example application of the San Francisco Bay Area road network, which contains approximately 12,000 edges representing the local and highway road segments, 33,000 nodes representing the intersections, and 1557 highway bridges. We evaluate the hazard consistency at all bridge locations as well as the consistency with selected performance measures of social and technical resiliency. Our results indicate that with a reduced suite of reweighted ground motion intensity maps, the resulting exceedance curves are reasonably consistent with those obtained from the benchmark method, a conventional Monte Carlo Simulation of both the ground motion intensity and selected network performance measures. While we have used the seismic reliability of the San Francisco Bay Area urban road network as a case study, our framework can be generalized for assessing risk from a wide range of perils that threaten infrastructure networks.

2 METHODS

2.1 *Generating a large suite of candidate ground motion intensity maps*

A block diagram of our overall methodology for efficient seismic risk analysis of infrastructure networks is provided in Figure 3. The methodology starts with building a Monte Carlo suite of spatially-correlated ground motion intensity values at each site location over a large range of earthquake rupture events. For purposes of illustration, we will describe building this suite of maps using the spectral acceleration at a period of $T = 1s$, which is denoted by $Sa(T = 1s)$ or simply Sa.

First, we generate ground motion scenarios from a seismic source model, which estimates the rates that earthquakes of specified magnitudes, locations, and faulting types will occur.

Second, for each ground motion scenario in the seismic source model, we use ground motion prediction equations to predict the mean, $\overline{\ln Sa(M_k, R_{ik}, V_{s30,i}, T, \ldots)}$ and spectral acceleration intra- and inter-event residual standard deviations, which are denoted by σ_{ik} and τ_k respectively, for location $i = 1 \ldots n$ in ground motion scenario $k = 1 \ldots o$, where M_k is the moment magnitude of scenario k, R_{ik} is the closest horizontal distance from the surface projection of the fault plane (R_{JB}) to location i, $V_{s30,i}$ is the average shear wave velocity down to $30m$ at location i, and $T = 1s$ is the oscillator period of interest.

Then, for each of the o ground motion scenarios, we simulate q realizations of the spatially-correlated ground-motion residual terms. The result is a suite of $m = o \times q$ ground motion intensity maps for all scenarios. Readers are referred to Han and Davidson (2012) for a survey of sampling methods.

Equation 1 summarizes this process; the total log spectral acceleration at a site $i = 1 \ldots n$ for a ground motion map $j = 1 \ldots m$ is the linear combination of the mean log spectral acceleration and the residual terms, precisely

$$\ln Sa_{ij} = \overline{\ln Sa(M_j, R_{ij}, V_{s30,i}, T, \ldots)} + \sigma_{ij}\epsilon_{ij} + \tau_j\eta_j. \tag{1}$$

where ϵ_{ij} is the intra-event residual in $\ln Sa$ representing location-to-location variability and η_j is the inter-event residual in $\ln Sa$ and the other parameters are defined above.

Note that the resulting ground motion intensity maps show correlation from two sources, from correlation in the mean log spectral acceleration and correlation in the residual terms as discussed in Baker and Miller (2011).

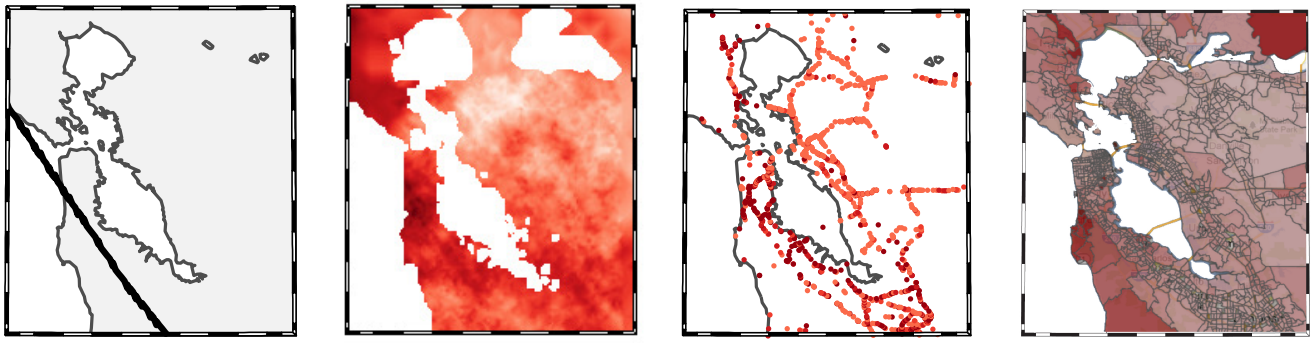

(a) Step 1: Earthquake scenario

(b) Step 2: Ground motion intensity

(c) Step 3: Bridge damage

(d) Step 4: Performance measure

Figure 1: 4-step framework for risk analysis of infrastructure networks. As an example, this figure depicts one earthquake rupture scenario with one realization each of ground motion intensities, bridge damage states and loss in community accessibility based on a variable-demand activity-based traffic model. Darker colors indicate more severe values.

2.2 *Reducing the number of ground motion intensity maps through optimization*

Although the previous section suggests a method to build an arbitrarily large suite of ground motion intensity maps, most applications require a finite catalog of maps.

2.2.1 *Definition and notation*

Each ground motion map contains ground motion measurements at a set of n different sites of interest. The set of possible ground motion intensity maps contains m realizations, $\ln \mathbf{Sa}_1, \ldots \ln \mathbf{Sa}_m \in \mathbb{R}^n$, each with a corresponding annual rate of occurrence, w_j.

In this section, we explain the algorithm of Han and Davidson (2012), which uses mixed integer optimization to choose a re-weighted subset of ground motion intensity maps. That is, the weighting returned by this algorithm minimizes the errors at each site between the 'true' hazard curve and one estimated from a subset of re-weighted ground motion intensity maps on a training sample for a given set of parameters. The ground motion intensity hazard curve error is indicated by e_{ir}^+ and e_{ir}^- in Figure 2 where the parameters $i = 1 \ldots n$ denote the sites at which the objective function will be minimized and the parameters $r = 1 \ldots R$ denote the corresponding return periods. The authors propose weighting the errors by the return period (in years) to more highly penalize errors at longer return periods, which typically correspond to ground motion intensity maps causing higher damage and likely of higher engineering interest. In Section 3.4, we will discuss how to choose the optimization parameters so as to yield a subset of ground motion intensity maps that also yields a low error for the performance measures.

2.2.2 *Evaluation of results*

We evaluate the error in the estimation of exceedance rates between the reduced set of ground motion intensity maps and an extensively-sampled suite to choose parameters for the optimization. Han and Davidson (2012) defined a mean hazard curve error for the ground motion intensity as follows:

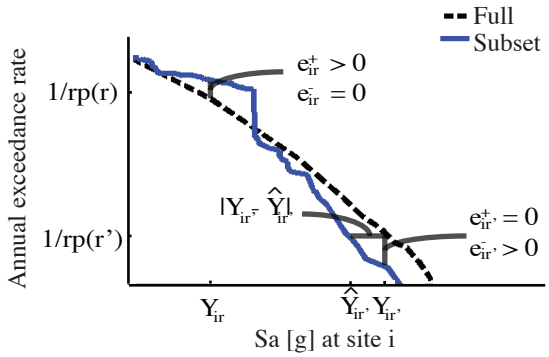

Figure 2: Relationship between error variables used in the optimization formulation (e_{ir}^+, e_{ir}^-) and error variables used in evaluating the results ($Y_{ir}, \hat{Y}_{ir}$) with the empirical ground motion intensity hazard curves at a sample site i for return periods (rp) at indices r and r'.

$$MHCE = \frac{1}{IR} \sum_{r=1}^{R} \sum_{i=1}^{n} \left| \frac{Y_{ir} - \hat{Y}_{ir}}{Y_{ir}} \right| \tag{2}$$

where $\hat{Y}_{ir}$ is the reduced set estimated ground motion intensity associated with return period $r = 1 \ldots R$ for site $i = 1 \ldots n$ and the ground motion estimated from an extensively-sampled suite is Y_{ir}. Figure 2 illustrates the relationship of $Y_{i,r} - \hat{Y}_{i,r}$ to the objective function used for optimization and the ground motion intensity curve.

We extend this error metric to evaluate the error in performance measure exceedance curves between candidate subsets of ground motion intensity maps and an extensively-sampled suite of ground motion intensity maps. This error metric is defined as

$$MPMCE = \frac{1}{R} \sum_{r=1}^{R} \left| \frac{Z_{pr} - \hat{Z}_{pr}}{Z_{pr}} \right| \tag{3}$$

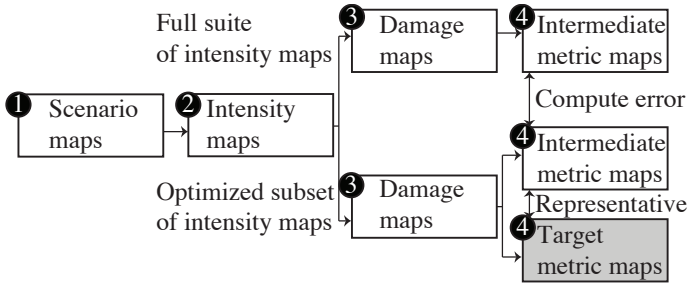

Figure 3: Proposed approach to seismic risk analysis of infrastructure networks using a reduced suite of ground motion intensity maps, the optimized subset.

where $\hat{Z}_{pr}$ is the reduced set estimated performance measure value associated with return period r for performance measure p and Z_{pr} is estimated from an extensively-sampled suite.

We evaluate our choice of maps in the subset by computing the error between the intermediate metric maps of the extensively-sampled suite and those of the subset, as shown in Figure 3 under Step 4.

3 APPLICATION

The case study focuses on the seismic risk of the road network of the 9-county San Francisco Bay Area, which covers an area of over 1000 km^2 (400 $miles^2$) and is home to over 7 million people.

3.1 *EARTHQUAKE SCENARIOS AND GROUND MOTION MAPS*

This work uses the Uniform California Earthquake Rupture Forecast, Version 2 (UCERF2) (Field et al. 2009), which computes the rate of occurrence of each ground motion scenario and is implemented under the software framework of OpenSHA (Field et al. 2003). The forecast is discretized into a list of identified faults and a stratified list of magnitudes and rupture locations for each. Excluding background seismicity, the forecast contains 1701 ground motion scenarios relevant for the case study region that are importance-sampled based on magnitude and each have an annualized rate of occurrence greater than or equal to 10e-5. We build the 'base case' input suite by combining the mean and spatially-correlated residual terms of the ground motion intensity according to Equation 1 using the models by Boore and Atkinson (2008) and Jayaram and Baker (2009) respectively. We will compare the results using a subset of the 'base case' ground motion intensity maps found using optimization with those using an extensively-sampled suite, which corresponds to 13,608 ground motion intensity maps from 8 residual realizations per ground motion scenario with importance-sampled magnitudes.

3.2 *NETWORK*

3.2.1 *Bridges*

We consider a portfolio of 1557 highway bridges in the San Francisco Bay Area, which include all unique state-managed bridges in current operation on our road network and were hand-matched to the road network. In this study, if a bridge collapses, the road it carries, as well as that which it crosses, are modeled as closed.

The serviceability of the bridges is modeled as full traffic capacity for bridges in what are defined as *none*, *slight*, or *moderate* damage states and no traffic capacity for *extensive* and *complete* damage states, as detailed in Werner et al. (2006) for the chosen time scale of 2-3 days after an earthquake event. Thus, the problem can be reduced to considering these two groups of damage states, which will be designated as *minor* and *major*.

Fragility functions relate the ground motion intensities to the post-earthquake bridge condition, which are expressed in terms of discrete damage states, and have the following form:

$$P(DS \geq ds_s | Sa = x, b) = \Phi\left(\frac{\ln x - \lambda_{s,b}}{\xi}\right) \qquad (4)$$

where DS is a discrete random variable, whose values represent the damage state of the bridge. The events s in the sample space S for the point in time two days after an event are ds_{minor} and ds_{major}. Caltrans provided the fragility function values used in this study using the method detailed in Basoz and Mander (1999). For each of 28 primary bridge classes $b = 1 \ldots 28$, defined based on structural characteristics, $\lambda_{s,b}$ and ξ are respectively the mean and standard deviation of the $\ln Sa$ value necessary to cause the s^{th} damage state to occur or be exceeded. Φ is the standard normal cumulative distribution function. The structural capacities of individual bridges are modeled as uncorrelated.

For each ground motion intensity map and bridge, we sample the damage state from the distributions of $P(DS \geq ds_s | Sa = x, b)$. Figure 1(c) illustrates an example of the *component damage map*. Here one realization of the bridge damage state is computed per ground motion map.

3.2.2 *Highway network*

The highway network for this case study is represented by a directed graph $G = (V, E)$, where V is a finite set of vertices representing intersections and the set E, whose elements are edges representing road links, is a binary relation on V. This study considers the latest (2010) version of the San Francisco Bay Area transportation network where $(|V|, |E|) = (11921, 32858)$ including centroidal links and $(|V|, |E|) = (9635, 24404)$ without. The model includes both highways and main local roads as well as the relevant trip demand data.

For each *component damage map*, we create a *network damage map*, which is a damaged network with the traffic flow capacity equal to zero on links with bridges in the *major* damage state (Werner et al. 2006) and provided values for length, traffic flow capacity per direction, lanes and free-flow speed on the other links. We do not consider damage on links without bridges.

3.3 *TARGET PERFORMANCE MEASURE AND INTERMEDIATE PERFORMANCE MEASURE*

The aim of performance measures is to gain insight into the performance of the network and impact on the human and/or natural environment. Metrics proposed in the literature include connectivity (Rokneddin et al. 2012), traffic flow capacity (Lee et al. 2011), fixed-demand traffic delay (Shiraki et al. 2007, Jayaram and Baker 2010, Stergiou and Kiremidjian 2006), economic impacts from increased travel time and bridge repairs (Stergiou and Kiremidjian 2006) and percentage change in average travel time (Rodier et al. 2010).

Here, we consider the percentage change in average morning travel time assuming fixed travel demand, which is highly correlated with the fixed-demand traffic delay. We evaluate the change in travel time by implementing the iterative traffic assignment (ITA) (Chen and Alfa 1991), which has been shown by Wang et al. (2012) to be consistent with results from the User Equilibrium (UE) method and leveraging the network analysis module NetworkX in Python (with some base algorithms in C and FORTRAN). We have assumed fixed traffic demand values, which is a simplification common in current literature.

We then find the most representative computationally-efficient intermediate performance measure by determining which of the intermediate performance measures is most highly correlated with the target performance measure. We perform our 4-step framework with an extensively-sampled suite of ground motion intensity maps, the target performance measure, and the candidate performance measures, namely percent of bridges damaged, connectivity, traffic flow capacity, and weighted-shortest path between locations of interest. We find that the percent of bridges damaged is the intermediate performance measure most highly correlated with the target performance measure, as illustrated in Figure 4. While we have evaluated the target performance matrix for an extensively-sampled suite of ground motion maps to verify our choice of an intermediate welfare metric, we recommend that readers use our proposed intermediate welfare metric, percent of bridges damaged, for similar applications. For different applications, readers can minimally sample ground motion intensity maps, the target performance measure, and the candidate performance measures to estimate the candidate performance measure most highly correlated with the target performance measure.

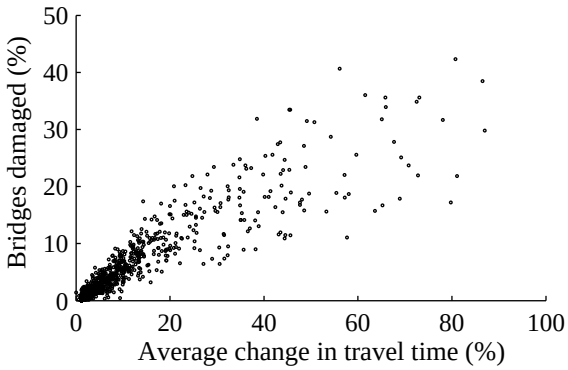

Figure 4: The intermediate performance measure *percent of bridges damaged* versus the target performance measure of greater interest, *percentage change in average morning travel time* for a suite of simulated damaged networks.

3.4 *TUNING THE OPTIMIZATION FORMULATION*

As detailed in Section 2.2.1, there are up to $2m + 2nR$ total variables in the optimization problem, which we have implemented in the python interface to the ZIB Optimization Suite. To achieve convergence, it is advantageous to choose new smaller values m', n', and R'. We do this by performing a sensitivity analysis on these three main parameters and comparing the seismic hazard curve error and performance measure exceedance curve error as described in Section 2.2.2 that we evaluate over all 1557 bridge locations and over 50 total return periods between 50 and 2500 years, which is a range of engineering interest. The number of return periods and sites are both an order of magnitude larger than in previous work.

We must first chose the maximum cardinality of the desired final subset of ground motion intensity maps. This parameter value depends on the computational efficiency of the target performance measure and available computational resources. For this study, $m_{red} = 200$. Next, we reduce the number of sites from $n = 1557$ to n'. The sites are clustered using K-Means clustering in Matlab using the Euclidian distance between points as the distance metric. We find that the reduction of the error with increasing number of sites begins to taper after about 10 or 12 sites. Thus, we select 12 centroidal sites and the resulting 12 clusters of bridges are illustrated in Figure 8(c). Third, we select the number of input ground motion intensity maps, m'. We find that one ground motion map for each of the 1701 discrete ground motion scenarios provides a good balance of running time, required maps, and error reduction. Finally, we test the sensitivity of the results to the number of return periods of interest $r = 1 \dots R'$ within the selected return period range, 50 to 2500 years as illustrated in Figure 5. By looking only at the mean ground motion intensity hazard curve error, one might choose $R' = 20$ return

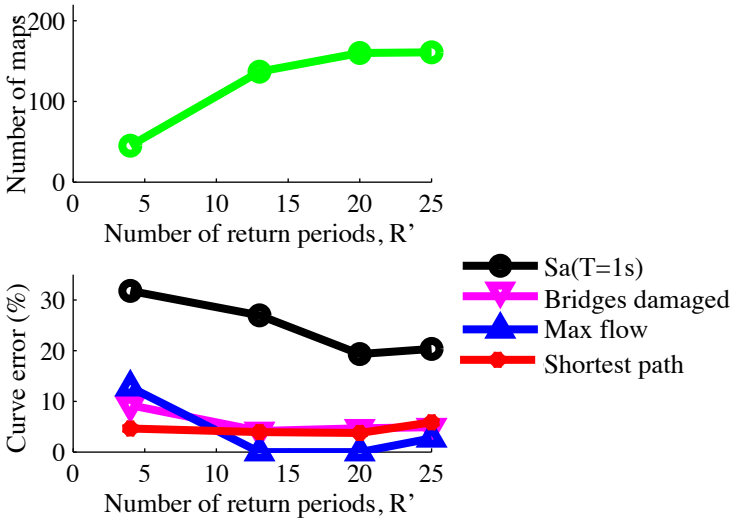

Figure 5: Sensitivity of the number of ground motion intensity maps required (where new annual exceedance rate $\geq 5e-5$) (top) and error (bottom), as measured by modified mean hazard curve error (MHCE) and mean performance measure exceedance curve error (MPMCE), to changes in the number of return periods of interest used in the optimization, R'. m' and n' are 1701 and 12 respectively.

periods. However, by examining the intermediate performance measure identified in Subsection 3.3, we see that $R' = 13$ return periods is also a candidate choice, which actually has a lower error for our intermediate performance measure while still corresponding to a reasonable hazard curve error. In this application, $R' = 13$ is preferable because it corresponds to a 15% reduction in the total number of maps of interest required in the final catalogue (137 vs. 161 final maps given the other parameter values). For the following two sections, we will use $(m', n', R') = (12, 13, 1701)$ to illustrate our 2-stage evaluation procedure.

3.5 *EVALUATING CONSISTENCY WITH GROUND MOTION INTENSITY EXCEEDANCE RATES*

To evaluate how representative a suite of ground motion maps is of the ground motion intensity hazard, one must consider a number of factors including source distribution, magnitude exceedance rates, correlation structure and marginal and joint distributions of the ground motion intensity.

A potential pitfall when selecting a reduced suite of ground motion maps is grossly over-representing a main fault or two, such as the Northern San Andreas fault in the San Francisco Bay Area. As might be expected, the procedure preferentially samples from the Northern San Andreas fault. However, the overall distribution is relatively balanced with a strong contribution from Hayward Fault events and in fact 41 unique fault segments are included in the final ground motion suite. We also spot-check seemingly similar events in the final subset suite, like some Hayward events, and have found that they show distinguishably different

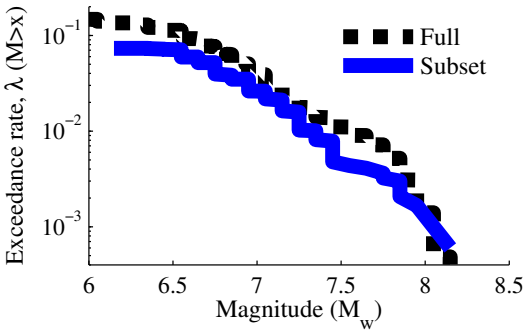

Figure 6: Annual exceedance curves of earthquake magnitudes for the extensively-sampled ('Full') suite of ground motion intensity maps and the subset ('Subset'). The weight for each scenario is the annual recurrence rate.

patterns of regional ground motions.

We also analyze the exceedance curves of the moment magnitudes (M_w). Figure 6 compares the exceedance curves of the earthquake magnitude between the extensively-sampled suite of ground motion intensity maps and the selected subset. We see a reasonably close match. At lower magnitudes, as expected, the subset somewhat underestimates the rate because these magnitudes aren't contributing as much as other magnitudes to the Sa range of engineering interest. One contributing factor to the apparent underestimate is that the residual values of the ground motion intensity tend to have a mean greater than zero, corresponding to a higher distribution of positive residual values, and we have optimized based on the ground motion intensity.

Two key evaluations of ground motion intensity consistency are comparing the marginal and the joint distributions of the ground motion intensity. Figure 7 compares the marginal distributions of ground motion intensity between the extensively-sampled suite of ground motion intensity maps and the subset. For both the centroidal sites on which we minimized the MHCE in the optimization (top row) and the non-centroidal sites not part of the optimization (bottom row), we see that the selected subset reasonably reproduces the expected ground motion intensity hazard curves. While the MHCE hovers around 20-30% when evaluated over all sites and a large range of return periods, Figure 7 illustrates that the hazard curves generally show a reasonably tight correspondence. Figure 8(b) illustrates the joint distribution of Sa between two particular locations that represent an expected worst case because their ground motion intensity hazard is each dominated by different sources (Hayward and San Andreas faults respectively). Even in this case we see that the Sa are well distributed, neither showing extremely high correlation or very low correlation and capturing much of the variability in the underlying extensively-sampled suite of ground motion maps.

(a) Ground motion hazard curves at 5 random centroidal locations

(b) Ground motion hazard curve at 5 random non-centroidal locations

Figure 7: Comparison of the marginal distributions of ground motion intensity at both centroidal (used in the optimization) and non-centroidal locations between the extensively-sampled ('Full') suite of ground motion intensity maps and the subset ('Subset').

3.6 *EVALUATING CONSISTENCY WITH PERFORMANCE MEASURE EXCEEDANCE RATES*

After considering the consistency with the ground motion intensity, we evaluate consistency with the performance measure exceedance rates.

Figure 9 illustrates how the performance measure exceedance curve is reasonably consistent between the subset and extensively-sampled suite of ground motion intensity maps for the intermediate performance measure, percentage of bridges damaged.

Using the chosen subset of ground motion intensity maps and generating corresponding component and network damage maps, we estimate the exceedance curve of the target performance measure as shown in Figure 9. By comparing the exceedance curve estimated by the subset of maps with the curve from a extensively-sampled suite, we see a tight correspondence. While we have provided the results of the extensively-sampled suite here to test our method, future work can benefit from saving multiple orders of magnitude of computational expense by tuning the final suite of maps using the intermediate performance measure and evaluating the target performance measure using only the reduced tuned suite of ground motion intensity maps.

We also find that our 2-stage evaluation procedure yields a lower MPMCE for the target performance measure than simply choosing a set of maps and parameters by minimizing the MHCE. For example, if instead of $(m', n', R') = (12, 13, 1701)$ as we have chosen here, we choose $(m', n', R') = (8, 20, 1701)$ because it has a lower MHCE (23% instead of 27%), we would yield a subset of maps with a higher MPMCE for the intermediate metric (9.1% instead of 4.1%) and significantly, higher MPMCE for the target metric (4.6% instead of 3.7%). This indicates the value of considering the intermediate metric in selecting suites of ground motion maps.

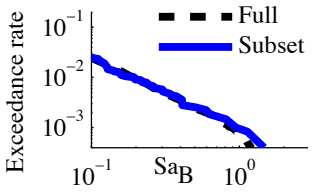

(a) Ground motion hazard curve at site B

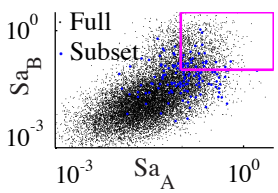

(b) Ground motion hazard at sites A and B

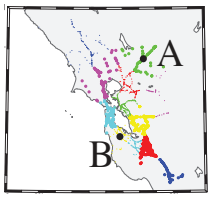

(c) Map of sites A and B; other sites colored by cluster

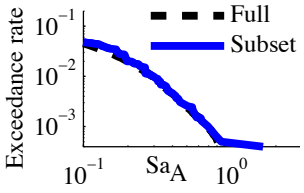

(d) Ground motion hazard curve at site A

Figure 8: Comparison of the marginal and bivariate distributions of ground motion intensity at two locations between the extensively-sampled ('Full') suite of ground motion intensity maps and the subset ('Subset'). The contributions to the hazard at the two sites are expected to be rather different. The box in subfigure 8(b) marks the axis limits of the subfigures 8(a) and 8(c).

4 CONCLUSIONS

We have presented a method for selecting a suite of ground motion intensity maps and evaluating how consistent this suite of maps is with the exceedance curves of both the ground motion intensity and network performance. For illustration, we have applied our method to a case study of the road network of the San Francisco Bay Area. We have shown common pitfalls in selecting a suite of maps that can be

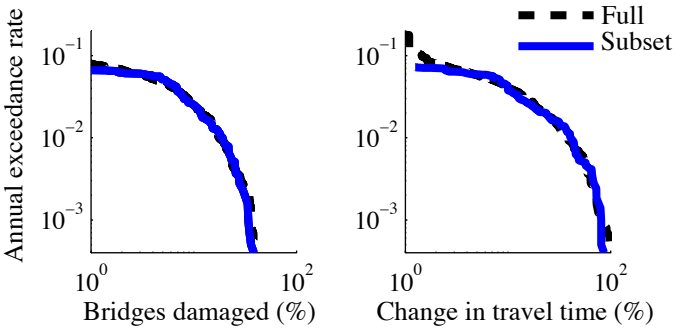

Figure 9: Comparison of the performance measure exceedance curves between the extensively-sampled ('Full') suite of ground motion intensity maps and the subset ('Subset'). The left figure is the intermediate performance measure, *percentage of bridges damaged*, and the right figure is the target performance measure, *percentage change in average morning travel time*.

avoided by checking for consistency with magnitude and source distributions as well as the marginal and joint distributions of ground motion intensity. Additionally, we have proposed tuning the optimization results with an efficient performance measure that is representative of the target performance measure. In the case study, the method produces a suite of 137 ground motion intensity maps.

Our framework is agnostic to the choice of intensity measures and performance measures and thus is broadly applicable to a variety of perils and types of spatially-distributed portfolios. For the 4-step simulation framework, future users need a stochastic catalog of possible events (Step 1), a model for the intensity measure at each location of interest given each event (Step 2), a vulnerability model (Step 3), and a model to evaluate the network performance measure given a damaged network (Step 4). Thus, we expect this simulation framework and evaluation procedure to be be an appropriate method for selecting event scenarios even as computational efficiency increases and new intensity measures or simulation techniques become the norm. For example, as computational limitations are overcome, one direction for future work includes explicitly considering computationally-intensive target performance measures in the optimization formulation instead of intermediate performance measures as we have done here. Work is in progress to test the results of this framework with performance measures such as changes in accessibility as modeled by an agent-based model. This paper further suggests future investigations into interpreting the results of the analysis to inform decision-making such as prioritizing bridge retrofits, quantifying supply chain risk, or identifying and increasing the resiliency of communities particularly at risk to losses in accessibility.

5 ACKNOWLEDGEMENTS

We thank Dave Ory at the Metropolitan Transportation Commission (MTC) and Tom Shantz and Loren Turner at Caltrans for motivating discussions and providing all the case study data. We also thank Jessica Jacobo for painstakingly integrating the Caltrans and MTC data sources. The first author gratefully acknowledges the support of the Stanford Graduate Fellowship and the National Science Foundation Graduation Research Fellowship. This work was supported in part by the National Science Foundation under NSF grant number CMMI 0952402. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

Baker, J. W. & M. Miller (2011). Effects of earthquake source geometry and site conditions on spatial correlation of earthquake ground motion hazard. Santa Barbara, CA. 4th IASPEI/IAEE International Symposium on Effects of Surface Geology on Seismic Motion.

Basoz, N. & J. Mander (1999). Enhancement of the highway transportation lifeline module in HAZUS. Final Pre-Publication Draft (#7) prepared for the National Institute of Building Sciences (NIBS).

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

Bruneau, M., S. E. Chang, R. T. Eguchi, G. C. Lee, T. D. ORourke, A. M. Reinhorn, M. Shinozuka, K. Tierney, W. A. Wallace, & D. von Winterfeldt (2003). A framework to quantitatively assess and enhance the seismic resilience of communities. *Earthquake Spectra 19*(4), 733–752.

Chen, M. & A. S. Alfa (1991). A network design algorithm using a stochastic incremental traffic assignment approach. *Transportation Science 25*(3), 215–224.

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

Field, E. H., T. E. Dawson, K. R. Felzer, A. D. Frankel, V. Gupta, T. H. Jordan, T. Parsons, M. D. Petersen, R. S. Stein, R. J. Weldon, & C. J. Wills (2009). Uniform california earthquake rupture forecast, version 2 (UCERF 2). *Bulletin of the Seismological Society of America 99*(4), 2053–2107.

Field, E. H., T. H. Jordan, & C. A. Cornell (2003). OpenSHA: a developing community-modeling environment for seismic hazard analysis. *Seismological Research Letters 74*(4), 406–419.

Han, Y. & R. A. Davidson (2012). Probabilistic seismic hazard analysis for spatially distributed infrastructure. *Earthquake Engineering & Structural Dynamics 41*(15), 2141–2158.

Jayaram, N. & J. W. Baker (2009). Correlation model for spatially distributed ground-motion intensities. *Earthquake Engineering & Structural Dynamics 38*(15), 1687–1708.

Jayaram, N. & J. W. Baker (2010). Efficient sampling and data reduction techniques for probabilistic seismic lifeline risk assessment. *Earthquake Engineering & Structural Dynamics 39*(10), 1109–1131.

Lee, Y.-J., J. Song, P. Gardoni, & H.-W. Lim (2011). Post-hazard flow capacity of bridge transportation network considering structural deterioration of bridges. *Structure and Infrastructure Engineering 7*(7-8), 509–521.

Rodier, C., J. E. Abraham, B. Dix, & J. D. Hunt (2010). Equity analysis of land use and transportation plans using an integrated spatial model. Washington D.C. Preprint for the 89th Annual Meeting of the Transportation Research Board.

Rokneddin, K., J. Ghosh, L. Dueas-Osorio, & J. E. Padgett (2012). Bridge retrofit prioritisation for ageing transportation networks subject to seismic hazards. *Structure and Infrastructure Engineering: Maintenance, Management, Life-Cycle Design and Performance*, 1–17.

Shiraki, N., M. Shinozuka, J. E. Moore, S. E. Chang, H. Kameda, & S. Tanaka (2007). System risk curves: Probabilistic performance scenarios for highway networks subject to earthquake damage. *Journal of Infrastructure Systems 13*(1), 43–54.

Stergiou, E. & A. S. Kiremidjian (2006). *Treatment of uncertainties in seismic risk analysis of transportation systems*. Ph. D. thesis, Stanford University.

Wang, P., T. Hunter, A. M. Bayen, K. Schechtner, & M. C. Gonzlez (2012). Understanding road usage patterns in urban areas. *Scientific Reports 2*, 1–6.

Werner, S., C. Taylor, S. Cho, & J. Lavoie (2006). REDARS 2 methodology and software for seismic risk analysis of highway systems (technical manual). Oakland, CA. Seismic Systems & Engineering Analysis for MCEER.