Bayesian network for post-earthquake infrastructure risk assessment and decision

Michelle T. Bensi
U.S. Nuclear Regulatory Commission, Washington DC, USA

Armen Der Kiureghian
University of California, Berkeley, CA, USA

Daniel Straub
Technical University of Munich, Munich, Germany

ABSTRACT

A framework for post-earthquake risk assessment and decision making for infrastructure systems is developed. Use is made of the Bayesian network methodology for modeling earthquake hazards and system performance, as well as for probabilistic updating in light of an evolving state of information gained from ground motion sensors and observations of the system component states. The Bayesian network is extended by addition of decision and utility nodes to construct an influence diagram, which is used for post-earthquake decision making regarding the type of inspection to perform or for setting the performance level of components. A value-of-information heuristic is used to determine the best sequence of component inspections. The methodology is demonstrated by its application to a hypothetical model of a segment of the California high-speed rail system.

Keywords: Bayesian network, earthquakes, influence diagram, infrastructure

1. INTRODUCTION

The resilience of modern communities in face of the earthquake hazard strongly depends on the functioning of infrastructure systems, such as transportation, communication, water, and power networks. In the immediate aftermath of a damaging earthquake, decisions must be made regarding the deployment of emergency personnel and equipment, evacuation of people, inspection, closure or opening of facilities, and other actions to assure safety of people and mitigate losses. Furthermore, soon after the event, selections must be made among alternative actions to restore functionality to vital infrastructure services. The key ingredient for such decision-making is information about the nature and characteristics of the earthquake, the states of the system and its components, and the consequences of various decision alternatives. In the chaotic aftermath of a major earthquake, the available information is usually incomplete, highly uncertain, and rapidly evolving in time. Therefore, a probabilistic approach to assessing the system state and to decision making is essential.

In this paper we present a framework for post-earthquake risk assessment and decision making by use of the tools of Bayesian network (BN) and influence diagram (ID). A BN is a graphical model describing a set of random variables and their interdependencies (Pearl, 1988; Jensen and Nielsen, 2007). We use it to model the earthquake hazard as well as the infrastructure system and its components. An important advantage of this tool is its ca-

pability for probabilistic updating in light of acquired information, e.g., measurements from sensors placed on the ground or on the components of the system, or observations of damage states of system components. Another advantage is its transparent modeling paradigm, which allows verification of modeling assumptions by individuals who know the hazard and system characteristics, but may not be experts in probabilistic analysis. An ID is a BN extended by addition of decision and utility nodes. It allows modeling of various decision alternatives and codification of the consequences, measured in terms of utilities, for all possible combinations of decision alternatives and outcomes. The end result of analysis by the BN-ID model is a ranking of decision alternatives based on the fundamental principle of expected utility decision making. Of course, as the available information evolves, so may the ranking of the decision alternatives. Thus, the BN-ID framework serves as a dynamic decision support system for post-event decision making.

The paper begins with a brief introduction to the BN and ID methodologies. This is followed by the development of BN models for the seismic hazard, the system components, and the infrastructure system. These models are then enhanced by decision and utility nodes to yield the ID for decision making at the component and system levels. The methodology is applied to a hypothetical model of a segment of the proposed California high speed rail system. The example demonstrates how the probabilistic characterization of the system state and the preference ordering of decision alternatives change with an evolving state of information about the hazard and the states of the infrastructure components.

2. BRIEF ON BAYESIAN NETWORK AND INFLUENCE DIAGRAM

A BN is a directed acyclic graph consisting of a set of nodes representing random variables and a set of directed links representing probabilistic dependencies. Consider the simple BN in Figure 1. The directed links from X_1 and X_2 to X_3 indicate that the distribution of X_3 is defined conditioned on X_1 and X_2. In the BN terminology, random variable X_3 is said to be a *child* of random variables X_1 and X_2, while the latter are the *parents* of X_3. Similarly, X_4 is a child of X_1, while X_5 is a child of X_4. Aiming at near real-time applications, we assume the random variables are discrete or discretized, each having a finite set of outcomes or states. Exact inference algorithms are available for such *discrete* BNs (see, e.g., Jensen and Nielsen, 2007), which facilitate rapid computations.

Attached to each node of the BN is a *conditional probability table* (CPT), providing the conditional probability mass function (PMF) of the random variable given each of the mutually exclusive states of its parents. For *root* nodes that have no parents, e.g., X_1 and X_2 in Figure 1, a *marginal probability table* is assigned.

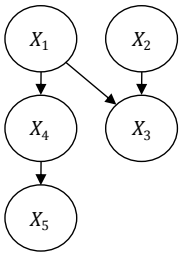

Figure 1. A Simple BN Model

The set of CPTs and marginal PMFs for root nodes completely define the probability distribution of the random variables. In fact, the joint PMF of the random variables can be written as

$$p(x_1, \dots, x_n) = \prod_{i=1}^{n} p(x_i|\mathrm{Pa}(X_i)) \qquad (1)$$

where $\mathrm{Pa}(X_i)$ is the set of parents of node X_i, $p(x_i|\mathrm{Pa}(X_i))$ is the CPT of X_i and n is the number of random variables (nodes) in the BN. BNs are particularly useful for answering probabilistic queries when one or more variables are observed. For exam-

ple, suppose the observations $X_3 = x_3$ and $X_4 = x_4$ have been made and the conditional PMF of X_2 given this information is of interest. This is computed by *marginalizing* the joint PMF in (1) to obtain $p(x_2, x_3, x_4)$ and $p(x_3, x_4)$, from which the desired conditional PMF is calculated as $p(x_2|x_3, x_4) = p(x_2, x_3, x_4)/p(x_3, x_4)$. The marginalization involves summation of the joint PMF in Eq. (1) over all states of variables not included in the set of interest. In a large BN, this process can be extremely time consuming. For this reason, special algorithms have been developed to perform the summation in an optimal manner using the decomposition in the right-hand side of Eq. (1) (see Jensen and Nielsen, 2007).

An ID is a BN to which *decision* nodes (shown as rectangles) and *utility* nodes (shown as diamonds) have been added. Nodes representing random variables (shown as circles) are now called *chance* nodes. The states of a decision node are decision alternatives. In our application, these are either action alternatives, e.g., shut down a system component, or test alternatives, e.g., inspect a system component. A test alternative facilitates the gathering of information before making a final action decision. Links coming into the decision node describe the state of information available to the decision maker. Thus, an incoming link from a chance node indicates that the decision maker knows the state of that random variable, whereas an incoming link from another decision node indicates that the decision is made with knowledge of the selected alternative of the preceding decision node. A utility node has no states; instead, it is assigned a utility value (e.g., monetary units) as a function of its parent decision and chance nodes, reflecting the utility associated with each combination of decision alternative and chance outcome. For example, a utility node may represent the cost of inspection, the cost of losing the service of a closed component, or the cost of liability if a damaged component is kept open. A utility node cannot have children.

As an example, consider the simple IDs in Figure 2 concerning the decision whether or not to shut down component i of a system. Chance node S_i indicates the demand on the component and chance node C_i, which is a child of the demand node, indicates the damage state of the component. In the ID on the left, the decision is made with no prior information (there is no incoming link into the decision node). The value of the utility node depends on the state of the component and the selected decision alternative. Hence, if the component is in an intact state and decision is made to shut it down, there will be a loss associated with interruption of service, whereas if the component is damaged and decision is made to keep it open, there will be a loss associated with liability from making a potentially unsafe decision. The ID in the middle includes a link from node C_i into the decision node. Here, the decision is made with perfect information about the state of the component. In the ID on the right an imperfect observation is made of the component state, which is represented by chance node O_i. This may represent the result of a visual inspection of the component, providing an imperfect assessment of its state. This imperfect information is available to the decision maker. Note that the CPT of node O_i is identical to the so called *test likelihood matrix* (see Benjamin and Cornell, 1970).

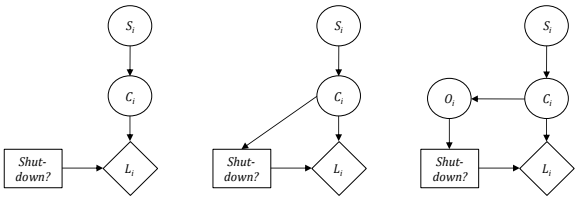

Figure 2. ID Models of Shutdown Decision for Component i

The IDs in Figure 2 include only one decision node with action-type alternatives, i.e., whether or not to shut down the component. A more comprehensive exam-

ple with multiple decisions is shown in Figure 3. Here, there is a precedence of decision nodes. First we must decide whether or not to inspect the component. The alternatives for such a decision might be not to inspect, to make a visual inspection, or to make an inspection using instruments. Each alternative has an associated inspection cost, as reflected in the utility node IC_i. The observation node O_i is now a child of the inspection decision node since the outcome of the observation will depend on the selected method of inspection. Furthermore, the decision whether or not to shut down the component is now made with knowledge of the type of inspection conducted and as well as the result of the inspection. These are respectfully indicated by the links from nodes $Inpsect?$ and O_i into the decision node.

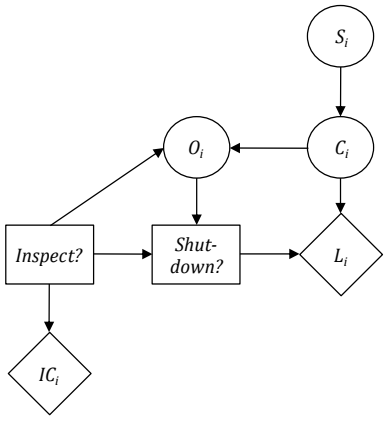

Figure 3. ID Model of Inspection-Shutdown Decision for Component i

3. BN MODEL OF SEISMIC DEMAND

A series of increasingly refined BN models of seismic hazard, including the effects of ground shaking, liquefaction and fault rupture, are presented in Bensi et al. (2011a). Here, we only present one conceptual model for ground shaking without providing the mathematical details, which can be found in the above reference. All continuous random variables are discretized in order to facilitate the use of exact inference algorithms. To enhance clarity of the presentation, we make use of BN objects, which are portrayed as rectangular nodes with rounded corners and represent BNs that operate in the background with specified input and output nodes.

The intensity of ground shaking at the site of a component i during an earthquake is described by a predictive model (commonly referred to as a ground motion prediction equation or attenuation law) of the form

$$\ln S_i = f(M, R_i, \mathbf{V}_i) + \epsilon_M + \epsilon_{R_i} \quad (2)$$

where S_i denotes a measure of the intensity of ground shaking, such as spectral acceleration at a selected frequency, M is the magnitude of the earthquake, R_i is the closest distance to the earthquake source, $\mathbf{V}_i$ denotes additional site-specific variables that may influence the intensity at the site (e.g., properties of the soil at the site), the function $f(M, R_i, \mathbf{V}_i)$ describes the logarithmic mean of the intensity and ϵ_M and ϵ_{R_i} are model error terms: ϵ_M represents the inter-event error term and describes the variability of the model error from event to event; ϵ_{R_i} represents the intra-event error term and describes the variability of the model error from location to location for the same event (see, e.g., Abrahamson et al., 2008). These are both zero mean normal random variables with variances that depend on M. In an earthquake, ϵ_M is the same for all locations, while ϵ_{R_i} vary from location to location and are spatially correlated.

Figure 4 shows the global BN model of the seismic hazard for ground shaking. The model has a single root node, *Source*, whose states describe the potential earthquake sources, i.e., the seismogenic faults in the spatial domain affecting the infrastructure system. These are modeled as straight lines or connected segments of straight lines, though more refined models can be developed. One child of this node is the *Fault Geometry*, which describes the geometric coordinates of each source fault. Another child is node M, which describes the magnitude of the earthquake. The dis-

tribution of M depends on the type and size of the fault, thus the reason for the dependence on the *Source* node. We assume the earthquake manifests as a rupture along the fault, originating at a random point X_{epi}, the epicenter. This is modeled by node X_{epi}, which is a child of the *Fault Geometry* node. All four nodes are parents to a BN object named *Rupture Location*.

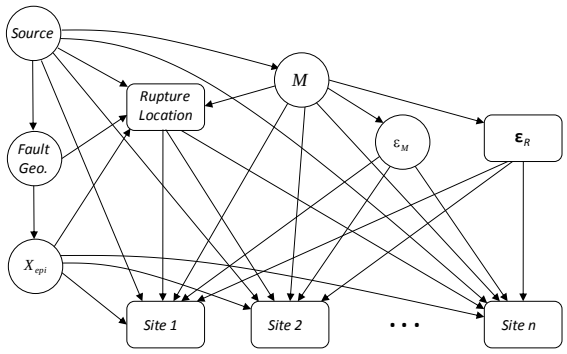

Figure 4. BN Model of Seismic Demand

Figure 5 shows the *Rupture Location* object. Inputs into this object are the four nodes mentioned earlier; they are highlighted with dashed borders. Given the *Fault Geometry* and M, one can determine the length L of the potential rupture from a regression model of the form $\ln L = a + bM + \epsilon_L$, where a and b are the regression coefficients and ϵ_L is the random model error term (Wells and Coppersmith, 1994). Usually, L is no greater than half the total fault length, hence the dependence of node L on the *Fault Geometry*, as depicted in Figure 5. In the BN object, nodes a, b and ϵ_L together with the *Fault Geometry* and M define the rupture length node L. We assume the rupture occurs randomly along the fault on either or both sides of the epicenter without extending beyond the known ends of the fault. This determines the distribution of the location of the rupture, which we characterize by specifying the coordinates of the end points of the rupture, X_{left} and X_{right}, as children of nodes *Fault Geometry*, rupture length L and epicenter location X_{epi}. The output nodes from this object are X_{left} and X_{right}. This is highlighted by the bold borders of the corresponding nodes. These are used as inputs to every site object, as described below.

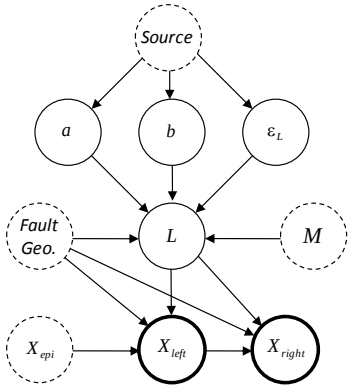

Figure 5. *Rupture Location* BN Object

The object $\boldsymbol{\epsilon}_R = (\epsilon_{R_1}, \epsilon_{R_2}, \ldots, \epsilon_{R_n})$ shown in Figure 4 represents the intra-event error terms for the considered sites. Since these are correlated random variables, the BN represented by this object should include links between all pairs of nodes ϵ_{R_i} and ϵ_{R_j}, $i \neq j$. This means that at least one of these nodes will have n parents, with n denoting the number of component sites. Naturally, for a system with a large number of component sites, the CPT associated with that node becomes unfeasibly large. This is a drawback of the BN methodology in modeling correlated random variables. In Bensi et al. (2011b), we have addressed this problem by developing a method to eliminate unimportant links, while minimizing the error in the representation of the correlation matrix. The interested reader should consult the above reference.

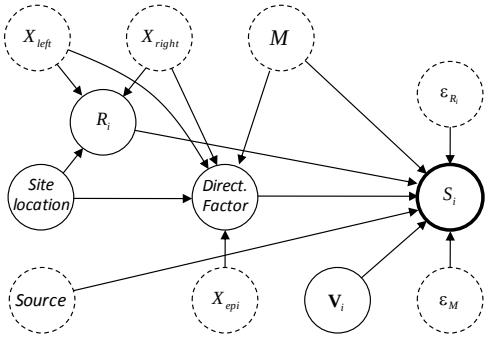

Figure 6. BN Model of Site Object

Finally, Figure 6 shows the BN object for the ith site. Input nodes into this object are the magnitude, M, the coordinates X_{left} and X_{right} of the end points of the rupture, and the error terms ϵ_M and ϵ_{R_i}. The *Source* node may also be an input, since the regression formula $f(M, R_i, \mathbf{V}_i)$ may depend on the type of faulting. The object also includes a node for the additional site-specific variables $\mathbf{V}_i$, e.g., the shear-wave velocity of the site. Given the coordinates of the rupture and the *Site Location*, the distance R_i from the site to the nearest point on the rupture is computed. Nodes M, R_i, $\mathbf{V}_i$, ϵ_M, ϵ_{R_i} and $Source$ then determine the measure of ground motion intensity at the site, which is represented by the output node S_i.

Many refinements to the above model can be made. As an example, here we consider the effect of directivity of the fault rupture for a near-fault site. As is well known, under certain conditions, when the fault rupture propagates towards a site, constructive interference of seismic waves arriving from intermittent segments of the rupture may occur, resulting in a large-amplitude, long period velocity pulse at the site. This is known as the forward directivity effect (Somerville et al., 1997). Conversely, when the rupture propagates away from the site, the ground motion at the site is likely to have a smaller amplitude but longer duration. In the current practice, these effects are accounted for by applying a correction factor to the attenuation law in Eq. (2) (Somerville et al., 1997; Abrahamson, 2000). These factors in general depend on the geometry of the site relative to the fault rupture, principally the angle θ between the fault line and the line connecting the site to the epicenter, and the length of the rupture propagating towards the site. The BN model in Figure 6 includes this factor. As can be seen, the object also include X_{epi} as an input node. This is necessary not only for determining the angle θ, but also for determining the length of the segment of the rupture that propagates towards the site. The *Directivity Factor* node is now an additional parent of node S_i.

4. BN MODEL OF COMPONENT PERFORMANCE

Suppose component i of the system can be in one of m_i states, depending on the demand to which it is subjected. Following the standard approach in earthquake engineering, we define the conditional probabilities associated with the states of the component in terms of a set of fragility functions. Each fragility function describes the conditional probability of the component reaching or exceeding a certain state, given the seismic demand (intensity of the ground motion). Figure 7 exemplifies the set of fragility functions for a component having $m_i = 4$ states (e.g., intact, light damage, heavy damage, failed). The probability that the component is in any particular state is the difference between the two bounding fragility curves. Since these probabilities are given conditional on the intensity of ground motion, S_i, they are identical to the probabilities needed for the CPT of a component node that is a child of the seismic demand node. Figure 8 shows the global BN model of the system with the seismic demand and system performance models shown as objects. Of course a discretization of the seismic demand and the corresponding fragility functions is necessary to yield a BN with discrete nodes.

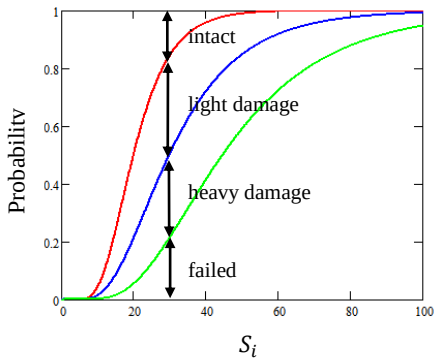

Figure 7. Component Fragility Functions

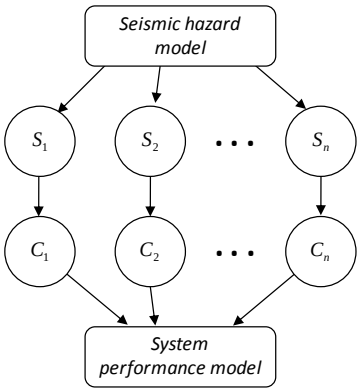

Figure 8. Global BN Model of Infrastructure System Subject to Seismic Demands

5. BN MODEL OF INFRASTRUCTERE SYSTEM

The state of a system is a function of the states of its components. For example, given the states of bridges, tunnels, and roadways of a highway system, one can determine if travel between any set of origin and destination points is possible. Using this relationship, a simple BN model describes the system performance node as a child of all the component nodes. This formulation, however, is inefficient since for a large system the system performance node will have many parents. For two-state systems with two-state components, a more efficient formulation is to introduce intermediate nodes representing minimum-link sets (MLSs) or minimum-cut sets (MCSs) of the system. A MLS (MCS) is a minimum set of components whose joint survival (failure) constitutes survival (failure) of the system (Pagès and Gondran, 1986). However, with growing number of components, the sizes and numbers of MLSs or MCSs also grow, yielding a formulation that is also inefficient. In Bensi et al. (2011a, 2013), we have developed a topology optimization scheme for constructing efficient BN models of two- and multi-state systems. Using MLSs (MCSs), the BN is modeled by introducing survival (failure) path events, which are a chain of events such that the state of each event depends on the state of the preceding event in the chain as well as the state of an associated component. The result is a chain-like BN model, where nodes have few parents. More details can be found in the cited references.

6. ID OF INFRASTRUCTURE SYSTEM

An ID at the system level is needed in order to support post-earthquake decision making for the infrastructure. In particular, we wish to make decisions on whether or not to inspect each component, and whether or not to reduce or shut down the operation of each component to avoid further losses, with or without first inspecting the component. Since usually there are limited resources for inspection of components, it is also necessary to prioritize the order of component inspections.

There are several types of IDs depending on how the precedence between decision nodes is handled. A *perfect recall ID* is the conventional form, in which the temporal sequence of all decisions is prescribed, i.e., there exists a directed path through the ID that contains all the decision nodes. Such IDs are based on a "no forgetting" assumption (Jensen and Nielsen, 2007), i.e., when making each decision the decision-maker remembers all preceding observations, outcomes and decisions. For our application, the temporal sequence of decisions is not known in advance. In fact, determining the order of inspections is a desired outcome of the decision problem. *Limited memory IDs* (LIMIDs) and *unconstrained IDs* (UIDs) relax the requirement that there be an explicit temporal ordering of the decision nodes. A LIMID drops the no-forgetting assumption and instead assumes that only nodes that are explicitly represented as parents to a decision node are known at the time the decision is made. The LIMID solves decision problems within smaller domains; therefore, its solution is likely to be suboptimal (Jensen and Nielsen, 2007). UIDs address decision problems in which not only the optimal choice for each decision, but also the best temporal ordering of

the decisions, are of interest (Jensen and Vomlelova, 2002). This is relevant to our problem. However, UIDs experience exponential growth in complexity with the number of decision nodes when using exact solution algorithms. Hence, they are not practical for large systems, particularly in near-real time applications. In this study, we utilize LIMIDs. However, we invoke a value-of-information heuristic to decide on the temporal ordering of the inspections. Efficient algorithms and software for analysis of LIMIDs are available (e.g., see DSL, 2007; Hugin Expert A/S, 2008).

When solving a LIMID, we first determine a *policy* for each decision node that maximizes the expected utility for any given configuration of the states of its parent nodes. The set of utility-maximizing policies for all decision nodes in a LIMID is referred to as a *strategy*. Because LIMIDs do not require temporal ordering of the decisions, locally optimal solutions are obtained using an iterative procedure known as *single policy updating*. The algorithm begins with an initial strategy that is typically random. A cycle of the algorithm updates policies for all decision nodes in the LIMID. Converges is achieved when the expected utilities associated with successive cycles remain the same (Lauritzen and Nilsson, 2001).

Figure 9 shows the system-level LIMID for decisions concerning component inspection and performance level. For the sake of clarity, the models of the seismic hazard and system performance are shown as BN objects. The ID for each component is similar to that shown in Figure 3, except that a component capacity node, Cap_i, has been added as a child of the component state and the *Shutdown?* decision node. The latter node includes decision alternatives concerning the operational level of the component. The alternatives might be to operate at full capacity, to operate at reduced level, or to shut down. The system BN object is shown as a child of the component capacities. This object corresponds to an efficient chain-like BN, as described in the preceding section. Attached to the node describing the state of the system (the output from the system BN object) is a utility node. This node contains the utility values associated with each performance level of the system. Of course these utilities are exclusive of the costs associated with the states of individual components, which are separately accounted for in the local utility nodes attached to each component. For example, if we are dealing with a highway system, the system utility node may describe the cost of not being able to travel between selected origin and destination nodes, while individual component utility nodes describe the local cost of closing down a bridge or having to repair it. Observe that the LIMID does not include precedence links between the inspection decision nodes. Hence, we need to develop an approach to prioritize the temporal ordering of the component inspections.

In decision analysis, the value of information (VoI) is commonly used to quantify the benefit gained from acquiring additional information before making an action decision (Raiffa and Schlaifer, 1961; Straub and Der Kiureghian, 2010a, 2010b). The VoI is defined as the gain in expected utility from obtaining the additional information, e.g. through an inspection or a test. It is calculated prior to obtaining the outcome of the inspection or test, i.e., it is an expected value with respect to the observation outcome. The value of the additional information lies in the enhanced likelihood of selecting a better decision alternative.

Here, we use the VoI concept to determine a temporal ordering of the component inspections. First consider the component-level decision problem depicted in Figure 3. Let $EU|NI$ be the maximum expected utility for the available decision alternatives given no information, i.e., when the decision node $Inspect?$ in Figure 3 is set equal to the no-inspection

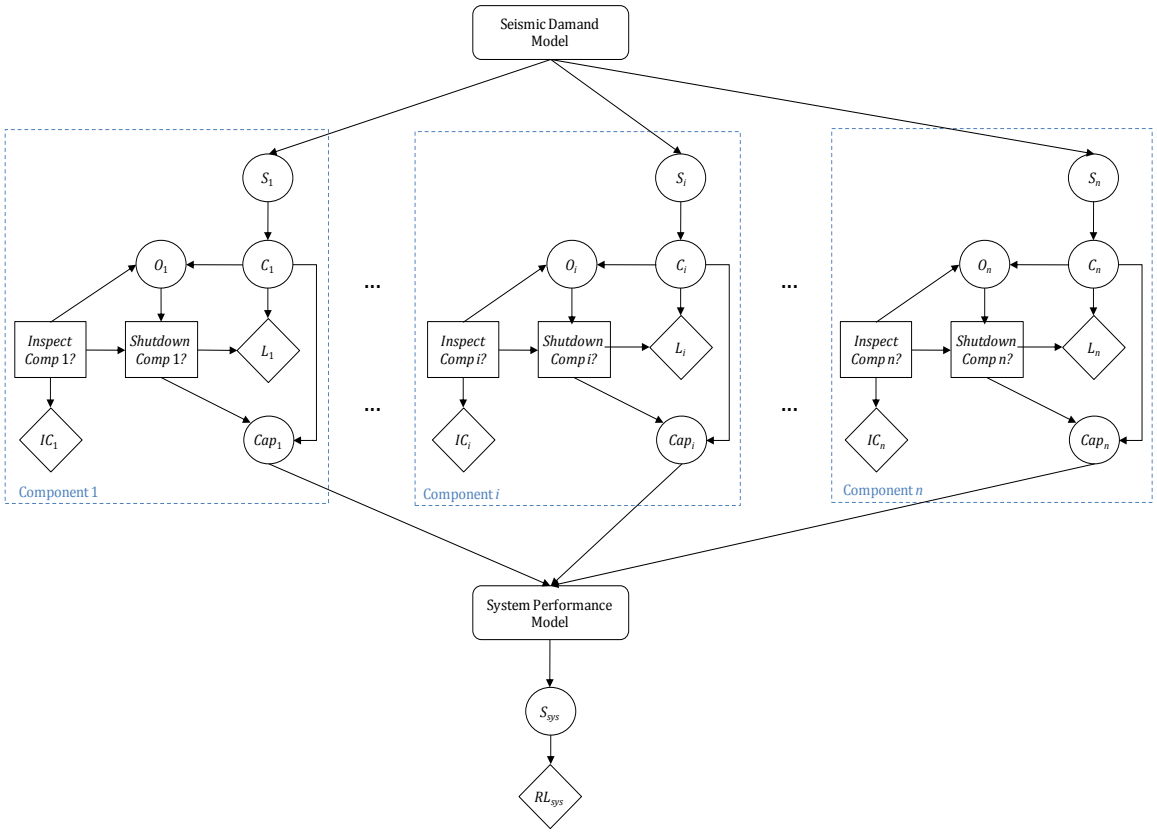

Figure 9. ID of Infrastructure System

option. Similarly, define $EU|I$ as the maximum expected utility for the available decision alternatives given an inspection option, excluding the cost of the inspection, i.e., when setting $Inspect?$ in Figure 3 to an inspection option and setting the value of node IC_i equal to zero. The expected VoI for the inspection option then is

$$VoI = EU|I - EU|NI \tag{3}$$

An inspection is worth performing if its cost is less than its VoI. Among a set of inspection alternatives, the one that has the largest positive difference between the VoI and the cost of inspection is the optimal alternative, i.e., the one with the largest net benefit of inspection (NBI).

To help prioritize post-earthquake inspections of components at the system level, we propose a VoI heuristic. The heuristic is defined such that, at each stage, the decision-maker looks for the next best component to inspect. This is determined as the component that has the highest NBI at the system level. The system-level component NBI values are computed as follows: First, the LIMID in Figure 9 is analyzed to determine the (locally) optimal inspection option for each component. The solution may include cases where the optimal decision is not to inspect the component (i.e., the cost of inspection outweighs the potential benefits of gaining additional information, such as would occur when the available information about the hazard suggests high likelihood for one particular state of the component) and cases where the optimal decision is to inspect the component using one of the available inspection options. Next, the total system-level expected utility is computed, while all inspection decisions are set to their locally optimal values. This corresponds to the system-level utility associated with making the (locally) optimal inspection decision for each component. Next, for component i requiring inspection, the inspec-

tion decision node is set to the "no inspection" state and the corresponding cost IC_i is set to zero, while the remaining components are left at their locally optimal inspection decisions. The expected utility for the system is again computed. The system-level NBI for component i is the difference between the two expected utility values. The procedure is repeated until the system-level NBI for each component is determined. These system-level NBI values are then used to rank the temporal order of component inspections.

Using the LIMID methodology, the inspection prioritization order evolves as new information becomes available. Sources of information may include measurements of ground motion intensity and data from structural health monitoring sensors. As inspections are performed and decisions are made regarding the operational level of components, this information is entered into the BN-ID model. The information propagates through the network to provide an updated probabilistic characterization of the system model and decision alternatives. At any stage, the recommendations for component inspections may differ from those made previously, e.g., a component previously deemed to require an inspection may no longer need one, and vice versa. Thus, the LIMID provides the decision-maker with guidance on optimal decisions relating to inspection and component closure, at any point in time, based on all available information up to that time.

It is noted that the above heuristic does not consider all possible orderings of the inspection and shutdown decisions explicitly, a task that is computationally expensive and possibly intractable for large systems or for near-real time decision making. Therefore, the proposed approach may arrive at a suboptimal solution. However, considering the need to repeatedly solve the decision problem after each inspection, possibly in near-real time, the proposed heuristic facilitates a computationally efficient algorithm that is likely to produce a solution near the global optimum.

7. EXAMPLE APPLICATION

As an illustration of the proposed BN-ID framework, we consider post-earthquake risk assessment and decision making for a hypothetical model of a segment of the proposed California high-speed rail (HSR) system, which is under development. Our model is a highly idealized and simplified one; the specifics of the system components are imagined rather than being based on existing structures. For this reason, the application and the results derived from it should be considered as a conceptual illustration of what can be achieved with the proposed framework rather than viewed as results relevant to the actual HSR system, if and when it becomes a reality. Furthermore, for the sake of brevity, only novel aspects of the model are described here; details of other, standard models can be found in Bensi et al. (2011a).

We consider the northern segment of the proposed HSR system, from San Francisco to Gilroy. This segment is situated in a highly seismic region with several active faults, including the San Andreas, Hayward and Calaveras, in the vicinity. Figure 10 shows the configuration of the system and the seismic environment. Models consistent with the current state of practice are used to describe the seismic sources, rupture-length/magnitude relations, ground motion prediction equations, and spatial correlation of ground motion intensity, as described in Bensi et al. (2011a). The hypothetical system consists of different types of components along the route from San Francisco to Gilroy, including tunnels, embankments and aerial structures (bridges). These are idealized into 2 point-site components and 17 distributed-site components. Ground motion prediction points (GMPPs) along the path, shown in Figure 10, define the locations of the point-site components and end points of the distributed components. As an exam-

ple, component 1 is a tunnel having 4km length and stretching from GMPP 1 to GMPP 2. The two point-site components are located at GMPPs 12 and 14. Three states for each component are considered: undamaged, slightly damaged, and moderately/severely damaged. Undamaged components are assumed to be fully operational. The slightly damaged state is associated with reduced performance, i.e., trains must slow down when traversing the component. Moderate or severe damage implies complete loss of service. Fragility models for these component states are adapted from the literature, as described in Bensi et al. (2011a). We emphasize that these fragility models are hypothetical and do not reflect the seismic capacities of the future HSR system components.

We define the system performance in terms of the ability to travel from San Francisco to Gilroy and from there on to Southern California. However, we assume that if a northern segment of the system is not passable, passengers can use alternative transportation means to bypass the damaged segment and board the train further down the line to complete their journey to Southern California. With this in mind, four system performance criteria are considered: ability to travel with normal/reduced speed (1) from San Jose to Gilroy, (2) from Palo Alto to Gilroy, (3) from Millbrae to Gilroy, and (4) from San Francisco to Gilroy. Travel at normal speed between San Jose and Gilroy is possible if none of the components 15-19 is damaged; travel must be at reduced speed if any component is slightly damaged; and ability to travel is lost if any component is moderately or severely damaged. Travel from Palo Alto to Gilroy similarly depends on the states of components 8-14, but also on the state of travel from San Jose to Gilroy. Similarly, travel from Millbrae to Gilroy depends on the states of components 5-7 as well as the state of travel from Palo Alto to Gilroy, and that from San Francisco to Gilroy depends on the states of components 1-4 and the state of travel from Millbrae to Gilroy. It should be clear that the system cost of a damaged component farther down the line is greater because of the inability to travel to Southern California.

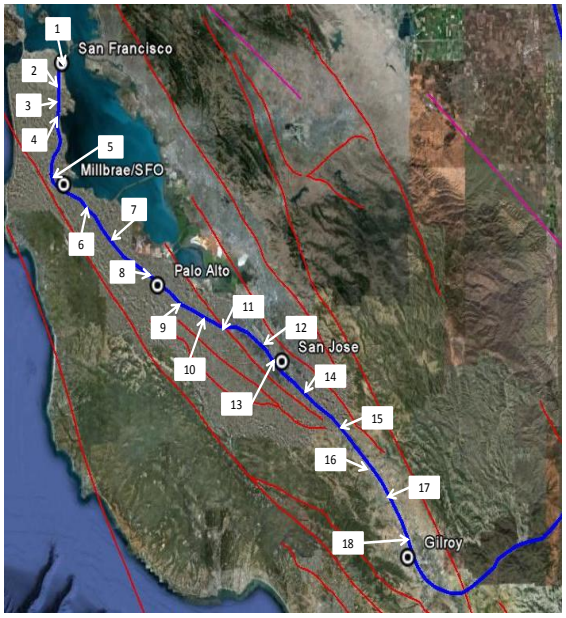

Figure 10. Configuration of the HSR System, the Seismic Environment and Ground Motion Prediction Points (map generated using Google Earth)

A main objective of this study is to demonstrate how the BN-ID framework allows updating of the seismic hazard model and component/system states and decision making under an evolving state of information. For this purpose, we consider four evidence cases (ECs) as follows:

1. EC 1: Shortly after the earthquake event, we learn from the Berkeley Seismological Laboratory that the earthquake had a magnitude $M = 6.8$ and an epicenter located on Hayward fault, 30 km from its north end.
2. EC 2: A little later we learn that a recording at GMPP 2 indicated a spectral acceleration of $S_2 = 0.45\text{-}0.50\text{g}$ at 1Hz frequency.
3. EC 3: A little later we learn that the ground motion at GMPP 9 was "pulse-like." This observation would suggest that the rupture on the fault propagated southward from the epicenter.

4. EC 4: A little later we learn that component 17 has experienced severe damage.

The above evidence cases are designed to monotonically provide "bad" news. Hence, one would expect that component and system failure probabilities progressively increase with the evidence cases. Of course this assumed trend is not necessary, but it is helpful in the demonstration.

We first consider the influence of the evolving information on the predicted seismic demand at selected GMPPs. These are obtained by entering evidences at the appropriate nodes of the BN (e.g., *Source*, M, X_{epi}, S_2) and updating the marginal distributions of S_i at selected sites. Figure 11 shows the updated distributions of the (discretized) seismic demand at GMPPs 1, 6, 11 and 16. First examine the distributions at GMPP 1. Under EC 1, the most likely value (mode) of the distribution is 0.1-0.15g (when disregarding the tail probability for the intensity interval 0.6-∞). After EC 2, the probability mass sharply moves towards higher values so that the most likely value now is 0.25-0.30g. This is because the observation of a high intensity at a nearby site (GMPP 2) provides evidence for a higher intensity at GMPP 1 due to the spatial correlation among the intra-event error terms ϵ_{R_i}. EC 3 does not affect the distribution of the intensity at GMPP 1. This is because, for GMPP 1, the distance from the site to the rupture is not affected by the direction of propagation of the rupture; furthermore, the position of the site excludes the possibility of a forward or backward directivity effect. For EC 4, we see further shifting of the probability mass towards higher values by a small amount. This is due to the observation of severe damage of component 17. Since that component is located far from GMPP 1, the effect is not due to the spatial correlation, but due to the correlation arising from the inter-event error term ϵ_M, which is common to all GMPPs.

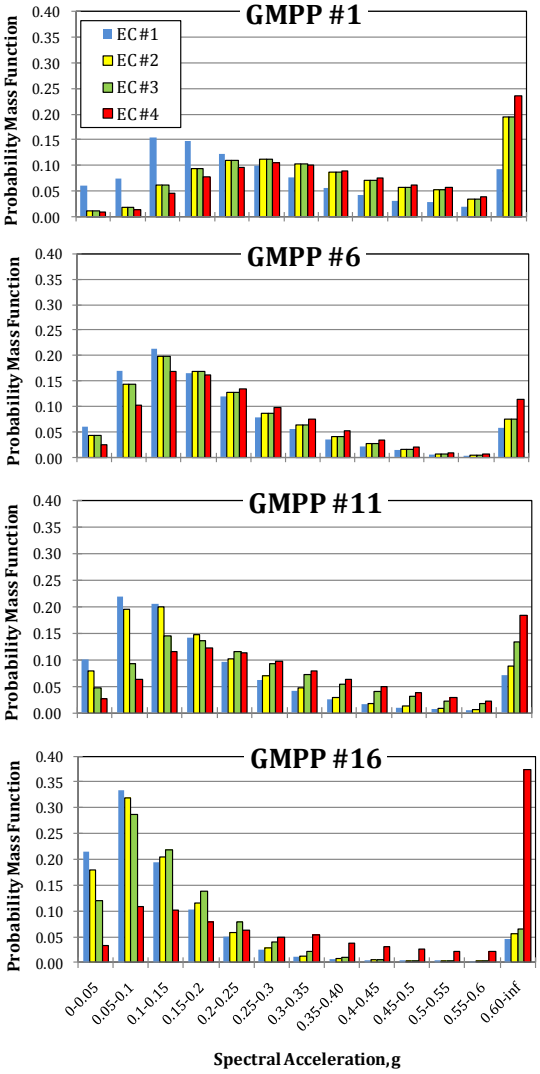

Figure 11. Updated PMFs of Spectral Acceleration at GMPPs 1, 6, 11 and 16

Next consider GMPPs 6, 11 and 16. For EC 2, we see shifts of probability distributions towards higher values, but the shifts are smaller than that for GMPP 1. This is because these sites are far from GMPP 2, where the high-intensity recording was made. In fact, under EC 2, these shifts are due to the effect of the common ϵ_M error term, not due to the spatial correlation of ϵ_{R_i}. For EC 3, strong shifts of the probability mass towards higher values are observed for GMPP 11 and 16, but not for GMPP 6. This is because the former sites are affected by forward directivity of the rupture. Finally, after EC 4, we see systematic shifts of the probability masses towards higher values at all sites (due to the

effect of the inter-event error term ϵ_M) and a particularly large shift for GMPP 16. This is because observation of severe damage of component 17 suggests high intensity ground motion at GMPP 16, which is located at one end of the damaged component 17.

Next, we examine the probabilities of moderate or severe damage of each component under the evolving state of information, shown in Figure 12. In light of EC 2, the damage probabilities sharply increase for components 1-4, which are located near GMPP 2, but less so for components 5-19, which are farther away. The sharp increase for components 1-4 is primarily due to the effect of spatial correlation of the ground motion, while the smaller increase for the other components is due to the common ϵ_M term in the seismic demand model. Under EC 3, we see sharp increases of the damage probabilities for components 8-19, which are located in the forward directivity sites. Under EC 4, we see sharp increases of the damage probabilities for all components, especially those near component 17. The probability of moderate or severe damage of component 17 is of course 1.0 (beyond the limit of the graph)

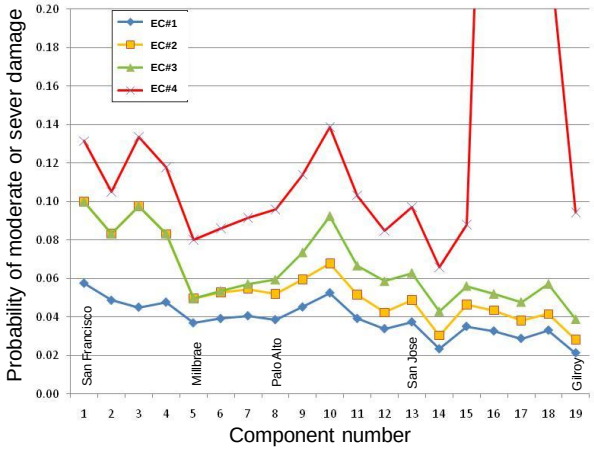

Figure 12. Updated Probabilities of Component Damage Under Evolving State of Information

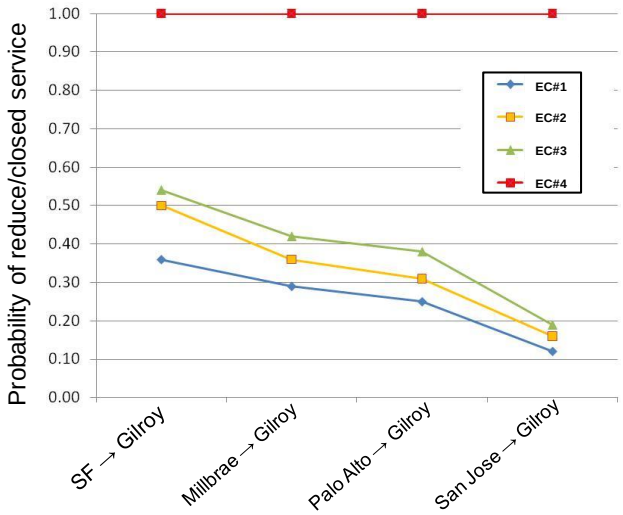

Figure 13. Updated Probabilities of Reduced or Closed Service Under Evolving Sate of Information

Figure 13 shows the probabilities of reduced or closed service from each city to Gilroy for the evolving state of information. Naturally, the probability of not having normal service to Gilroy decreases if one starts from a city further south from San Francisco. Note also that probabilities increase with increasing "bad" news. In particular, in light of EC 4, travel from any of the cities to Gilroy is impossible due to severe damage of component 17, which affects the section of the track between San Jose and Gilroy.

Now imagine that after the earthquake we need to decide whether to continue normal operation of each component, reduce the train speed while traversing it, close the component down, or delay this decision until the component is inspected. We also need to determine the optimal type and sequence of inspections. Since closing a component implies closing the entire segment between two cities, we first make decisions at the "Link" level, where Link 1 = San Francisco to Millbrae, Link 2 = Millbrae to Palo Alto, Link 3 = Palo Alto to San Jose, and Link 4 = San Jose to Gilroy. For this application, we consider two levels of inspection: perfect and imperfect, the latter specified by a test likelihood matrix.

To make the above decisions, it is necessary to specify utility values for all pairs of decision alternative and outcome, both at the component and link levels. We also need to specify inspection costs. Assumed values for these utilities are reported in Bensi et al. (2011a). These include costs for perfect and imperfect inspection options for each component, liability costs associated with keeping a damaged component open, and losses associated with reducing the train speed over a link or closing the link, all in a common unit of utility.

Evaluating the LIMID for each link and invoking the VoI heuristic at the link level we obtain the NBI for each link and each evidence case. For the assumed utility values, the recommended decisions are as follows: Under EC 1 and EC 2, inspect Links 1-3, do not inspect Link 4, but reduce the train speed over it. Under EC 3, inspect Links 1 and 2, do not inspect Links 3 and 4, but reduce speed over them. Under EC 4, inspect Links 1-3, close Link 4. The no-inspection cases correspond to situations, where the state probabilities for component states for the link are narrowly concentrated in one set of states and, hence, the cost of inspection outweighs the value gained from inspecting the link.

Once the priority ranking of link inspections is determined, we again use the VoI heuristic to rank inspection of components within each link. As an example, the recommended order of inspection of the components of Link 1 under EC 2 is: 1, 3, 4 and 2. Of course these results depend on the assumed inspection, liability and loss of service costs and any change in these values may affect the recommended order.

8. SUMMARY AND CONCLUSIONS

A framework for post-earthquake risk assessment and decision making for infrastructure systems is developed. The framework makes use of Bayesian network (BN) and influence diagram (ID) as tools for probabilistic modeling, information updating, and decision making. This modeling approach is transparent, allowing the possibility for verification of the model by disciplinary experts who may not be knowledgeable in probabilistic methods. It is also general in the sense that it can describe general probabilistic dependencies and system configuration. Most importantly, the BN allows rapid processing of information to update probabilistic characterization of the hazard and system states.

IDs for decision making at the component and system level are developed. These allow decision making for setting the operational level of the component, or to inspect the component before making a decision on the operation level. At the system level, there is a need to determine the priority ranking of component inspections. For this purpose, a value-of-information heuristic is proposed, which essentially identifies the "next best" component to inspect. Although it may potentially generate sub-optimal solutions, the heuristic allows rapid calculations by use of available algorithms for limited-memory IDs.

The proposed framework is illustrated by its application to a hypothetical model of the California high-speed rail system. Since the system is still under development, highly idealized models of the system components are used. The results should be viewed as demonstrative of the proposed framework rather than indicative of the reliability and risk of the future high-speed rail system, if and when it becomes a reality. The application clearly shows the power of the BN-ID framework to process information and to prioritize decision alternatives based on updated probabilistic information.

Although the earthquake hazard was the focus in this paper, the proposed framework is equally applicable to other natural and man-made hazards. In each case, one needs to develop an appropriate BN model of the hazard and incorporate it together with models of the system and its components, similar to what was presented here.

ACKNOWLEDGMENTS

The authors acknowledge financial support from California Department of Transportation through Pacific Earthquake Engineering Research Center and from the National Science Foundation under Grant No. CMMI-1130061.

DISCLAIMER

Any opinions, findings and conclusions expressed in this paper are those of the authors and do not necessarily reflect the views of the United States Nuclear Regulatory Commission or the funding organizations.

REFERENCES

Abrahamson, N., 2000. "Effects of rupture directivity on probabilistic seismic hazard analysis," *Proceedings of the 6th International Conference on Seismic Zonation*, Earthquake Engineering Research Institute, Palm Springs.

Abrahamson, N., Atkinson, G., Boore, D., Bozorgnia, Y., Campbell, K., Chiou, B., Idriss, I.M., Silva, W., and Youngs, R., 2008. "Comparisons of the NGA ground-motion relations," *Earthquake Spectra*, 24, 45-66.

Benjamin, J. and Cornell, C.A., 1970. *Probability, Statistics and Decision for Civil Engineers*, McGraw-Hill, New York, NY.

Bensi, M., Der Kiureghian, A., and Straub, D., 2011a. "A Bayesian network methodology for infrastructure seismic risk assessment and decision support," *Report No. 2011/02*, Pacific Earthquake Engineering Research Center, University of California, Berkeley, March.

Bensi, M., Der Kiureghian, A., and Straub, D., 2011b. "Bayesian network modeling of correlated random variables drawn from a Gaussian random field," *Structural Safety*, 33, 317-332.

Bensi, M., Der Kiureghian, A., and Straub, D., 2013. "Efficient Bayesian network modeling of systems," *Reliability Engineering and System Safety*, to appear.

DSL, 2007. *GeNIe 2.0 by Decision Systems Laboratory*, http://genie.sis.pitt.edu/.

Hugin Expert A/S, 2008. *Hugin Researcher API 7.0*, Denmark: Hugin Expert A/S, www.hugin.com.

Jensen, F.V. and Nielsen, T.D., 2007. *Bayesian networks and decision graphs*, 2nd Edition, Springer Science+Business Media, LLC, New York, NY.

Jensen, F. V. and Vomlelova, M., 2002. "Unconstrained influence diagrams," in Darwiche, A., and N. Friedman (Eds.), *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, San Francisco, CA, 234–241.

Lauritzen, S.L. and Nilsson, D., 2001. "Representing and solving decision problems with limited information," *Management Science*, 47, 1235-1251.

Pagès, A. and Gondran, M., 1986. *System reliability*, Springer-Verlag, New York, NY.

Pearl, J., 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*, Morgan Kaufmann Publishers, San Mateo, CA.

Raiffa, H., and Schlaifer, R., 1961. *Applied statistical decision theory*, Cambridge University Press, Cambridge, U.K.

Somerville, P. G., Smith, N.F., Graves, R.W., and Abrahamson, N., 1997. "Modification of empirical strong ground motion attenuation relations to include the amplitude and duration effects of rupture directivity," *Seismological Research Letters*, 68, 199-222.

Straub, D., and Der Kiureghian, A., 2010a. "Bayesian network enhanced with structural reliability methods: Methodology," *Journal of Engineering Mechanics*, ASCE, 136, 1248-1258.

Straub, D., and Der Kiureghian, A., 2010b. "Bayesian network enhanced with structural reliability methods: Application," *Journal of Engineering Mechanics*, ASCE, 136, 1259-1270.

Wells, D., and Coppersmith, K., 1994. "New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement," *Bulletin of the Seismological Society of America*, 84, 974-1002.