

Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds

MULTI-OBJECTIVE OPTIMIZATION FOR BRIDGE RETROFIT TO ADDRESS EARTHQUAKE HAZARDS

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

Sandia National Laboratories
 PO Box 5800
 Albuquerque, NM 87185-1188, USA

Linda K. Nozick
 Natalia Romero
 Ningxiong Xu

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

ABSTRACT

Protecting infrastructures against natural hazards is a pressing national and international problem. Given the current budgetary climate, the ability to determine the best mitigation strategies with highly constrained budgets is essential. This paper describes a set of computationally efficient techniques to determine optimal infrastructure investment strategies, given multiple user objectives, that are consistent with an underlying earthquake hazard. These techniques include: optimization methods for developing representative events to characterize the hazard and the post-event condition of infrastructure components, a simulation model to characterize post-event infrastructure performance relative to multiple user objectives, and a multi-objective optimization algorithm for determining protection strategies. They are demonstrated using a case study of the highway network in Memphis, Tennessee.

1 INTRODUCTION

Ensuring the continued operations of infrastructures after natural disasters is a pressing national and international problem. For example, Kunreuther et al. (2009) reports that economic losses from natural disasters have increased 15-fold since the 1950s. While it is clear that there is a need to protect infrastructures against the effects of natural disasters, identifying effective strategies when resources are scarce is challenging.

Figure 1 outlines an analysis process to identify optimal mitigation investment strategies. The first step is to characterize the underlying hazard using appropriate physical models. This step results in a set of earthquake scenarios that match the underlying hazard. The next step is to characterize the damage to each infrastructure component under each scenario which reflects component-level impacts based on a selected mitigation. Once the vulnerabilities associated with the infrastructure components are understood, a system model is used to assess the performance characteristics of the network. Finally, an optimization engine is used to determine the best investment strategy for a given set of objectives. It is important to notice that as the selection of package of protection strategies changes, the system-level behavior of the infrastructure changes requiring system-level impacts to be reassessed. Based on the assessment of these impacts, the mitigation package is modified. This process is repeated until an optimal solution is found.

The main challenge with the approach described above is that, for most real-world applications, it presents significant computational challenges. First, the post-event condition of the infrastructure is inherently uncertain, even under a single known earthquake event. Second, there are always multiple, potentially competing, user objectives. Third, a computationally challenging system-level model is required to

assess post-event performance. Finally, determining the optimal protection strategy requires many assessments of component vulnerabilities and the resulting system-level performance for each package of mitigation strategies. Without efficient computational strategies for each step of this process, determining the optimal protection is computationally impractical.




Figure 1: Analysis Structure of Mitigation Investment

Several authors have focused on pieces of this problem. Chang, et al. (2000) randomly generates 10 consequence scenarios for each single event (where about 30 events are used to describe the hazard) using the probabilistic information derived for each bridge to analyze the impacts of earthquakes on the Los Angeles highway system. Shiraki et al. (2007) extend the analysis in Chang et al. (2000) to consider 47 earthquake events using 10 Monte Carlo samples of damage realizations of each event. Our analysis can be viewed as a complementary analysis in which we apply optimization to generate the realizations of each event in lieu of Monte Carlo simulation. Jayaram and Baker (2010) apply Monte Carlo sampling to the San Francisco highway system. They show that about 150 ground motion intensity maps are needed to represent the hazard. For each ground motion map they use simulation to assess system performance using fragility curves. A novel element of their analysis is the use of data reduction techniques to reduce the number of ground motion maps. As in Shiraki et al. (2007), this analysis can also be viewed as complementary to that given in Jayaram and Baker (2010) in that we propose a method to generate damage maps using optimization instead of Monte Carlo simulation of ground motion maps. Chang et al. (2012) apply Monte Carlo sampling to simulate a single earthquake's intensity, resultant damage to bridges and finally resultant bridge residual capacity. Based on this, they solve 1,000 (one for each scenario for which residual capacity is identified for each bridge) mixed integer programs where each of 616 bridges can be selected for retrofit subject to a budget constraint. They then rank the bridge retrofit options.

Other authors focused more heavily on the optimization of retrofit decisions and take the hazard modeling/component modeling more as given. For example, Liu et al. (2009) employ a generalized Benders decomposition solution strategy to identify the bridges to retrofit to minimize post-event travel time under a static system-optimal assumption for driver behavior with a budget constraint. They illustrate their approach on a case study with up to 13 bridges that are candidates for retrofit and as many as 6 earthquake scenarios. Peeta et al. (2010) focuses on optimizing network connectivity post-event. They develop a novel solution approach that is shown to be effective on a representation of Istanbul's highway network, which has about 30 links. Furuta et al. (2011) focus on optimizing inspect and repair strategies for bridges in a highway network when there are multiple objectives of concern including life cycle costs with seismic risks. They develop a multi-objective genetic algorithm and illustrate it on a case study with 10 bridges.

Our focus is similar to those described above, however our intent is to treat the spectrum of issues that arise. We focus on problems with hundreds of bridges that are candidates for retrofit using dynamic estimates of travel time impacts, on problems that cover a large geographic area, under a stochastic representation of the hazard when there are multiple objectives of interest. Given this, we focus on the use of heuristic methods to construct the solution approach. We illustrate this approach through a case study of the highway network in Memphis, Tennessee. While this example focuses on determining the optimal set of highway bridges to reinforce to protect against the seismic hazards posed to a highway network, these methods could be adapted to address other natural hazards and other infrastructure systems.

2 MODELING STEPS

This section describes the computational approaches illustrated in Figure 1, in the context of a case study focused on the highway network in Memphis, Tennessee, which is subject to substantial seismic hazards. 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.

For this study we focus on the Memphis road network where the key vulnerabilities are in the 335 highway bridges built prior to 2000, included in the National Bridge Inventory produced by the Federal Highway Administration (2011). Of these 335 bridges, 286 can be retrofitted to reduce their seismic vulnerabilities. Therefore, the goal is to select a set of these bridges to reinforce that optimizes several objectives. Figure 2 illustrates the Memphis highway network, the bridges of interest, and the traffic analysis zones.




Figure 2. Memphis Highway Network, Traffic Analysis Zones and Bridges (squares).

Section 2.1 describes the three objectives that were considered in this study. Two of these objectives measure aspects of the post-event performance of the highway network. Measuring these two objectives is challenging since they require characterizing the uncertainty in the seismic hazard and resulting bridge damage (Sections 2.2-2.4), as well as simulating the post-event traffic flow over the network (Section 2.5). Section 2.6 describes the algorithm that was used to determine the optimal investment strategy which minimizes these objectives. Section 2.7 presents results for the Memphis highway network.

2.1 Objectives

For this study, three objectives were considered. The first objective is to minimize the increase in travel time for typical daily travel during the morning commute. The second objective is to control the increase in travel time needed for people to reach their nearest hospital. The third is to minimize the cost of the ret-

rofit investments. It is important to notice that the performance of any package of mitigation investments (where a package identifies what retrofit, if any, will be undertaken for each bridge) under each of the first two objectives, is described by a probability distribution. Because there are a range of events that can occur and, for each event, there is uncertainty in the underlying damage that will be done to each bridge, a probability distribution is needed to describe the performance in each dimension.

2.2 Identifying Earthquake Scenarios and their Hazard consistent Probabilities of Occurrence

The seismic hazard in NMSZ is modeled by identifying a small set of earthquake events and an associated probability of occurrence which are representative of the seismic hazard in the region. While seismic hazards can be characterized for each point within the region, we use a scenario approach since the highway network is spatially distributed and capturing the spatial impacts of the seismic hazard is essential. We identify these scenario events using Vaziri (2012). The core idea in that method is to use optimization to select a small set of earthquake events from a larger collection of candidate events which, in aggregate, represent the seismic vulnerability in the region. For this analysis we characterize earthquake events by peak spectral acceleration (PSA), which is the most relevant term for assessing damage to bridges. We use two sources of candidate earthquakes scenarios. The first source is the 433 earthquakes in the Central-East United States earthquake scenarios catalog from the USGS (2008). In addition, we use 20 synthetic events on 5 synthetic faults created by USGS to represent the hazard in New Madrid. The 20 scenarios correspond to each of the 4 possible magnitudes (7.3, 7.5, 7.7 and 8) for ruptures in the 5 different branches described in Petersen et al. (2008). Figure 3 shows the location and magnitude of the 8 events selected for this analysis. Observe that 6 events from the historical catalog as well as two event from the synthetic faults generated by the USGS.




Figure 3. Map showing the location and magnitude of the 6 earthquake scenarios selected through optimization

2.3 Component Damage Modeling

We use HAZUS, a loss estimation methodology developed by FEMA (2010), to create estimates of damage to specific bridges for each earthquake scenario. HAZUS categorizes damage to highway bridges into five classes: no damage, slight, moderate, extensive and complete. The core mechanism to express the relationship between damage state and PSA is a fragility curve (which is defined for each type of highway bridge). Such a curve is illustrated in Figure 4 and is from the HAZUS Manual (2010). A fragility curve

is 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. Bridge mitigation changes the fragility curve which implies a generally lower probability of having higher levels of damage.




Figure 4: Example Fragility Curves for a Class of Highway Bridge (non-California, built prior to 1990 of conventional design with length >150 m).

2.4 Consequence Scenario Construction

Since the damage state is uncertain under a given earthquake scenario, we use an optimization model to generate a family of consequence scenarios for each earthquake event. This optimization model is an extension of that described in Gearhart et al. (2011) which first introduced the notion of creating a family of consequence scenarios to describe the joint distribution of earthquake damage for spatially distributed infrastructure systems. Gearhart et al. (2013) extended that idea to include correlation in damage. In the interest of space, see Gearhart et al. (2011 and 2013) for a thorough discussion of other research that addresses the creation of consequence scenarios.

We extend that formulation to construct consequence scenarios to include the impacts of mitigation. That optimization model is given by the equations below where D is the number of damage states from the loss methodology (5 in the case of HAZUS), K is the number of bridges, and J is the desired number of consequence scenarios for the event interest. Further, let m be 0 and 1 to indicate that a bridge is not mitigated or is mitigated, respectively. Also let π_{kmd} be the probability that bridge k ($k = 1, \dots, K$) falls into damage state d ($d = 1, \dots, D$) given the bridge mitigation state is m ($m = 0, 1$). For HAZUS this would correspond to the probability of being in the relevant damage state from the fragility curve. We let $\sigma_{kmk'm'}$ be the covariance in the damage between pairs of bridges k and k' given bridge k mitigation state is m and bridge k' mitigation state is m' , where $m, m' = \{0, 1\}$.

The objective is to select bridge damage states for each scenario and determine the associated consequence scenario probabilities so the implied statistical measures match the actual statistical measures as closely as possible. In this case, the model seeks to match both the marginal probability for each damage state for each bridge and the covariance in those damage states. The objective function (1) measures the sum squared errors of the deviations from the marginal probabilities (e_{kmd}) and the covariance terms ($\zeta_{kk'mm'}$). The term β is used to assign priority to each of these objectives. Equation (2) computes the deviation in the damage state probabilities for each bridge based on the consequence scenarios constructed and these probabilities as given by the loss estimation methodology. Equation (3) gives the differences

between the actual and implied bridge damage state covariance terms. Equation (4) ensures that each bridge is in exactly one damage state for each scenario, under each mitigation strategy. Equation (5) sets upper and lower bounds on the scenario probabilities. These bounds can be used to limit the s_j variables to values between 0 and 1, however more restrictive bounds can be used. Equation (5) ensures that the sum of scenario probabilities is equal to 1. Equation (6) requires that for each consequence scenario, each bridge is either in a specific damage state or not in that damage state.

$$\sum_{kd} e_{kmd}^2 + \beta \sum_{kk'mm'} \zeta_{kk'mm'}^2 \quad (1)$$

$$\sum_j s_j b_{jkmd} + e_{kmd} = \pi_{kmd}, \quad m, k, d, \quad (2)$$

$$\sum_j \left(\sum_d db_{jkd} - \sum_{jd} db_{jkd} s_j \right) \left(\sum_d db_{jk'd} - \sum_{jd} db_{jk'd} s_j \right) s_j + \zeta_{kk'} = \sigma_{kk'} \quad \forall k, k' | k \neq k', m, m', \quad (3)$$

$$\sum_d b_{jkmd} = 1, \quad m, k, j, \quad (4)$$

$$s_{min} \leq s_j \leq s_{max}, \quad (5)$$

$$\sum_j s_j = 1, \quad (6)$$

$$b_{jkmd} \in \{0,1\}, \quad m, j, k, d. \quad (7)$$

We use an iterative heuristic algorithm to solve the above formulation. That algorithm iterates between identifying the composition of each consequence scenarios (b_{jkmd}) and optimizing the scenario probabilities. We use local search to optimize the variables b_{jkmd} and the nonlinear solver in MATLAB to identify the scenario probabilities, s_j .

Only two of the six earthquake scenarios illustrated in Figure 3 were used since only they produce enough ground shaking to cause sufficient damage to the highway system. For each earthquake, 20 consequence scenarios were constructed as the basis for the highway network performance assessment. Since each consequence scenario has an associated damage level with and without mitigation, the post-event states of network components can be “looked-up” for a given mitigation strategy without the need to generate a new set of consequence scenarios. The efficiency provided this scenario-based approach is essential for being able to accomplish the next two steps: evaluating the network performance and optimizing the investment strategy.

2.5 Assessing Performance Objectives Under Specific Mitigation Strategies

As previously stated, for this case study, there are three objectives to be minimized: cost to reinforce bridges, total peak period travel time, and travel time to the closest hospital. For the this study the reinforcement costs were assumed to be 10% of the replacement cost of the bridge (which is an element in the National Bridge Database). Assessing the other two objectives is more complicated since they are based on the results of simulating the performance of the highway network under each consequence scenario. A dynamic traffic assignment (DTA) algorithm was used to simulate the post-event performance of the highway network.

The results of this model were then used to assess the last two objectives. We use Dynamic User Equilibrium (DUE) as the underlying behavioral assumption in the DTA algorithm we employ. The underlying behavioral assumption behind DUE is that each individual traveler dynamically chooses a route that minimizes their travel time to the destination. That is, as travel conditions change on the network, drivers adapt so as to minimize their travel time. DUE is a generalization of user equilibrium (UE), which is a static analysis which simply assumes that each traveler selects a fixed path that minimizes their travel time. The generalization of UE to DUE allows for a more accurate representation of the traffic dynamics as they unfold over time. However, it also requires the development of an origin-destination (OD) table

for each time period during the planning horizon. We use 15 minute time periods over a three-hour time span (6AM-9AM) to represent travel during the peak period, about 560,000 vehicle trips. The translation of the OD table for the peak period into a series of OD tables (one for each period) is accomplished based on data in the Memphis Metropolitan Planning Organization Long-Range Plan (2011). The DTA algorithm employed is discussed at length in Li et al. (2012) and a comprehensive review of the DTA literature can be found in Peeta and Ziliaskopoulos (2001).

The consequence scenarios described in Section 2.4 give the physical damage to each bridge in the form of a damage level (or state). Before the DTA can be run for a post-event consequence scenario, it is necessary to translate this physical damage into a characteristic of the transportation network. This translation consists of modifying the capacities of the links that cross over or under the affected bridges. For roadways that are on bridges, the capacities are reduced to 50% when there is slight damage and 0% when there is moderate, extensive, or complete damage. For roadways under a bridge, the capacity is reduced to 0% when the damage level is moderate or greater. In the event that a roadway interacts with several bridges, the most restrictive capacity is used.

The DTA algorithm calculates the time required to travel from each origin to each destination by departure time, or gives an indication that the destination can no longer be reached. For the hospital connectivity measure we use the 75th percentile of the travel time distribution for each origin to the nearest hospital. The travel time to the nearest hospital from a given origin is a distribution, because under each consequence scenario the damage to the network is different and the travel times vary over the morning peak. It is also possible that for a given origin, the closest hospital changes as the traffic fluctuates over time, even within the same consequence scenario. In computing the hospital connectivity measure, we use a piecewise linear penalty term that is weighted by zone importance and applies increasingly larger penalties to discourage longer travel times. For example, travel times between 0-8 minutes are penalized with a weight of 0.1 since this is an acceptable travel time to reach a hospital. Travel times that are between 8-16 minutes are penalized with a weight of 1, and this weighting increases as the travel times increase. When no hospital can be reached for a given origin, a large penalty of 10,000 is used in place of the travel time.

2.6 Optimization Algorithm for Retrofit Selection

Given the identified earthquake scenarios and associated consequence scenarios (and their hazard-consistent probabilities of occurrence), an optimization algorithm using the objective functions described was developed to estimate the efficient frontier of retrofit policies. The goal of this optimization model is to select a set of bridges to reinforce so that the three objective functions are minimized. Since it is difficult to assign priorities to each of these objective functions *a priori* this algorithm takes the approach of constructing Pareto frontiers so that decision makers can understand the tradeoffs. This algorithm combines a Genetic Algorithm (GA) with local search methods to create a set of Pareto optimal points similar to the techniques referenced in Konak et al. (2006) and Deb (2012).

The gene used for this algorithm is simply a vector of binary decision variables, where each entry in the array corresponds to a bridge ID. Each individual in the initial population is initialized by randomly choosing the reinforcement decision for each bridge. For each individual, the performance attributes are assessed using the DTA algorithm. Each new solution can be checked against the set of Pareto optimal points P , and the set P can be updated as necessary by including non-dominated solutions as defined in Deb (2012). Non-Pareto solutions are also stored with their fitness measured by calculating their smallest Euclidean distance from any of the Pareto points in P using the three objective values as the solution space coordinates.

Each new pair of individuals is created via genetic crossover of two parents selected from the solution pool. A randomly selected cut point in each parent's gene determines the new child genes by combining the first part of the first parent's gene with the second part of the second parent's gene for the first child and vice versa for the second child. The selection of parents is random but is biased towards points near the frontier. Points on the frontier are three times as likely to be selected as points furthest from the fron-

tier, with the likelihood for intermediate points varying linearly between these two extremes based on the fitness score.

Mutation is employed to add diversity to the solution and counter the tendency for crossover to produce homogenous populations as described in Sait (1999). Using a custom implementation of the variable-rate of mutation strategy, mutations are based on the level of consistency amongst bridge reinforcement strategies across the population of solutions. The percentages T_H and $100-T_H$ are defined to represent thresholds for which a bridge is considered to be reinforced or not reinforced “most of the time”, respectively, where T_H was selected to be 25%. A *similarity ratio* is used to determine the percent of bridges that fall into either of these categories for the given population and provides a measure of the consistency of bridge reinforcement strategies across the population. As this ratio increases, the mutation rate increases up to a maximum value of 5%, if there is complete consistency across all bridges. The number of mutations that will occur is determined by multiplying the population size, the number of bridges and the mutation rate. Mutations are performed by iteratively selecting random members of the population with equal probability. Then a bridge is selected for mutation (state inversion from reinforced to not reinforced or vice versa) by randomly selecting from the consistent bridges in their consistent state. If no bridges meet this criteria, then any random bridge is selected for mutation. This process is repeated until all mutations have been made. To prevent a mutation from being negated, a bridge whose state has been modified cannot be changed in the same mutation cycle.

The standard GA is modified to include a local search step after each new individual has been created which is repeated multiple times until the desired number of solutions is obtained. After assessing a solution, the model explores neighboring solutions by selecting a single bridge and changing its reinforcement strategy. Since changing the reinforcement strategy for a single bridge is unlikely to change the damage state in all scenarios, only those scenarios where the damage state changes need to be reevaluated. This technique of *scenario differencing* provides a roughly 15% improvement in the efficiency of the algorithm. In order to maximize the chances of reaching a better solution given the large number of possible permutations (2^{286}), the choice of bridge to select is biased by the number of scenarios that will be impacted by changing the reinforcement strategy. For example, a bridge reinforcement change which impacts 30 scenarios is more likely to be selected over one that affects only 20 scenarios. This approach encourages the local search algorithm to focus on modifications that will have the largest potential impact on the network.

2.7 Investment Strategy Optimization Results

On a single compute node (8 Nehalem X557 processors running at 2.93 GHz with 12 GB of RAM), the DTA requires 12 minutes to evaluate a single solution with 40 scenarios. In order to generate a significant number of solutions in a reasonable amount of time, 100 processes (across 50 nodes) are executed in parallel. The initial estimate of the Pareto frontier is created by evaluating the 2 extreme strategies: all bridges reinforced and no bridges reinforced. Each process evaluates a single strategy across all 40 scenarios then does a local search where it generates 10 additional neighboring solutions. The master process updates its Pareto Frontier with all solutions collected from each child process. The collection of new solutions are ranked based on their Pareto fitness (as described in Section 2.6) which is used to bias the selection of solutions used in the generation of new mitigation strategies via crossover and mutation. Each sub-process is sent a new mitigation strategy, along with the updated frontier, so that another evaluation can be executed. This procedure is repeated 10 times by the optimization engine for a total runtime of 18 hours and 22 minutes. Figure 9 illustrates the frontier across all 3 objectives where the area of each data point is proportional to the mitigation cost. In order to illustrate the results for a few select solutions, the following four solutions are indicated on Figure 9: no bridges mitigated, all bridges mitigated, a solution where 83 bridges were mitigated, and a solution where 185 bridges are mitigated.

Table 1 compares the objective values for each of these solutions and gives the pre-event travel time and hospital objective when there is no damage to the network. The no mitigation and all-mitigated cases

give an upper and lower bound on these two objectives, respectively. The all-mitigated case also shows the lowest possible deviation from the pre-event case given the current mitigation options. The travel time objective for the all-mitigation case is nearly twice the pre-event case, implying that when all bridges are mitigated the cumulative travel time of all vehicles is still doubled, post-event.




Figure 9: Pareto Frontier, Cost, Travel Time and Hospital Connectivity Objectives

The results in shown in Table 1 and Figure 9 provide several insights when considering an investment strategy. First they help the decision maker bound the potential impact by showing consequences of the “invest in everything” and “invest in nothing” strategies. A decision maker could also use these results to understand the best possible improvements for a given budget or conversely the minimum budget that is required to meet target performance levels in one or both of the remaining objectives. Finally, for a fixed cost level the decision maker could explore the tradeoffs between the two remaining objectives.

Table 1: Comparison of Objective Values for Four Select Mitigation Strategies.

Objective	Pre-Event (No Damage)	No Mitigation	83 Bridges	185 Bridges	All Mitigated
Cost (\$ thousands)	N/A	\$0	\$49,650	\$100,400	\$194,500
Travel Time (million vehicle minutes)	6.0	13.5	12.7	12.3	12.1
Hospital Connectivity	1,106.16	3,278.25	1,505.10	1,315.08	1,253.13

While the meaning of the cost and travel time objectives are intuitive, the meaning of the hospital objective is more difficult to interpret. Table 2 illustrates how the hospital objective relates to individuals. The first column shows bins of travel times. Since the travel time to the nearest hospital varies from 6AM to 9AM, the 75th percentile of travel times over that timeframe was used to represent the travel time for each zone. The pre-event column shows the population that can reach a hospital within each interval of travel times, based on population data available from the U.S. Census Bureau. The columns for each of the four mitigation strategies show the change in the number of people who will be able to reach the hospital for each travel time interval after the event. Observe that for all four mitigation strategies, the population that is able to reach the hospital in 0 to 8 minutes is decreased, and is increased for all other time

intervals. As more bridges are reinforced, the population in 0 to 8 minute time interval increases, whereas it decreases in the other time intervals. When no bridges are mitigated, nearly 95,000 people are unable to reach a hospital in 0 to 8 minutes. Mitigating all bridges reduces this number to about 27,000 people.

Table 2: Comparison of Hospital Connectivity of Four Mitigation Strategies to Pre-Event Connectivity.

75 th Percentile of Travel Times to Hospital	Expected Population	Difference from No Damage Case, Post-Event			
		Pre-Event	No Mitigation	83 Bridges	185 Bridges
0-8 Minutes	853,855	-94,868	-72,440	-49,280	-26,656
8-16 Minutes	216,712	+75,791	+56,684	+40,071	+20,207
16-24 Minutes	69,646	+8,918	+8,171	+5,562	+4,654
24-32 Minutes	17,379	+6,933	+5,924	+3,096	+1,489
32-40 Minutes	3,046	+1,572	+1,292	+365	+129
40+ Minutes	0	+172	+141	+15	+10
Disconnected	0	+1,366	+112	+55	+52

Finally, Figure 10 compares the expected traffic flow for the no mitigation case to the case when 185 bridges are reinforced, for a small section of the road network near the Mississippi River. The green lines with black outlines represent the roadways. The thickness of these lines corresponds to the expected total traffic flow over each link during the given timeframe. The solid and hollow black squares represent bridges that are and are not reinforced, respectively, for the given solution. A red line is drawn along Interstate 55 to indicate the roadway of interest. In the solution where 185 bridges are mitigated, nearly all bridges along this route are mitigated, which increases the expected flow on this section of roadway.




Figure 10: Comparison of the Expected Traffic Flow for a Section of Interstate 55 for the No Mitigation (left) and 185 Bridge Mitigation Solutions (right), over the AM peak period.

3 CONCLUSIONS

In order to provide decision makers with a useful spectrum of options for choosing how to maximize infrastructure investments across multiple objectives, the development of a Pareto frontier is essential. While the techniques describe above are illustrated for a particular hazard and infrastructure, they were developed with the intention that they could be applied to other infrastructures (i.e. electric power) and hazards (i.e. hurricanes) as long as a systems model and efficient techniques for characterizing the hazard

and resulting consequences exists. This paper demonstrates that even when infrastructure simulation models are required to assess system performance, the combination of an efficient hazard model and a standard GA with an intelligently-biased local search and variable rate of mutation, can generate a useful set of Pareto solutions on a large-scale, real-world problem.

REFERENCES

Chang, L., Peng, F., Ouyang, Y., Elnashai, A., and Spencer, B., 2012, "Bridge Seismic Retrofit Program Planning to Maximize Postearthquake Transportation Network Capacity", *Journal of Infrastructure Systems*, 18, 75-88.

Chang, S., M. Shinozuka, and J. Moore. 2000. Probabilistic Earthquake Scenarios: Extending Risk Analysis Methodologies to Spatially Distributed Systems, *Earthquake Spectra* 16, 557-572.

Deb, K., "Advances in Evolutionary Multi-objective Optimization". 2012. *SSBSE 2012*. Edited by G. Fraser. Springer-Verlag, Berlin, Heidelberg (2012).

Federal Highway Administration, 2011. "National Bridge Inventory (NBI)". Accessed April 2013. <http://www.fhwa.dot.gov/bridge/nbi.cfm>.

FEMA (Federal Emergency Management Agency). 2010. HAZUS User and Technical Manual, Department of Homeland Security. Accessed March 14, 2011. http://www.fema.gov/plan/prevent/hazus/hz_manuals.shtm.

Furuta, H., Frangopol, D., and Nakatsu, K., 2011. "Life-cycle Cost of Civil Infrastructure with Emphasis on Balancing Structural Performance and Seismic Risk of Road Networks", *Structure and Infrastructure Engineering: Maintenance, Management, Life-cycle Design and Performance*, 7:1-2, 65-74.

Gearhart, J., N. Brown, D. Jones, L. Nozick, N. Romero and N. Xu, 2011, "Optimization of Scenario Construction for Loss Estimation in Lifeline Networks". *Proceedings of the 2011 Winter Simulation Conference*. Edited by S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu.

Gearhart, J., N. Brown, D. Jones, L. Nozick, N. Romero and N. Xu, 2013, "Optimization-based Probabilistic Consequence Scenario Construction for Lifeline Systems", accepted, *Earthquake Spectra*.

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.

Konak, A., D. W. Coit, A. E. Smith. 2006. "Multi-objective optimization using genetic algorithms: A tutorial". *Reliability Engineering and System Safety*, available online as of January 2006.

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.

Li, A., L. Nozick, R. Davidson, B. Wolshon, N. Brown, and D. A. Jones. 2012. "An Approximate Solution Procedure for Dynamic Traffic Assignment". *Transportation Engineering*, accepted, available online as of Nov. 2012.

Liu, C., Fan, Y., Ordonez, F. 2009. "A two-stage stochastic programming model for transportation network protection". *Computers & Operations Research*, 36, 1582-1590.

Memphis Metropolitan Planning Organization, 2011. "Long-Range Transportation Plan (LRTP)". Accessed January 2011. http://memphismpo.org/index.php?option=com_content&view=article&id=66&Itemid=53.

Peeta, S., A. K. Ziliaskipoulos. 2001. "Foundations of dynamic traffic assignment: the past, the present and the future." *Networks and Spatial Economics*, 1, 233-265.

Peeta, S., F. Salman, D. Gunnec, K. Viswanath. 2010. "Pre-disaster investment decisions for strengthening a highway network." *Computers and Operations Research*, 37, 1708-1719.

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.

Sait, S. M. and H. Youssef. 1999. *Iterative Computer Algorithms with Applications in Engineering: Solving Combinatorial Optimization Problems*. Los Alamitos, CA: IEEE Computer Society.

Shiraki, N., Shinozuka, M., Moore, J.E., Chang, S.E., Kameda, H., and S. Tanaka, 2007. System Risk Curves: Probabilistic Performance Scenarios for Highway Networks Subject to Earthquake Damage. *Journal of Infrastructures Systems* 13, 43-54.

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.

USGS (United States Geological Survey), 2008. 2008 NSHM Software. U.S. Geological Survey Website. Accessed June 2010. <http://earthquake.usgs.gov/hazards/products/conterminous/2008/software/>.

Vaziri, P., Davidson, R., Apivatanagul, P., Nozick, L., "Identification of Optimization-Based Probabilistic Earthquake Scenarios for Regional Loss Estimation", *Journal of Earthquake Engineering*, 16(2), 296-315.

AUTHOR BIOGRAPHIES

NATHANIEL 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. 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 manager of the Operations Research & Computational Analysis departments. His emphasis is on the research, design and application of mathematical models for use in the analysis of complex systems. 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.