Framework for Post-Earthquake Risk Assessment and Decision Making for Infrastructure Systems

Michelle Bensi[1]; Armen Der Kiureghian, M.ASCE[2]; and Daniel Straub[3]

Abstract: The Bayesian network (BN) and influence diagram (ID) are used to develop a framework for post-earthquake risk assessment and decision making for infrastructure systems. The BN is used to model the earthquake hazard and component and system performance, and to update these models probabilistically in light of information gained from ground motion or structural health-monitoring sensors, or from observation of system or component states. The BN is extended by addition of decision and utility nodes to construct an ID, which is used for post-earthquake decision making regarding the type of inspection to perform or for setting the operational levels of components. A value-of-information heuristic is proposed to determine the optimal temporal sequence of component inspections. The methodology is demonstrated by its application to a hypothetical model of a segment of the proposed California high-speed rail system. Although the focus in this paper is on the earthquake hazard, the methodology can be used for infrastructure risk assessment and decision making for many other kinds of hazards. The main contribution of the paper lies in the integration of BN and ID models for the hazard and the infrastructure system and its components in a logical framework that allows rapid decision making to mitigate infrastructure losses after a major hazard.
DOI: 10.1061/AJRUA6.0000810.

Author keywords: Bayesian network; Decision making; Earthquakes; Hazard; Influence diagram; Infrastructure; Risk analysis; Systems.

Introduction

Infrastructure systems, such as transportation, water and power networks, are backbones of modern societies. Their resilience in the face of natural and man-made hazards is vital for the well-being of communities. In the immediate aftermath of such an event (or series of events), e.g., an earthquake, hurricane, flood, tornado, or an act of sabotage, 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 the safety of people and to mitigate economic 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: information about the nature and characteristics of the hazard, about the states of the system and its components, and about the consequences of various decision alternatives. However, in the chaotic aftermath of a hazard event, the available information is usually incomplete, highly uncertain, and rapidly evolving in time. A probabilistic framework can deal with incomplete and uncertain information. However, to be applicable for post-hazard infrastructure risk assessment and decision making, such a framework must provide results in near–real time, say within hours, to be useful to first responders. Hence, the algorithms for probabilistic inference must be sufficiently fast, and the resulting tools must work near-autonomously following an event, without a requirement for intervention by experts in hazard assessment, system modeling, or risk assessment.

In this paper a mathematical framework is proposed for post-hazard risk assessment and decision making for infrastructure systems by use of Bayesian network (BN) and influence diagram (ID). Although the approach is general and can be applied to all kinds of hazards, in this paper the focus is on the earthquake hazard. Earthquakes are a dominant hazard to infrastructure systems in many parts of the world. When they occur, they can affect large regions and cause damage to all kinds of constructed facilities and distributed systems and give rise to other hazards, such as fire, landslide, and flooding. Recent earthquakes, such as the ones in Christchurch, New Zealand (February 11, 2011) and Tōhoku, Japan (March 3, 2011), are reminders of the types of catastrophic losses that can result from this hazard.

A BN is a graphical model describing a set of random variables and their interdependencies (Pearl 1988; Jensen and Nielsen 2007). It is used in this study to model the earthquake hazard and the infrastructure system and its components. An important advantage of this tool is its facility for probabilistic updating in light of received information, e.g., measurements of sensors placed on the ground or on the components of the system, or observation of damage states of system components. Another advantage is its transparent modeling paradigm, which allows verification of modeling assumptions by individuals who know the system or hazard characteristics, but may not be experts in probabilistic analysis.

An ID is a BN with the addition of decision and utility nodes. It allows the modeling of various decision alternatives and codification of their consequences, measured in terms of utilities. The end result of analysis by the BN-ID model is a ranking of decision alternatives based on the fundamental principle of maximum expected utility decision theory (Neumann and Morgenstern 1944) and the updated state of information. As the available information evolves and the BN is updated, the ranking of decision alternatives may also evolve. Thus, the BN-ID framework serves as a dynamic decision-support system for post–event decision making. As

[1]Civil Engineer, U.S. Nuclear Regulatory Commission, Washington, DC 20555-0001. E-mail: michelle.bensi@nrc.gov

[2]Taisei Professor of Civil Engineering, Univ. of California, Berkeley, CA 94720 (corresponding author). E-mail: adk@ce.berkeley.edu

[3]Associate Professor, Engineering Risk Analysis Group, Technische Universität München, 80290 München, Germany. E-mail: straub@tum.de

Note. This manuscript was submitted on January 9, 2014; approved on August 11, 2014; published online on October 8, 2014. Discussion period open until March 8, 2015; separate discussions must be submitted for individual papers. This paper is part of the ***ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering***, © ASCE, 04014003(17)/$25.00.

discussed previously, it is important that such a system process information in near–real time. For this reason, the focus of this study is on the so-called *exact inference algorithms* (Jensen and Nielsen 2007), which are computationally robust and facilitate decision optimizations. Furthermore, as noted previously, it is important that the system have an interface that is convenient for use by decision makers, who may not be experts in probabilistic analysis or decision theory. In this regard, the graphical nature of the BN-ID tool provides an ideal framework.

BNs and IDs have been employed in a wide variety of disciplines; however, they have not been used to model and optimize post–hazard decision making. BN models of earthquake hazard and risk are reported in Bayraktarli and Faber (2011) and Kuehn et al. (2011). Bayraktarli et al. (2011) use the BN to perform deaggregation of the seismic hazard. An application of BNs to near-real-time tsunami warning based on observation of earthquakes is proposed in Blaser et al. (2011), and a BN for automated identification of inundated areas during and following floods is developed by Frey et al. (2012). Finally, Straub and Der Kiureghian (2010a) combine BNs with structural reliability methods to develop a tool for reliability analysis of infrastructure systems under evolving information. The aforementioned applications only employ BNs and do not leverage the BN-ID framework to explicitly model and optimize decisions. Moreover, none of these BN models assesses the performance of infrastructure systems with explicit consideration of the dependence among components that is introduced by the spatial correlation in the hazard model. To the authors' knowledge, this paper, which is an expansion of their earlier report (Bensi et al. 2011a), is the first systematic application of the BN-ID framework to the earthquake hazard for spatially distributed systems.

As with any tool, the BN-ID framework has its limitations. One is the large computational and memory demands when using exact inference algorithms with large numbers of dependent random variables. This poses a special difficulty in modeling spatially distributed hazards, such as the random field of earthquake ground motion. Another challenge arises when modeling systems with a large number of interdependent and interconnected components. The computational challenges arising from the use of the BN framework for modeling random fields and system performance have been addressed by the authors in earlier publications (Bensi et al. 2011b, 2013). In this paper, a refined BN model of the seismic hazard (including random field and directivity effects) is developed, combined with the infrastructure system model, and expanded into an ID to solve decision problems. In terms of the BN-ID framework, a major computational difficulty arises when the optimal sequence of several decision actions must be determined, e.g., the order of inspection of system components. A practical solution to this problem is proposed by use of a heuristic based on the concept of value of information (VoI).

Because BNs and IDs have rarely been used in earthquake engineering and infrastructure risk assessment, the readers who would most benefit from this paper may not be familiar with them. Hence, this 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, system components, and the infrastructure system. These are then combined and extended by decision and utility nodes to yield the ID for decision making at the component and system levels. Methods to address the optimal ordering of decision actions are reviewed and an approximate solution in terms of VoI is proposed. This methodology is then 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 damage states of the infrastructure components.

Brief on Bayesian Network and Influence Diagram

A BN is a directed acyclic graph that consists of a set of nodes representing random variables and a set of directed links representing probabilistic dependencies. Aiming at near-real-time applications based on exact inference, the random variables are assumed to be discrete or discretized, each having a finite set of outcomes or states. Considering the simple BN in Fig. 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, the random variable X_3 is said to be a *child* of random variables X_1 and X_2, whereas the latter are the *parents* of X_3. Similarly, X_4 is a child of X_1, whereas X_4 is the parent of X_5. Attached to each node of the BN is a *conditional probability table* (CPT), which provides the conditional probability mass function (PMF) of the random variable represented by the node, 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 Fig. 1, a *marginal probability table* is assigned.

The joint PMF of random variables in the BN, $p(x_1, \ldots, x_n) = \Pr(X_1 = x_1 \cap \cdots \cap X_n = x_n)$, is given as the product of conditional distributions, i.e.,

$$p(x_1, x_2, \ldots, x_n) = \prod_{i=1}^{n} p[x_i|\mathrm{Pa}(X_i)] \tag{1}$$

where $\mathrm{Pa}(X_i)$ = set of parents of node X_i; $p[x_i|\mathrm{Pa}(X_i)]$ = CPT of X_i; and n = number of random variables (nodes) in the BN. Thus, for the BN in Fig. 1, the joint PMF is

$$p(x_1, x_2, x_3, x_4, x_5) = p(x_5|x_4)p(x_4|x_1)p(x_3|x_1, x_2)p(x_1)p(x_2) \tag{2}$$

BNs are useful for answering probabilistic queries when one or more variables are observed. As an example, suppose for the BN in Fig. 1 the observations $X_3 = x_3$ and $X_4 = x_4$ have been made and the conditional distribution $p(x_2|x_3, x_4)$ is of interest. This posterior distribution is computed by first marginalizing the

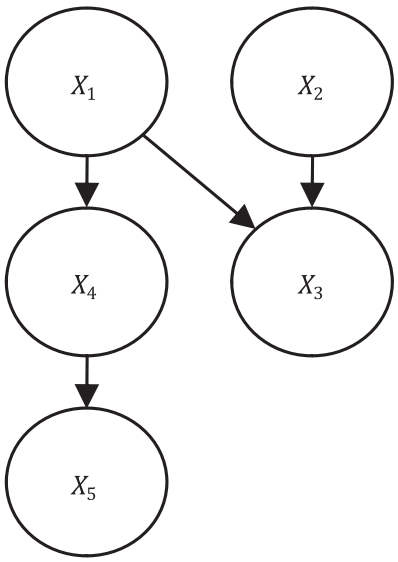

Fig. 1. Simple BN

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

joint distribution in Eq. (2) to obtain the joint distributions of the subsets of the variables

$$p(x_2, x_3, x_4) = \sum_{x_1, x_5} p(x_1, \ldots, x_5) \tag{3}$$

$$p(x_3, x_4) = \sum_{x_1, x_2, x_5} p(x_1, \ldots, x_5) \tag{4}$$

The desired conditional distribution then is $p(x_2|x_3, x_4) = p(x_2, x_3, x_4)/p(x_3, x_4)$. Although it is possible to obtain updated distributions by using this method, this is not a computationally efficient approach for nontrivial BNs. The efficiency of the BN stems from the decomposition of the joint distribution into local conditional distributions, as exemplified in Eq. (2). When summing over the joint distribution in Eqs. (3) and (4), no use is made of the decomposition. By writing the joint distribution in the product form of Eq. (1), it is possible to rearrange the summation and product operations because of their distributive and commutative properties. As an example, Eq. (3) is written as

$$\begin{aligned} p(x_2, x_3, x_4) &= \sum_{x_1} \sum_{x_5} p(x_5|x_4) p(x_4|x_1) p(x_3|x_1, x_2) p(x_1) p(x_2) \\ &= p(x_2) \sum_{x_1} p(x_4|x_1) p(x_3|x_1, x_2) p(x_1) \sum_{x_5} p(x_5|x_4) \end{aligned} \tag{5}$$

The summation operations can be interpreted as node eliminations. Because calculations are performed from right to left, Eq. (5) corresponds to an elimination of X_5 followed by the elimination of X_1. Solving the second line of Eq. (5) is more efficient than solving the first, because the summations are performed in smaller domains. The summation over X_5 is in the domain of X_4 and X_5 only (and can actually be omitted, because it results in 1). The summation over X_1 is in the domain of X_1, X_2, X_3 and X_4, for which it is required to establish a table, called a *potential*, whose number of entries is equal to the product of the number of states of these variables. The potentials, and consequently the efficiency of the inference algorithm, depend on the ordering of node eliminations. There exist computationally optimal elimination sequences, the domains of which are called *cliques*. The total size of the potentials associated with these cliques is a good measure of the computational demand for performing inference in the BN.

Several efficient algorithms for exact and approximate probabilistic inference in BNs have been developed (e.g., Lauritzen and Spiegelhalter 1988; Dechter 1996; Yuan and Druzdzel 2003, 2006; Madsen 2008; Langseth et al. 2009; Straub and Der Kiureghian 2010b). Although all exact inference algorithms aim at finding the optimal ordering of node eliminations, they follow different strategies for doing so. In particular, some algorithms optimize the elimination for a specific inference task, whereas others, such as the junction tree algorithm (Jensen and Nielsen 2007), optimize computations for general inference. With the latter, parts of the computations are reused, which is efficient when inference is performed repeatedly and at multiple nodes. Several of these algorithms are implemented in available software [e.g., Decision Systems Laboratory (DSL) 2007; Hugin Expert A/S 2008]. Because the focus of this paper is in near-real-time inference under an evolving state of information, the junction tree algorithm is preferred and is used for the reported application.

An ID is a BN to which *decision* nodes (shown as rectangles) and *utility* nodes (shown as diamonds) are added. Nodes representing random variables (shown as circles) are now called *chance* nodes to differentiate their probabilistic nature from the other nodes in the ID. The states of a decision node are the decision alternatives. In general, two types of decision alternatives exist: (1) action alternatives, e.g., shut down a system component, and (2) 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 at the time of making the decision, whereas an incoming link from another decision node indicates that the decision is made with knowledge of the selected alternative of that 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 alternatives and chance outcomes. 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.

To illustrate this description, consider the simple IDs in Fig. 2 concerning the decision of whether 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 S_i, indicates the state of the component (e.g., whether the component is intact) given the demand. In the ID in Fig. 2(a), the decision is made with no previous information (there is no incoming link into the decision

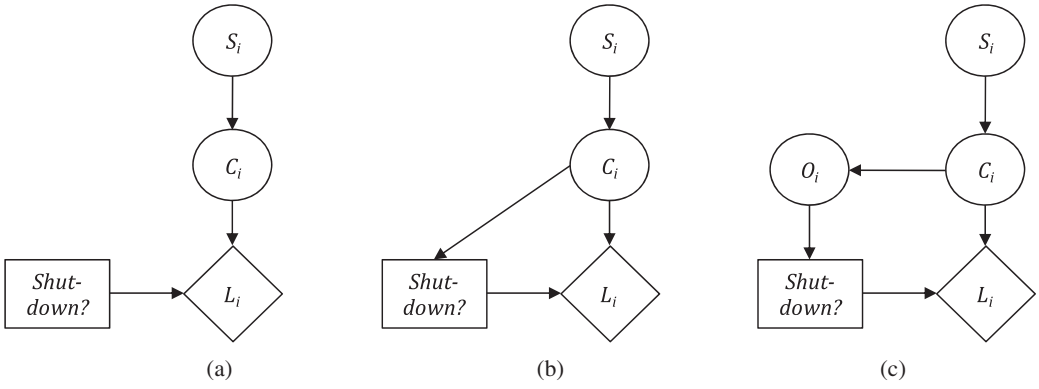

Fig. 2. IDs modeling shutdown decision for a system component: (a) no information; (b) perfect information; (c) imperfect information about state of component i

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

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 the decision is made to shut it down, there will be a loss associated with interruption of service, whereas if the component is damaged and the decision is made to keep it open, there will be a loss associated with the liability of making a potentially unsafe decision. The ID in Fig. 2(b) includes a link from node C_i into the decision node. In this case, the decision is made with perfect information about the state of the component. However, such cases of perfect information are not realistic. The ID in Fig. 2(c) describes a more realistic situation: An imperfect observation is made of the component state, which is represented by chance node O_i. This may represent, for example, the result of a visual inspection of the component that is providing an imperfect assessment of its state. This imperfect information is available to the decision maker. The observation does not affect the utility node. The CPT of node O_i is commonly known as the *test likelihood matrix*.

The IDs in Fig. 2 include only one decision node with action-type alternatives, i.e., whether to shut down the component. A more comprehensive example with multiple decision nodes is shown in Fig. 3, where there is a precedence of decision nodes. First, a decision is made about whether 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 because the outcome of the observation will depend on the selected method of inspection. Furthermore, the decision about whether to shut down the component is now made with knowledge of the type of inspection conducted and the result of the inspection. Later in this paper this model is incorporated into an ID for the infrastructure system as a whole. There, further aspects of IDs that require consideration when the order of decisions for system components becomes important are discussed. More in-depth treatment of IDs can be found in the text by Jensen and Nielsen (2007).

BN Model of Seismic Hazard

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). In this paper, one conceptual model is presented for ground shaking without providing the mathematical details, which are well established and can be found in the previous reference. To enhance the clarity of the presentation, BN objects are used, which are represented as rectangular nodes with rounded corners. These are BNs that operate in the background with specified input and output nodes.

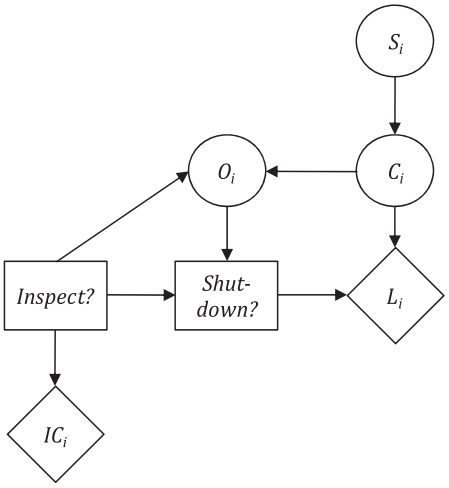

Fig. 3. ID modeling an inspection and shutdown decision of a system component

To facilitate the use of exact inference algorithms, all continuous random variables must be discretized. Methods for such discretization are described, e.g., in Kotsiantis and Kanellopoulos (2006), Neill et al. (2007), and Straub and Der Kiureghian (2010b). In this paper, the Monte Carlo simulation is used in conjunction with the method described in the latter reference to construct the CPTs. After determining the marginal outcome domain of each continuous variable, the domain is divided into a set of intervals, each representing a *discrete* state of the variable. To compute the CPT for a particular node, for each combination of the discrete states of the parent nodes, samples are generated within the corresponding intervals of the parent variables, and the corresponding realizations of the child node are computed according to the applicable mathematical relation, which can be either deterministic or probabilistic. If the parent is a root node, the marginal distribution is used to generate the sample. If the parent is not a root node, so that its marginal distribution is unknown, a uniform distribution within the interval is employed, or an exponential distribution if the interval extends to $+\infty$ or $-\infty$. Having computed the corresponding realizations of the child node, a normalized frequency diagram with the specified intervals of the node provides the column of the CPT for the particular selection of the parent node states. This process is repeated for all combinations of the parent node states. The error in this approximation clearly diminishes as the sizes of the intervals are reduced; however, this also increases the number of states of each node. Therefore, a balance between accuracy and computational efficiency must be sought. For the present work, a rigorous analysis of the discretization error has not been conducted; however, through a series of numerical experiments in the example application, the authors have made sure that the error in the representation is commensurate with the level of accuracy sought in seismic demand analysis.

The intensity of ground shaking at the site of a component i during an earthquake, i.e., the seismic demand on the component, is described by a predictive model (commonly referred to as a ground motion prediction equation or attenuation law and is obtained by regression analysis of observed data) of the form

$$\ln S_i = f(M, R_i, \mathbf{V}_i) + \epsilon_M + \epsilon_{R_i} \tag{6}$$

where S_i = measure of the intensity of ground shaking, such as spectral acceleration at the resonant frequency of the component; M = magnitude of the earthquake; R_i = closest distance to the earthquake source; $\mathbf{V}_i$ = additional site-specific variables that may influence the ground motion intensity at the site (e.g., properties of the soil at the site); the function $f(M, R_i, \mathbf{V}_i)$ = logarithmic mean of the intensity; ϵ_M and ϵ_{R_i} = model error terms. ϵ_M represents the inter-event error term and describes the variability of the model error from event to event, whereas 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 (e.g., Abrahamson et al. 2008). These are both zero-mean normal random variables with variances that may depend on M. In any given earthquake, ϵ_M is the same for all locations, whereas ϵ_{R_i} for different i vary from location to location and are spatially correlated. Eq. (6) is one example of the mathematical relation between a child node, S_i, and parent nodes, M, R_i, $\mathbf{V}_i$, ϵ_M and ϵ_{R_i}, as alluded to in the preceding paragraph.

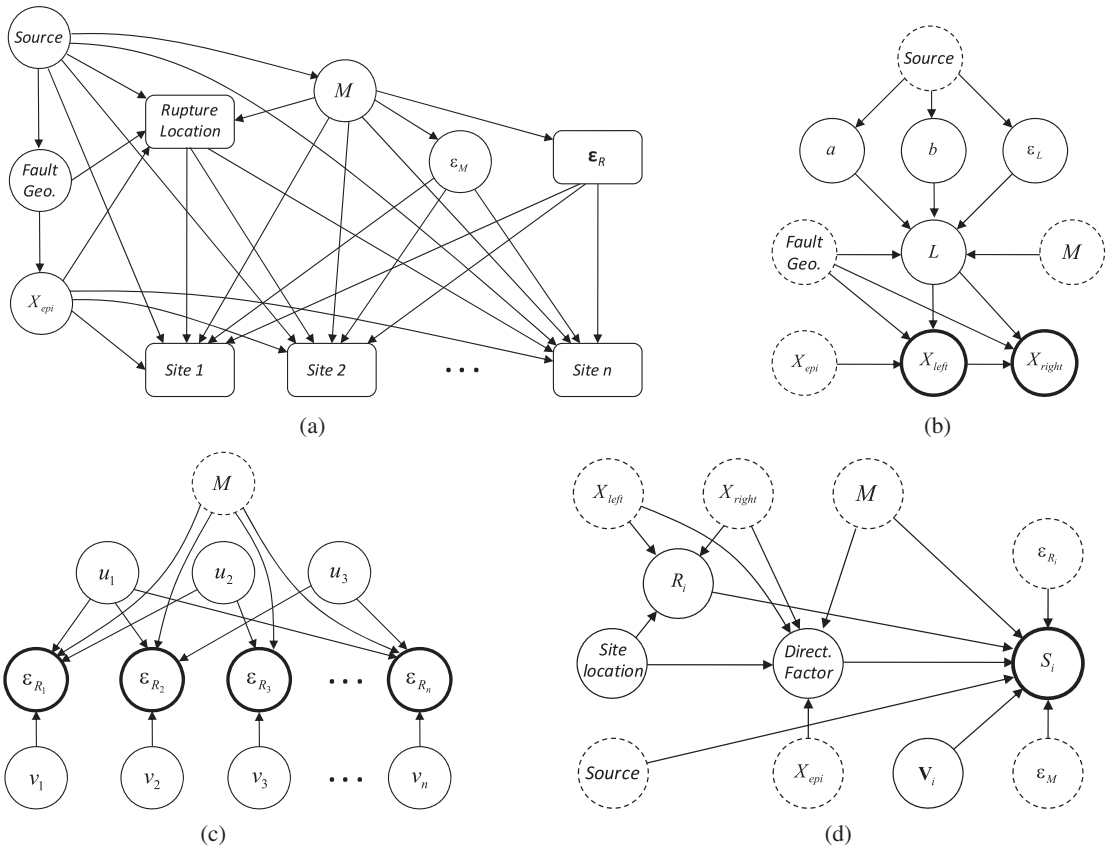

Fig. 4. BN model of seismic hazard: (a) global model; (b) rupture location object; (c) ϵ_R object; (d) site object

Fig. 4(a) 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, although more refined models can be developed. One child of this node is the *Fault Geometry* (labeled *Fault Geo.*), which describes the geometric coordinates of each source fault. Another child is node M, which describes the magnitude of the earthquake. The distribution of M depends on the type (e.g., strike-slip, dip-slip) and size of the fault, and thus the reason for the dependence on the Source node. It is assumed that the earthquake manifests as a rupture along the fault, originating at a random point X_{epi} on the fault, 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*. Fig. 4(b) shows this object. Inputs into this object are the four nodes just noted; 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), all three possibly depending on the source type. Usually, L is no greater than half the total fault length, hence the dependence of node L on the Fault Geometry, as depicted in Fig. 4(b). In the BN object, nodes a, b, and ϵ_L together with the Fault Geometry and M define the rupture length node L. The rupture is assumed to occur 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 is characterized by specifying the coordinates of the end points of the rupture, X_{left} and X_{right}, as children of the Fault Geometry nodes, rupture length L, and epicenter location X_{epi}. The output nodes from this object are X_{left} and X_{right}, which are highlighted by the bold borders of the corresponding nodes. These are used as inputs to every site object, as described subsequently. Knowing the coordinates of the end points of the rupture, one can determine the shortest distance R_i from each site to the fault rupture.

Now consider the BN object $\boldsymbol{\epsilon}_R$ of Fig. 4(a), which represents the intra-event error terms ϵ_{R_i}, $i = 1, \ldots, n$, for the considered component sites. Because 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 indicates that at least one of these nodes will have n parents, of which one is the magnitude node M and the remaining are the $n-1$ other components of $\boldsymbol{\epsilon}_R$. 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), this problem was addressed by expressing the vector of normal random variables $\boldsymbol{\epsilon}_R = [\epsilon_{R_1} \epsilon_{R_2} \ldots \epsilon_{R_n}]^T$ in the form $\boldsymbol{\epsilon}_R = \mathbf{T}\mathbf{u}$, where $\mathbf{u}$ is an $n \times 1$ vector of independent standard normal random variables and $\mathbf{T}$ is an $n \times n$ coefficient matrix obtained by eigenvalue analysis, Cholesky decomposition, or any other decomposition method that diagonalizes the covariance matrix of $\boldsymbol{\epsilon}_R$. One can show that in a BN model of the random vectors $\boldsymbol{\epsilon}_R$ and $\mathbf{u}$, nonzero elements of $\mathbf{T}$ correspond to links between the corresponding elements of the two vectors. To develop an efficient but approximate BN representation, the small-magnitude elements of $\mathbf{T}$ with zero are replaced incrementally, thus eliminating the corresponding links in the BN model,

while adjusting other non-zero elements of the matrix to minimize the error in the representation of the correlation matrix of $\boldsymbol{\epsilon}_R$. Zeroing all elements of column j of $\mathbf{T}$ is equivalent to eliminating the node corresponding to element u_j of $\mathbf{u}$. The final approximation is expressed in the form

$$\boldsymbol{\epsilon}_R \cong \mathbf{D}\mathbf{v} + \hat{\mathbf{T}}\mathbf{u} \tag{7}$$

where $\mathbf{v}$ = another $n \times 1$ vector of independent standard normal variables; $\mathbf{D}$ = $n \times n$ diagonal matrix; and $\hat{\mathbf{T}}$ = perturbation of $\mathbf{T}$ with many of its elements being zero. The role of the product $\mathbf{D}\mathbf{v}$ in Eq. (7) is to restore the variances of $\boldsymbol{\epsilon}_R$ after the approximation. The result of this process is an efficient, albeit approximate, BN model that is exemplified in Fig. 4(c). A few elements of $\mathbf{u}$ with sparse connections describe the correlations between ϵ_{R_i} nodes, whereas the v_i nodes restore the variances of individual ϵ_{R_i}. As demonstrated in Bensi et al. (2011b), this representation is far more efficient than the fully connected representation of the original variables $\boldsymbol{\epsilon}_R$. The input into this object is magnitude node M because the variances of ϵ_{R_i} may depend on M, whereas ϵ_{R_i} are the output nodes. $\mathbf{u}$ and $\mathbf{v}$ are latent variables and their nodes remain internal to the BN object shown in Fig. 4(a).

Next consider the ith site object in Fig. 4(d), which shows the BN behind the object *Site i* in Fig. 4(a). Input nodes into this object are the magnitude M, 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, because 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 readily 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 this model can be made. As an example, the effect of directivity of the fault rupture for a near-fault site is considered. As is well known, under certain conditions, when the fault rupture propagates toward 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. (6) (Somerville et al. 1997; Abrahamson 2000; Bray and Rodriguez-Marek 2004). These factors generally 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 toward the site. The BN model in Fig. 4(d) includes this factor. As shown, the object now must 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 toward the site. The *Directivity Factor* node is now an additional parent of node S_i.

BN Model of System Components

Suppose a 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, the conditional probabilities associated with these states of the component are defined 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 (i.e., the intensity of ground motion). Fig. 5(a) 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 obtained as the difference between the two bounding fragility curves. Because 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. Fig. 5(b) shows the global BN model of the system with the seismic demand and system performance models shown as objects. Of course, discretization of the seismic demand and the corresponding fragility functions is necessary to yield a BN with discrete nodes.

Without providing additional details, this BN model can be extended or enhanced to account for (1) dependence between the capacities of components, (2) the effect of parameter or model uncertainties, (3) components with multiple failure modes, and (4) distributed components (e.g., pipelines), where failure may occur randomly along the component. These issues are addressed in Bensi et al. (2011a).

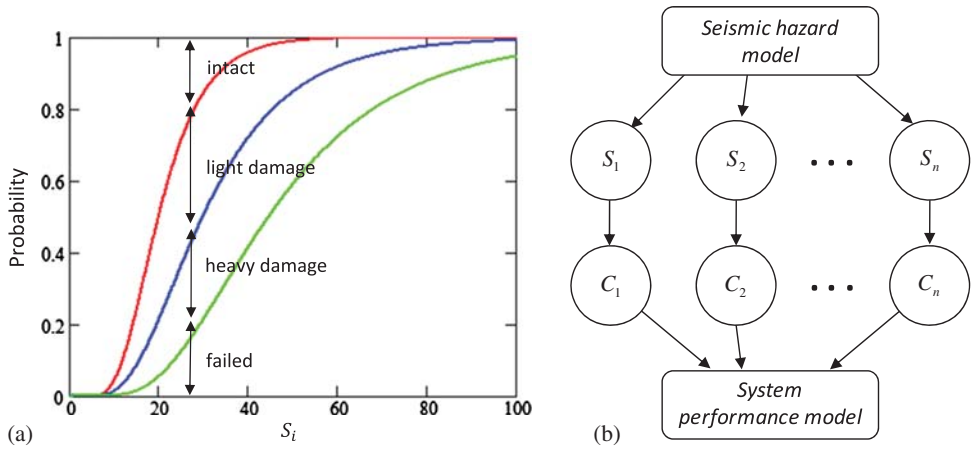

Fig. 5. Component performance model: (a) fragility functions for a four-state component; (b) BN model

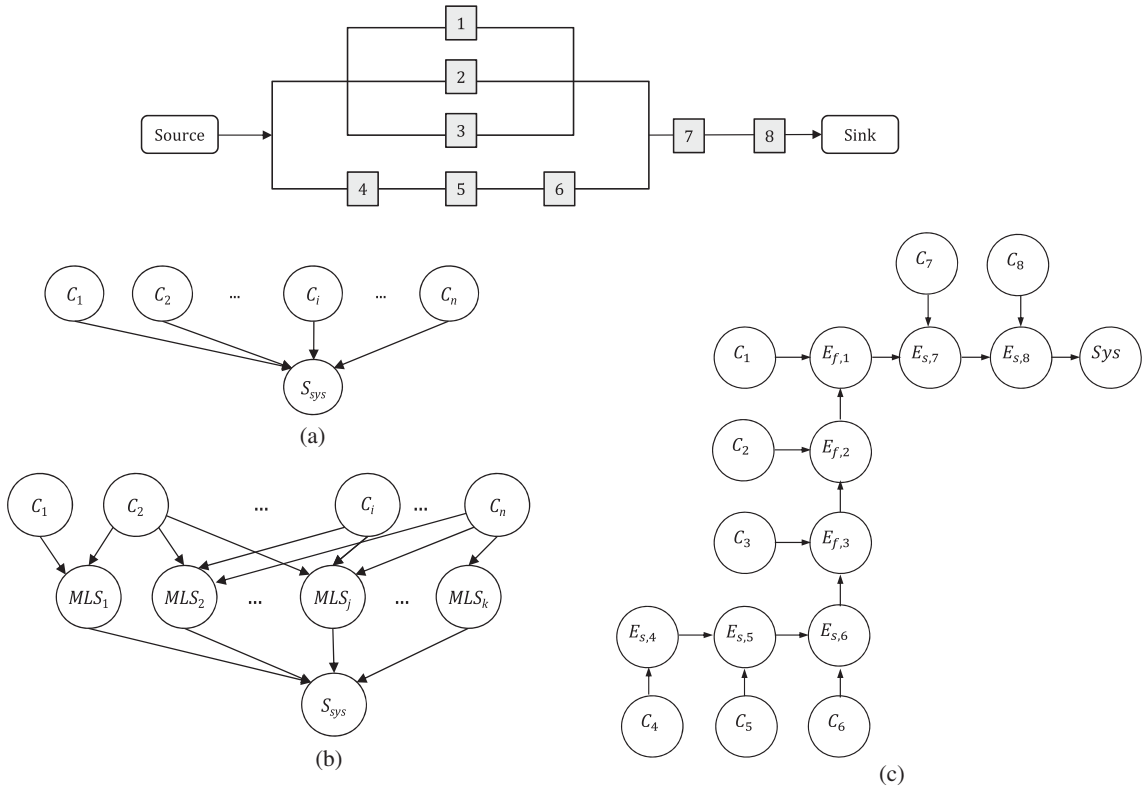

Fig. 6. System performance model: (a) naive formulation; (b) formulation in terms of minimum link sets; (c) efficient formulation for the simple system shown above in terms of survival and failure path events

BN Model of Infrastructure 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 whether travel between any set of origination and destination points is possible. Using this relationship, a simple BN model describes the system node as a child of all of the component nodes [Fig. 6(a)]. This formulation, however, is inefficient because, for a large system, the system node will have many parents. For two-state systems with two-state components, a more efficient formulation would be to introduce intermediate nodes representing minimum-link sets (MLSs) or minimum-cut sets (MCSs) of the system. An MLS (MCS) is a minimum set of components whose joint survival (failure) constitutes the survival (failure) of the system. A sample BN formulation in terms of MLSs is shown in Fig. 6(b). However, with a growing number of components, the sizes and numbers of MLSs or MCSs also grow, resulting in a formulation that is also inefficient.

In Bensi et al. (2013), a topology optimization scheme was developed for constructing efficient BN models of two- and multistate systems. Using MLSs or MCSs, the BN is modeled by a chain-like sequence of *survival and failure path events,* in which the state of each event depends on the state of the preceding event in the chain in addition to the state of an associated component. Fig. 6(c) shows an example of this formulation for a simple system with $E_{s,i}$ and $E_{f,i}$ denoting the survival and failure path events. This representation yields a chain-like BN model, in which nodes have few parents. See Bensi et al. (2013) for more details and for the extension of this method to multistate systems.

The computational resources required to support the proposed probabilistic framework are considerable for nontrivial systems. The computational challenge arises from the memory demands associated with performing inference in the BN by use of the junction-tree algorithm. Specifically, the growth in computational demand is associated with the size of cliques that are formed when performing inference. In this application, the topology of the BN leads to large clique sizes. This is because a construct involving common parents between component nodes is necessary to capture the spatial correlation in the ground motion (i.e., the random field aspect). Because these component nodes jointly determine the system performance, cliques become large. The presented efficient representations of the random field and the infrastructure system mitigate this problem only partly. For application of the framework to large infrastructure systems, new inference algorithms may need to be developed.

ID of Infrastructure System

An influence diagram (ID) at the system level is needed to support post-earthquake decision making for the infrastructure, specifically, making decisions on whether to inspect each component, and whether to perform mitigation actions, e.g., reduce or shut down the operation of each component to avoid further losses. Because usually there are limited resources for inspection of components, it is also necessary to prioritize the order of component inspections.

Thus, following the principles of expected utility decision analysis, the aim of this study is to solve the following optimization problem: Maximize the expected utility (or minimize the expected cost) of operating the infrastructure in the immediate aftermath of the earthquake hazard, in which the optimization parameters are the choice of the order and the type of component inspections

(including the option of not inspecting a component) and of the mitigation actions to perform. Because the inspections provide additional information, the optimization problem is adaptive: After each inspection outcome, the optimal decisions (on mitigation actions and/or on future inspections) can change. Such adaptive decision problems have exponential growth in computational demand with the number of decision alternatives and inspection outcomes (DeGroot 1969). If the number of components to inspect is n, the number of inspection types is m, each with k different possible outcomes, and the number of potential mitigation action alternatives is l, an exhaustive computation would require considering $n! \times (m \times k \times l)^n$ possible combinations of decision alternatives and inspection outcomes. For each of these combinations, the conditional probability of system performance must be computed. It should be evident that for numbers of components n larger than just a few, the computational demand of such an exact solution becomes intractable, particularly for consideration for a near-real-time decision support system. In the following, the authors first consider how such problems can be modeled by IDs. Thereafter, a computationally efficient heuristic is proposed for solving the optimization problem in an approximate manner.

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 of 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 of the preceding observations, outcomes, and decisions. For the present 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). For 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) proposed UIDs. They are relevant to this problem. However, as described previously, UIDs experience exponential growth in complexity with the number of decision nodes when using exact solution algorithms. Hence, they are not practical for application to large systems. In this study, LIMIDs are used; however, a value-of-information heuristic is invoked, and is applied *outside* the construct of the LIMID to decide on the temporal ordering of the inspections. Efficient algorithms and software for analysis of LIMIDs are available (e.g., DSL 2007; Hugin Expert A/S 2008).

When solving a LIMID, one first determines a *policy* for each decision node that maximizes the expected utility for any given configuration of the states of its parent nodes, i.e., it specifies a preferred action for every combination of the states of the parent nodes. A 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* (Lauritzen and Nilsson 2001). 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. Let q_i be the current strategy. Begin a cycle by updating decision D_1 to yield a new strategy q_{i+1}. This is done by updating the local maximum-utility policy for decision D_1, while keeping all other elements of q_i fixed. The algorithm next computes a locally optimal policy for decision D_2. When all policies for all decisions have been updated, the cycle is complete. The algorithm converges when the expected utilities associated with successive cycles remain the same (Lauritzen and Nilsson 2001).

Fig. 7 shows the system-level LIMID for decisions concerning component inspection and operational states. For the sake of clarity, the models for the seismic hazard and system performance are shown as BN objects. The ID for each component is similar to that shown in Fig. 3, except that a component capacity node, Cap_i, has been added as a child of the component state and the *Shut-down?* decision node. This latter node includes decision alternatives concerning the operational state of the component. For example, the alternatives might be to operate at full capacity, operate at reduced level, or shut down. The system BN object is shown as a child of the component capacities. Attached to the node describing the state of the system (the output from the system BN object) is a utility node RL_{sys}. This node contains the utility values (revenue lost) associated with each performance level of the system. These utilities are exclusive of the costs associated with individual components, which are separately accounted for in the local utility nodes representing liabilities (nodes labeled L_i) and inspection costs (IC_i) attached to each component i. For example, if a highway system is being used, the system utility node may describe the cost of not being able to travel between selected origin and destination points, whereas individual component utility nodes may describe the liability cost associated with making an unsafe decision (e.g., keeping a damaged bridge open) and the cost of inspecting the component. The LIMID does not include precedence links between the inspection decision nodes. This is important because prioritization of inspections is a desired outcome of the ID and an inspection order should not be assumed a priori. Hence, an approach to prioritize the temporal ordering of the component inspections needs to be developed.

Value of Information Heuristic

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, e.g., a decision to shut down a component or keep it in operation (Raiffa and Schlaifer 1961; Straub and Der Kiureghian 2010a). 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 before 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.

In this paper, the VoI concept is used to determine a temporal ordering of the component inspections using the output of the ID. First consider the component-level decision problem depicted in Fig. 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 Fig. 3 is set equal to the no-inspection 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 Fig. 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

$$\text{VoI} = \text{EU}|\text{I} - \text{EU}|\text{NI} \tag{8}$$

An inspection is worth performing if its cost is less than its VoI. Among a set of inspection alternatives, the optimal alternative is

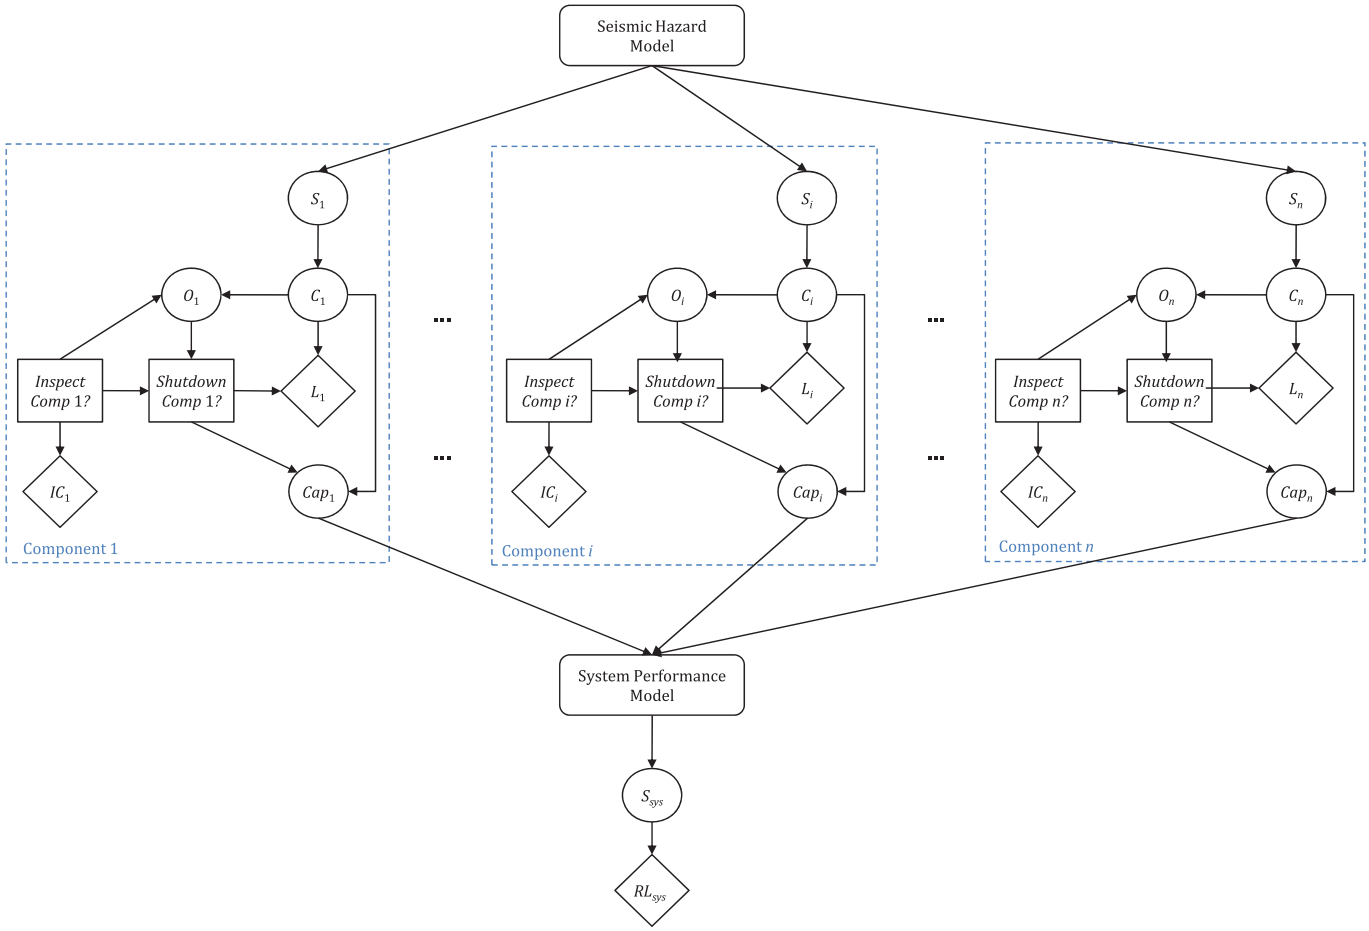

Fig. 7. LIMID model of seismic hazard

the one that has the largest positive difference between the VoI and the cost of inspection, i.e., the one with the largest net benefit of inspection (NBI).

To help prioritize post-earthquake inspections of components at the system level, a VoI heuristic is proposed. 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 Fig. 7 is analyzed to determine approximately the optimal inspection option for all components. The solution may include cases in which the optimal decision for a component is not to inspect. This occurs when the cost of inspection outweighs the potential benefits of gaining additional information, such as would occur when the available information about the hazard suggests a high likelihood for one particular state of the component, or when the cost of inspection is higher than the penalty or liability cost associated with making an unsafe decision. The corresponding total system-level expected utility is computed. Next, for component i requiring inspection, the inspection decision node is set to the *no inspection* state, whereas the remaining components are at their optimal inspection decisions (given no inspection of component i) with their corresponding inspection costs. The expected utility for the system is again computed. The system-level NBI for component i is the difference between the expected system utility with and without inspection of component i. The procedure is repeated until the system-level NBI for each component is determined. These NBI values are used to rank the component inspections and their recommended temporal order follows this ranking.

Using the LIMID methodology, the inspection prioritization order evolves as new information becomes available. As described previously, 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 also entered into the BN-LIMID. The information is propagated through the BN-LIMID to provide an up-to-date 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 the information available up to that time.

The formulation of the decision problem by a LIMID may arrive at a suboptimal solution. As discussed previously, an exact solution is computationally intractable for systems with more than just a few components because of the computational demands associated with probabilistic inference and the large number of possible combinations of decision alternatives when considering all possible orders of inspections. With the proposed approach, the number of probabilistic evaluations before each inspection is only a linear function of the number of components in the system. Hence, the approach is

applicable to large systems, while producing a solution that is at least locally optimal.

Application to a Hypothetical Model of the Proposed California High-Speed Rail System

As an illustration of the proposed BN-ID framework, post-earthquake risk assessment and decision making for a hypothetical model of a segment of the proposed California high-speed rail (HSR) system is considered. This system is still under development. Hence, the model is a highly idealized and simplified one; the specifics of the system components are imagined rather than being based on existing structures. In addition, simplified models of the seismic environment (e.g., only a subset of the existing active faults, simple fault geometries) are employed. For this reason, the application and 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 in this paper; details of other, more standard aspects can be found in Bensi et al. (2011a). The Hugin software (Hugin Expert A/S 2008) is used for all BN and ID computations in this application.

The northern segment of the proposed HSR system is considered, 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. Fig. 8 shows the configuration of the system and the seismic environment. Simplified models that are generally consistent with the current state of practice are used to describe the seismic sources for the three faults, rupture length–magnitude relations, ground-motion prediction equations, and spatial correlation of ground-motion intensity. 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 19 components with the types and lengths described in Table 1 (columns 2–4). Ground-motion prediction points (GMPPs) along the path define the end points of these components. These are shown in Fig. 8 and their correspondence with the components is given in Table 1 (columns 5 and 6). Thus, for example, Component 1 is a tunnel with 4-km length and stretching from GMPP 1 to GMPP 2. Distributed components are modeled assuming that damage occurs along the component according to a nonhomogeneous Poisson process with the rate defined as a linear function of the ground motion at the GMPEs on either end of the component. Three states for each component are considered: undamaged, slightly damaged, and heavily damaged. Undamaged components are assumed to be fully operational. Components in the state of

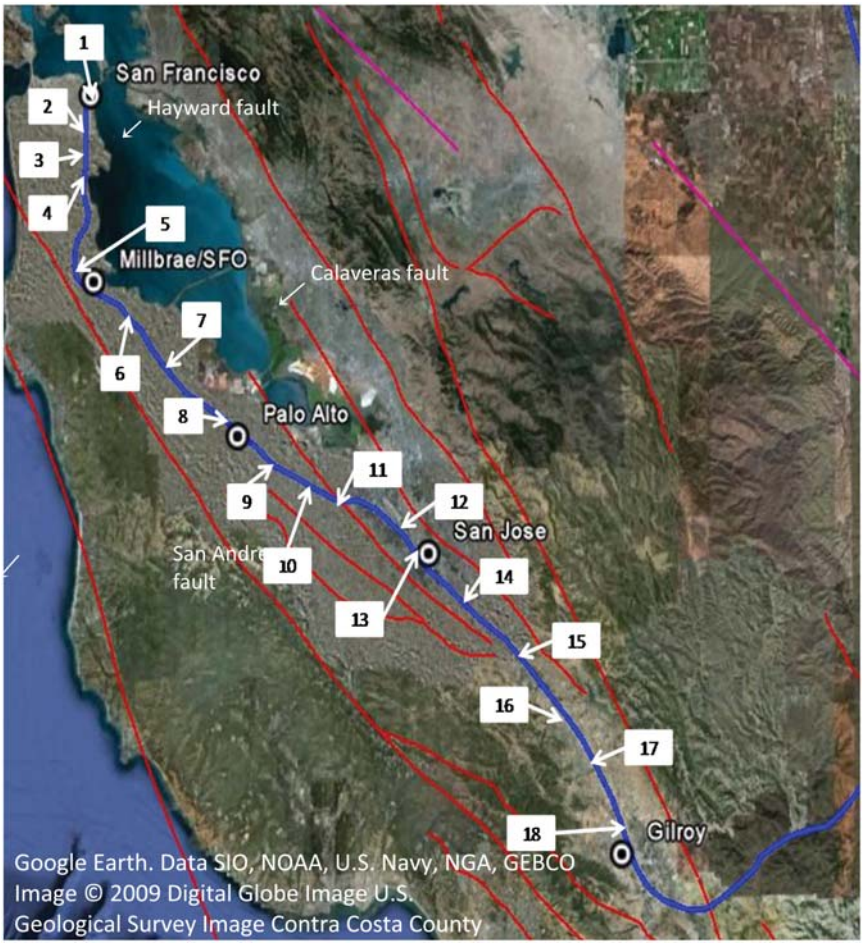

Fig. 8. Configuration of the system, seismic environment, and ground-motion prediction points [image © Google Earth, Data SIO, NOAA, U.S. Navy, NGA, GEBCO, image © 2009 Digital Globe Image U.S. Geological Survey Image Contra Costa County; alignment data from California High-Speed Rail Authority (2010); fault data from Bryant (2005)]

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

Table 1. Characteristics of System Components and Corresponding Inspection and Liability Costs

Link	Components	Type	Length (km)	GMPP (i)	GMPP ($i+1$)	Inspection cost visual	Inspection cost extensive	Liability cost[a] Normal SD	Liability cost[a] Normal HD	Liability cost[a] Reduced HD	Cost of reduced speed	Cost of closure
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
1	1	Tunnel	4.0	1	2	20	200	60	6,000	600	1,000	4,000
	2	Embankment	2.7	2	3	6.75	67.5	13.5	1,350	135		
	3	Tunnel	3.5	3	4	17.5	175	52.5	5,250	525		
	4	Embankment	8.8	4	5	22	220	44	4,400	440		
2	5	Embankment	8.7	5	6	21.8	118	43.5	4,350	435	1,000	3,000
	6	Embankment	8.4	6	7	21	110	42	4,200	420		
	7	Embankment	8.8	7	8	22	120	44	4,400	440		
3	8	Embankment	8.8	8	9	11	210	44	4,400	440	1,000	7,000
	9	Embankment	7.7	9	10	9.25	193	38.5	3,850	385		
	10	Embankment	6.9	10	11	7.25	173	35.4	3,540	354		
	11	Tunnel	6.0	11	12	30	300	90	9,000	900		
	12	Aerial	Point	12	12	75	750	25	2,500	250		
	13	Embankment	2.7	12	13	0.675	67.5	13.5	1,350	135		
	14	Aerial	Point	13	13	75	750	25	2,500	250		
4	15	Embankment	8.1	13	14	20.3	203	40.5	4,050	405	1,000	5,000
	16	Embankment	9.0	14	15	22.5	225	45	4,500	450		
	17	Embankment	8.9	15	16	22.3	223	44.5	4,450	445		
	18	Aerial	5.5	16	17	226	2,263	137.5	13,750	1,375		
	19	Embankment	9.5	17	18	23.8	238	47.5	4,750	475		

[a]SD = slight damage; HD = moderate or heavy damage.

slight damage are safe for reduced performance, e.g., slower train speed when traversing the component. For the state of heavy damage, the component should not be operated to assure safety. Fragility models for these component states are adapted from the literature. These fragility models are hypothetical and do not reflect the capacities of the future HSR system components.

The system performance is defined in this study in terms of the ability to travel from San Francisco to Gilroy, and from there to Southern California. However, it is assumed that if a segment of the system is not passable, passengers can use an alternative transportation means, such as a *bus bridge*, between selected cities 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 links along the path are considered for which alternative transportation means could be provided, if necessary. These are Link 1 between San Francisco and Millbrae, Link 2 between Millbrae and Palo Alto, Link 3 between Palo Alto and San Jose, and Link 4 between San Jose and Gilroy (Fig. 8 and column 1 of Table 1). Travel at normal speed along a link is possible if none of its components is damaged; travel should be at reduced speed if any component is slightly damaged; and ability to travel is lost if any component is heavily damaged. Costs at the system level occur when the train speed is reduced over a link, or when a link is closed and an alternative means of transportation is provided. In addition, liability costs incur if an unsafe decision is made, e.g., a decision is made to maintain the normal speed when a component is slightly damaged, or to keep a damaged component open. These costs are defined later in this section.

Fig. 9 shows the BN objects for the random field model of the intra-event error terms and for the system performance with the four travel segments indicated. The random field model of Fig. 9(a) is the result of applying the modeling approach described in Bensi et al. (2011b) to the locations of the HSR system components. The system performance BN of Fig. 9(b) follows the efficient modeling procedure presented in Bensi et al. (2013).

A main objective of the present study is to demonstrate how the BN-ID framework allows updating of the hazard model and the system state and decision making at the link and component levels under an evolving state of information. For this purpose, four instances of information are considered that may be acquired following an event, which are referred to as *evidence cases* (ECs), as follows:

1. EC 1: Shortly after the earthquake event, the Berkeley Seismological Laboratory (http://seismo.berkeley.edu/) learns that the earthquake had a magnitude $M = 6.8$ and an epicenter located on the Hayward fault, 30 km from its north end.
2. EC 2: A little later, a recording at GMPP 2 indicates a spectral acceleration of $S_2 = 0.45$–0.50 g at 1-Hz frequency.
3. EC 3: A little later information is received that suggests 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 information is received that indicates that Component 17 has experienced heavy damage.

Each evidence case is *added* to the evidence from the previous case by entering the information into the BN. These evidence cases are designed to incrementally provide "bad" news. Hence, one would expect that component and system failure probabilities progressively grow with the evidence cases. Of course, this trend is not necessary, but it is helpful in this demonstration.

First, the influence of the evolving information on the predicted seismic demand at selected GMPPs is considered. 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. Fig. 10 shows the updated (discretized) distributions of the seismic demand (spectral acceleration at 1-Hz frequency) 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 0.1–0.15 g (when disregarding the tail probability for the intensity interval 0.6–∞). After EC 2, the probability mass sharply moves toward higher values so that the most likely value now is 0.25–0.30 g and the upper tail probability increases significantly. This is because the observation of a high intensity at a nearby site (GMPP 2) provides evidence for a higher intensity at GMPP 1 because of the spatial correlation among the intra-event error terms ϵ_{R_i}. EC 3 does not affect the distribution of the intensity

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

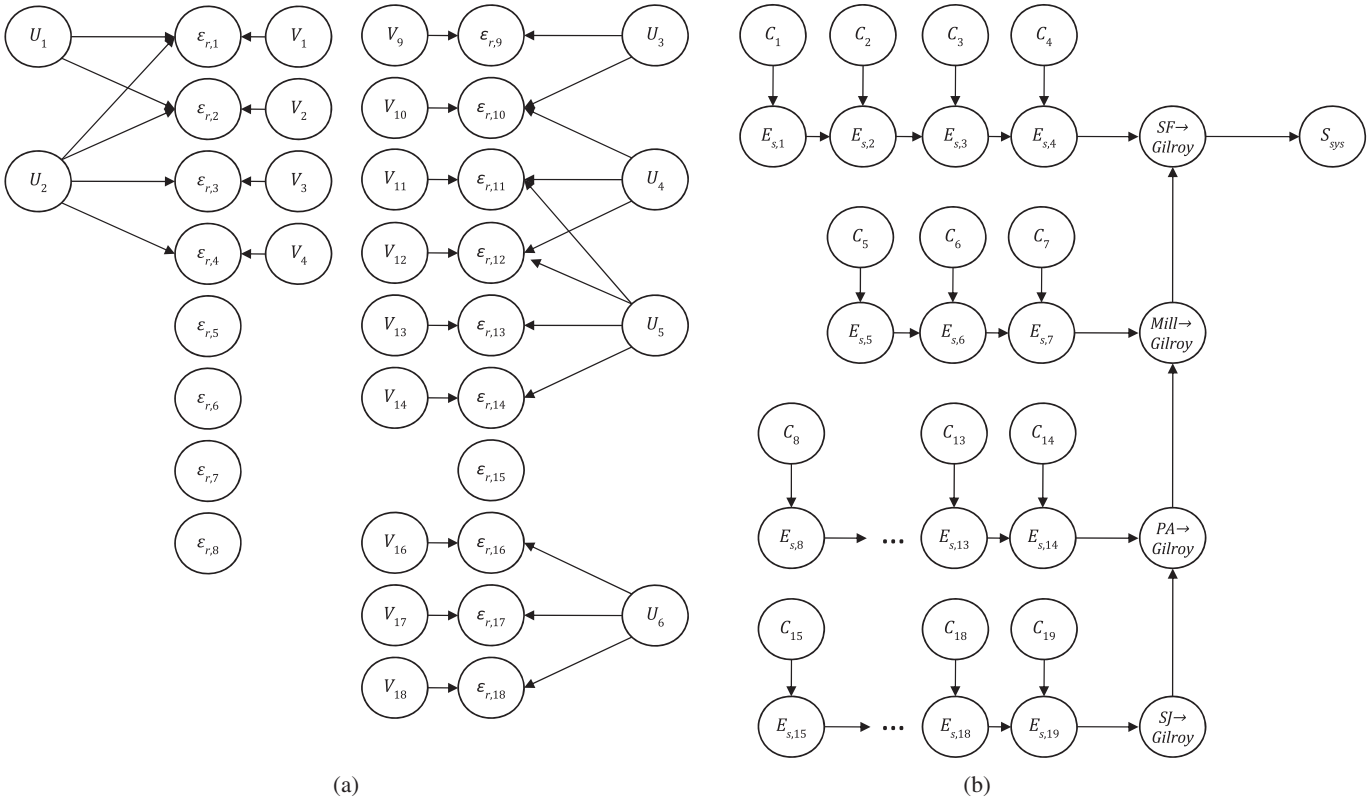

Fig. 9. BN objects of (a) random field model of the intra-event error terms associated with GMPEs; (b) system with four travel segments

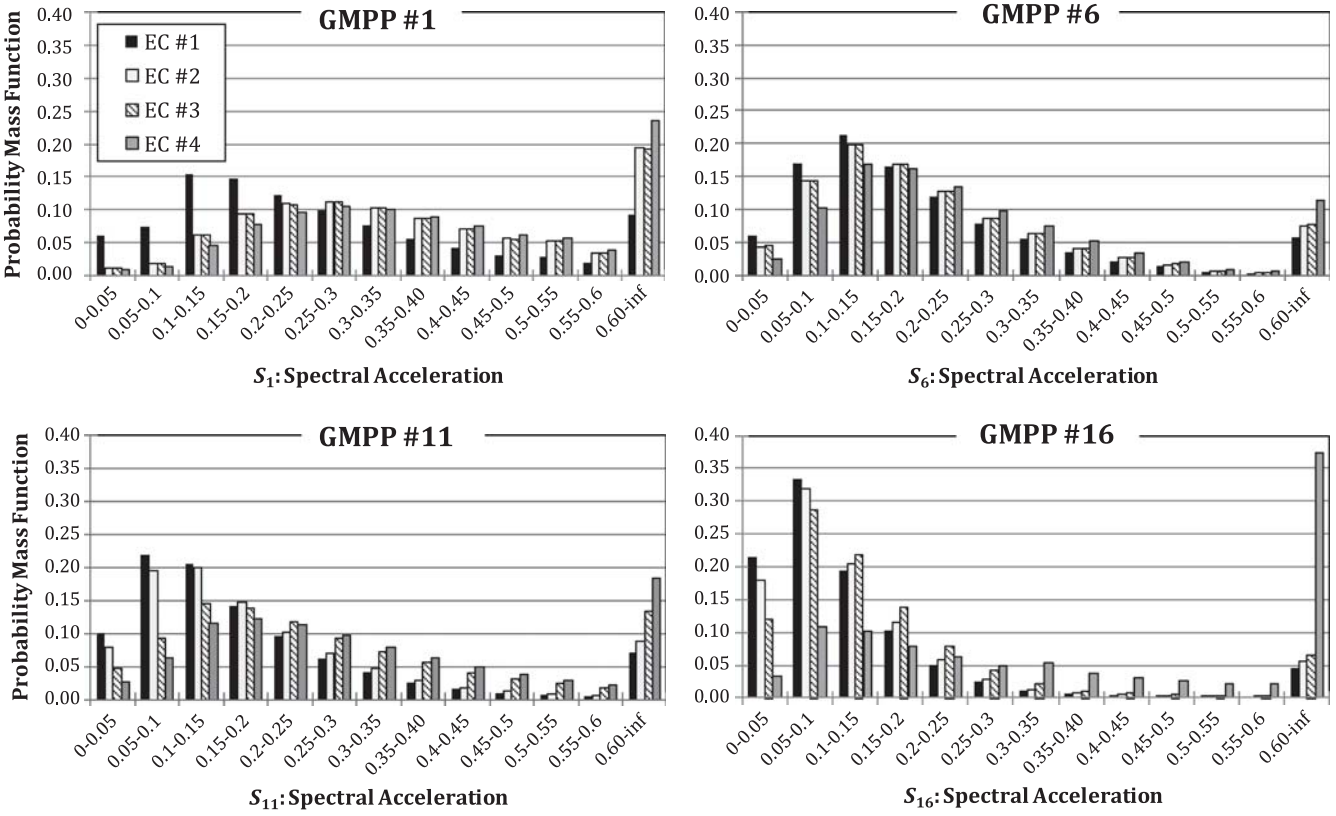

Fig. 10. Distributions of ground-motion intensities at GMPPs 1, 6, 11, and 16 for four evidence cases

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 this site relative to the fault and the epicenter excludes the possibility of a significant forward or backward directivity effect. For EC 4, further shifting of the probability mass toward higher values is shown by a small amount. This is because of the observation of heavy damage of Component 17. Because that component is located far from GMPP 1, the effect is not the result of the spatial correlation [in fact, the approximate random field model in Fig. 9(a) does not include a *path* in the BN between the GMPP 1 and GMPPs 15 and 16], but the result of the correlation arising from the inter-event error term ϵ_M, which is common to all GMPPs. Next consider GMPPs 6, 11, and 16. For EC 2, shifts of probability mass toward higher values is shown, 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 the result of the effect of the common ϵ_M, not because of the spatial correlation of ϵ_{R_i}. For EC 3, strong shifts of the probability mass toward higher values are observed for GMPPs 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, systematic shifts of the probability mass toward higher values at all sites are shown (because of the effect of the inter-event error term ϵ_M) and a particularly large shift for GMPP 16. This is because observation of heavy damage of Component 17 suggests high-intensity ground motions at GMPPs 15 and 16, which are located at its ends (Table 1).

Next, the probability of damage of each component under the evolving state of information is examined, as shown in Fig. 11. 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 from GMPP 2. Under EC 3, sharp increases of the damage probabilities for Components 8–19 are shown, which are located in the forward directivity sites. Under EC 4, sharp increases of the damage probabilities for all components are shown, especially those near Component 17. The probability of damage of Component 17 is of course 1.0 (outside the graph in Fig. 11).

Fig. 12 shows the probabilities of one or more components being in the states of slight or heavy damage on the path from each city to Gilroy for the evolving state of information. Naturally, this probability decreases if one starts from a city further south from San Francisco, because the path then includes fewer components. The probabilities increase with increasing "bad" news. In particular, in light of EC 4, the probability of interest is 1 for all cities because of the knowledge that Component 17 near Gilroy has experienced heavy damage.

After the earthquake, the decision needs to be made about whether to continue normal operation over each link, reduce the train speed while traversing it, close the link and provide alternative means of transportation, or delay this decision until the link is inspected. The optimal type and sequence of component inspections within each link also need to be determined. For this purpose, a separate BN-LIMID for each link is developed. As an example, Fig. 13 shows the graph for Link 1. (The decision model used in Fig. 13 differs slightly from that described in Fig. 7.) Because this is a series system, the decision to shut down any component or reduce the train speed over it is, by default, a decision to shut down/reduce speed over the entire link. As a result, a single shutdown/reduce-speed decision node at the link level is modeled. The shutdown/reduce-speed decision node for the link is a child of the link-inspection decision node, and nodes representing observations made of all of the components (using a chain-like structure) are consistent with the model for component-level shutdown

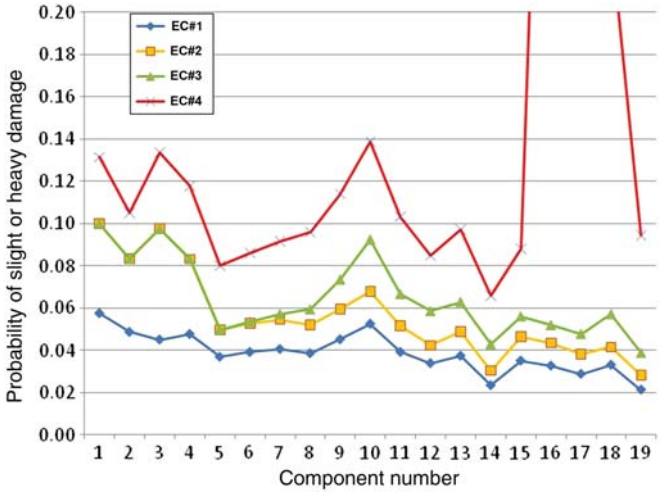

Fig. 11. Probabilities of component being slightly or heavily damaged for evolving evidence cases

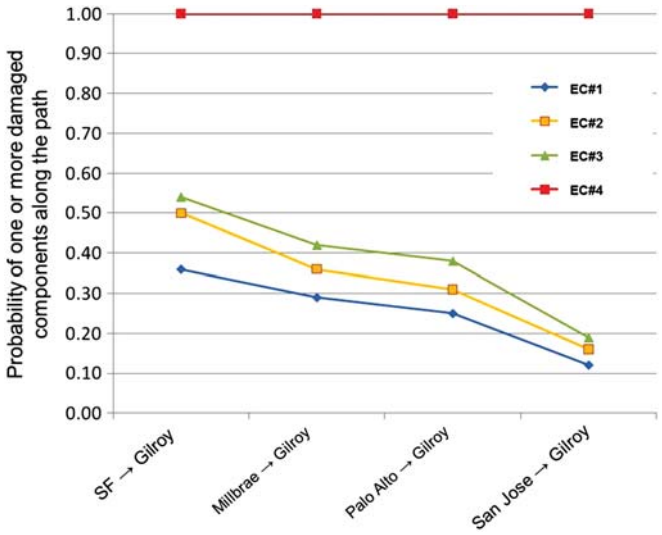

Fig. 12. Probabilities of one or more damaged components along the path from each city to Gilroy

decisions shown in Fig. 7. Furthermore, the link shutdown/reduce-speed decision node is a parent of each component utility node representing liability. Finally, a node has been added to represent the inspection decision at the link level. This node is a parent of the nodes representing component-level inspection decisions (to indicate decision precedence) and the component observation nodes.

Three inspection options are considered for each component: no inspection, visual inspection, and extensive inspection. Costs per component for the visual and extensive inspections are listed in columns 7 and 8 of Table 1. These are selected based approximately on the size and type of each component, and are expressed in an arbitrary cost unit. The no-inspection option has zero cost. The assumed likelihood matrices for the visual and extensive inspection options are listed in Table 2. The extensive inspection is assumed to yield perfect information. At the link level, two options are considered: Either inspect the link before making an action decision, in which case an optimal ordering of component inspections is

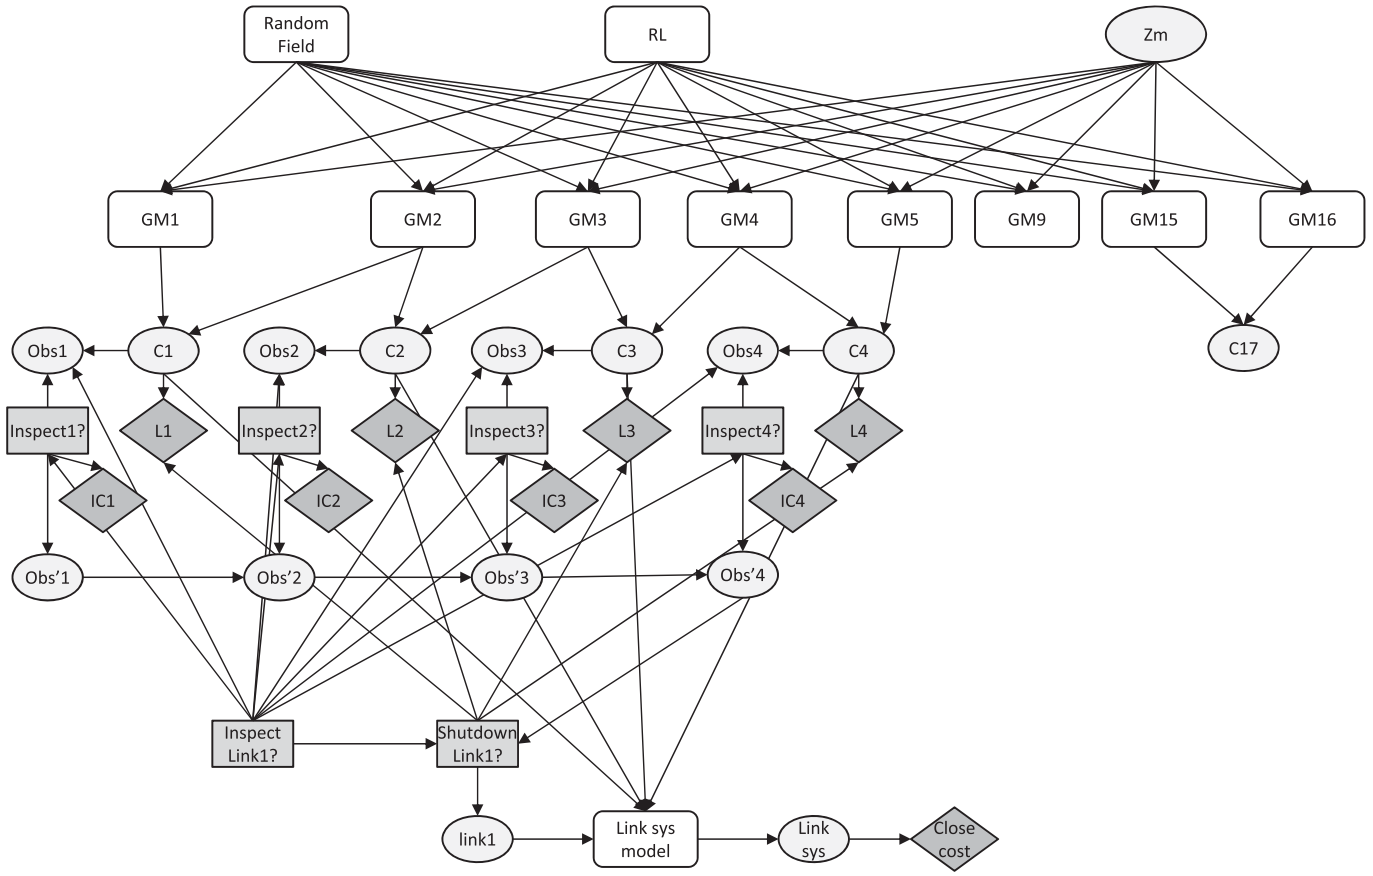

Fig. 13. LIMID for Link 1

Table 2. Assumed Likelihood Matrices for Visual and Extensive Inspections

Inspection outcome	Visual inspection			Extensive inspection		
	True state			True state		
	No damage	Slight damage	Moderate/heavy damage	No damage	Slight damage	Moderate/heavy damage
No damage	0.5	0.2	0.2	1	0	0
Slight damage	0.3	0.6	0.3	0	1	0
Moderate/heavy damage	0.2	0.2	0.5	0	0	1

determined, or do not inspect the link and take an action, e.g., leave it open with normal operation, leave it open but reduce the train speed over the link, or shut the link down. In the latter case, an alternative mode of transportation over the link will be provided. Assuming inspections will be conducted immediately, the cost of delay in operation as a result of inspection are not included.

To make decisions, it is necessary to specify utilities (or costs) for all pairs of decision alternatives and outcomes, both at the component and link levels. Columns 9–11in Table 1 list the assumed liability costs (in the same units as the inspection costs) for each component that is either slightly or heavily damaged and is left open with normal operation, or is heavily damaged and is operated at reduced speed. No liability costs are incurred for an undamaged component, or for a slightly damaged component that is operated with reduced speed. The cost of reducing the train speed over a link is listed in column 12, and the cost of closing a link is listed in column 13. The latter includes the cost of providing an alternative means of transportation in addition to lost revenue from passengers who select other options for their travel. These cost values are selected approximately based on the type and size of each component and the length of each link, but are otherwise arbitrary. Negatives of the cost values are interpreted as utility values.

The optimal decisions on inspections and operation of the links and components are determined in a two-level procedure. This procedure is invoked under each evidence case, so that the optimal decisions evolve as new evidence arises. First, a decision is made about which links to inspect and in what order. If no inspection is the optimal choice for a link, then a decision is made for the level of operation of the link (i.e., open, reduced speed, shutdown). For each link that is recommended for inspection, decisions are then made about the type of inspection of each component and their order.

The analysis at the link level for each evidence case is performed as follows: First the evidence is entered and propagated through the BN to update all probabilities. Next, the single-policy updating algorithm (Lauritzen and Nilsson 2001) is applied to arrive at the optimal policy for each link conditional on the evidence. This decision is binary on the inspection, plus an action decision if the

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

Table 3. Link-Level Decisions, NBI, and Priority Ranking in Case of Inspection

Link	EC 1	EC 2	EC 3	EC 4
1	Inspect, 455 (1)	Inspect, 351 (1)	Inspect, 351 (1)	Inspect, 237 (2)
2	Inspect, 148 (3)	Inspect, 306 (2)	Inspect, 319 (2)	Inspect, 596 (1)
3	Inspect, 23.3 (4)	Do not inspect, reduce speed	Do not inspect, reduce speed	Do not inspect, reduce speed
4	Inspect, 252 (2)	Inspect, 200 (3)	Inspect, 138 (3)	Do not inspect, close link

Note: Priority ranking is shown in parentheses.

Table 4. Component-Level Decisions on Type of Inspection, NBI, and Priority Ranking

Component	EC 1	EC 2	EC 3	EC 4
1	Visual, 247 (1)	Visual, 282 (1)	Visual, 282 (1)	Visual, 237 (1)
2	Visual, 202 (4)	Visual, 161 (4)	Visual, 161 (4)	Visual, 138 (4)
3	Visual, 213 (3)	Visual, 215 (2)	Visual, 215 (2)	Visual, 217 (2)
4	Visual, 217 (2)	Visual, 194 (3)	Visual, 194 (3)	Visual, 195 (3)

Note: Priority ranking is shown in parentheses.

choice is not to inspect, i.e., inspect the link, do not inspect the link + keep link open, do not inspect the link + reduce speed, do not inspect the link + close the link. When the optimal choice is to inspect the link, then the optimal inspection choice for each component is also determined. At the component level, there are three inspection choices: no inspection, visual inspection, and extensive inspection. Expected utility values for the optimal decision for each link and each evidence case are recorded. These correspond to the quantity EU|I defined in Eq. (8) plus the costs of inspection. Next, the link-level inspection decision node is set to *no inspection* and the analysis is repeated for each link and each evidence case. The expected utility of the optimal decision in this case represents the quantity EU|NI in Eq. (8). The difference between the two expected utilities is the net benefit of inspection (NBI) for the link. When the NBI for a link and evidence case is negative, an action decision is made without previous inspection. Otherwise, the NBI values are used to rank the priority of link inspections.

Table 3 summarizes these results for the links. Included are the optimal decision for each link, i.e., to inspect or not to inspect, and the NBI and priority ranking when the decision is to inspect the link. This table shows that the optimal decision for Links 1 and 2 is to inspect under all evidence cases. For Link 3, the optimal decision under EC 1 is to inspect and under ECs 2–4 is not to inspect, but reduce the train speed. For Link 4, the optimal decision under ECs 1–3 is to inspect, but under EC 4 is not to inspect and to close down the link. For links that are to be inspected, the priority rankings of inspections are listed in parentheses in Table 3. For EC 1, Link 1 has the highest priority for inspection followed in order by Links 4, 2, and 3. Under EC 2 and 3, the priority for inspection is still highest for Link 1, but now followed by Links 2 and 4. Under EC 4, the priority for inspection is Link 2 followed by Link 1. These results are related to the assumed costs, component failure probabilities, inspection performance (i.e., assumed test likelihood matrices) and the number of components along each link. There are too many factors involved to provide simple explanations of all results. However, some conclusions are obvious. For example, the decision not to inspect Link 3 under ECs 2–4 is because of the high cost of inspection of some of its components and its closure. The decision not to inspect Link 4 under EC 4 and to shut it down is caused by the knowledge that a component along this link has experienced heavy damage. The change in inspection rankings from EC 1 (Links 1–4–2–3) to EC 2 (Link 1–2–4) probably has to do with the observation of a high intensity ground motion at GMPP 2.

For each link that is to be inspected, the optimal inspection options for its constituent components and the corresponding rankings based on their NBI values are also determined. Table 4 summarizes the results for the components of Link 1. Included are the type of inspection, the NBI value, and the ranking (in parentheses) for each component under each evidence case. Visual inspection of all components is recommended under all evidence cases. Under EC 1, Component 1 has the highest priority for inspection followed by Components 4, 3, and 2. Under EC 2, the priority order changes to Components 1, 3, 4, and 2. Practically no change in the NBI values occurs under EC 3. This is because the components in Link 1 are not affected by the potential rupture directivity effect. Under EC 4, the NBI values change, but the priority ranking remains the same as under EC 2. Similar results are obtained for the components of Links 2–4. These include cases in which the optimal choice is to inspect some but not all of the components of a link. For the sake of brevity, the results for the other links are not reported in this study.

The results presented in this paper were obtained on a workstation with 32 GB of RAM and the Hugin software (Hugin Expert a/s 2008). In the context of near-real-time applications, computation times for each run required by the VoI heuristic were in the order of seconds (for the rail link with the fewest components) to an hour (for the link with the most components).

Summary and Conclusions

A framework for post-hazard risk assessment and decision making for infrastructure systems is developed. The framework makes use of Bayesian networks (BNs) and influence diagrams (IDs) as tools for probabilistic modeling, information updating, and decision making. Specific BN models are developed for the case of earthquake hazard, including effects of spatial correlation of ground motion, and for component and system performance. This represents a novel application of the BN-ID framework. The resulting 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 fully describe probabilistic dependencies and model general system configurations. Most importantly, the BN allows rapid processing of information to update probabilistic characterization of the hazard and system states and incorporates this information into subsequent decision making.

IDs for decision making at the component and system level are developed. These allow decision making for setting the operational level of a component or link, or to inspect or not inspect the component before making a decision on the operation level. At the

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

system level, there is 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* link or component to inspect. Although it may generate potentially suboptimal solutions, the heuristic allows relatively 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. Because the system is still under development, highly idealized models of the system components and hazard environment 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 shows the power of the BN-ID framework to process information and to prioritize decision alternatives based on updated probabilistic information for an infrastructure system immediately after a hazard.

Although the proposed framework provides a novel, transparent, and efficient mechanism for post-event risk assessment and decision support, it is important to acknowledge that challenges remain, particularly with respect to computational (memory) demands for application to real-life, large infrastructure systems. Additional work is needed to leverage new or different inference algorithms or develop even more efficient modeling approaches to improve the feasibility of near-real-time applications involving large infrastructure systems. Research along these lines is ongoing.

Although the earthquake hazard was the focus in this paper, the proposed framework is equally applicable to other natural and artificial hazards. In each case, what needs to be done is 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 done in this paper. The system performance and decision models developed in the paper are hazard-independent. Recent tragic events, such as Hurricane Sandy that caused massive failure of transportation and power systems in the East Coast, highlight the need for full-scale development of such decision-support systems to aid emergency responders and those responsible for managing the post-hazard decision-making process.

Acknowledgments

This research was supported by the State of California through the Transportation Systems Research Program of the Pacific Earthquake Engineering Research Center (PEER). Additional support was provided by the Taisei Chair in Civil Engineering at the University of California, Berkeley. The first author also gratefully acknowledges the support from a National Science Foundation graduate research fellowship while a doctoral student at UC Berkeley.

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 those of the funding agencies.

References

Abrahamson, N. (2000). "Effects of rupture directivity on probabilistic seismic hazard analysis." *Proc., 6th Int. Conf. on Seismic Zonation*, Earthquake Engineering Research Institute, Oakland, CA.

Abrahamson, N., et al. (2008). "Comparisons of the NGA ground-motion relations." *Earthquake Spectra*, 24(1), 45–66.

Bayraktarli, Y. Y., Baker, J., and Faber, M. (2011). "Uncertainty treatment in earthquake modelling using Bayesian probabilistic networks." *Georisk*, 5(1), 44–58.

Bayraktarli, Y. Y., and Faber, M. (2011). "Bayesian probabilistic network approach for managing earthquake risk of cities." *Georisk*, 5(1), 2–24.

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

Bensi, M., Der Kiureghian, A., and Straub, D. (2011b). "Bayesian network modeling of correlated random variables drawn from a Gaussian random field." *Struct. Saf.*, 33(6), 317–332.

Bensi, M., Der Kiureghian, A., and Straub, D. (2013). "Efficient Bayesian network modeling of systems." *Reliab. Eng. Syst. Saf.*, 112, 200–213.

Blaser, L., Ohrnberger, M., Riggelsen, C., Babeyko, A., and Scherbaum, F. (2011). "Bayesian networks for tsunami warning." *Geophys. J. Int.*, 185(3), 1431–1443.

Bray, J. D., and Rodriguez-Marek, A. (2004). "Characterization of forward-directivity ground motions in the near-fault region." *Soil Dyn. Earthquake Eng.*, 24(11), 815–828.

Bryant, W. A. (2005). "Digital database of quaternary and younger faults from the fault activity map of California, version 2.0." ⟨http://www.consrv.ca.gov/CGS/information/publications/QuaternaryFaults_ver2.htm⟩.

California High-Speed Rail Authority. (2010). "California high-speed rail authority website." ⟨http://www.cahighspeedrail.ca.gov/⟩ (Mar. 29, 2010).

Dechter, R. (1996). "Bucket elimination: A unifying framework for probabilistic inference." *Proc., 12th Int. Conf. on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers, San Francisco, 211–219.

Decision Systems Laboratory (DSL). (2007). *GeNIe 2.0 by Decision Systems Laboratory*, ⟨http://genie.sis.pitt.edu/⟩ (Nov. 30, 2007).

DeGroot, M. H. (1969). *Optimal statistical decisions*, McGraw-Hill, New York.

Frey, D., Butenuth, M., and Straub, D. (2012). "Probabilistic graphical models for flood state detection of roads combining imagery and DEM." *IEEE Geosci. Remote Sens. Lett.*, 9(6), 1051–1055.

Hugin Expert A/S. (2008). *Hugin researcher API 7.0*, Aalborg, Denmark.

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

Jensen, F. V., and Vomlelova, M. (2002). "Unconstrained influence diagrams." *Proc., 18th Conf. on Uncertainty in Artificial Intelligence*, A. Darwiche and N. Friedman, eds., Morgan Kaufmann, San Francisco, 234–241.

Kotsiantis, S., and Kanellopoulos, D. (2006). "Discretization techniques: A recent survey." *GESTS Int. Trans. Comput. Sci. Eng.*, 32(1), 47–58.

Kuehn, N. M., Riggelsen, C., and Scherbaum, F. (2011). "Modeling the joint probability of earthquake, site, and ground-motion parameters using Bayesian networks." *Bull. Seismol. Soc. Am.*, 101(1), 235–249.

Langseth, H., Nielsen, T. D., Rumí, R., and Salmerón, A. (2009). "Inference in hybrid Bayesian networks." *Reliab. Eng. Syst. Saf.*, 94(10), 1499–1509.

Lauritzen, S. L., and Nilsson, D. (2001). "Representing and solving decision problems with limited information." *Manage. Sci.*, 47(9), 1235–1251.

Lauritzen, S. L., and Spiegelhalter, D. J. (1988). "Local computations with probabilities on graphical structures and their application to expert systems." *J. R. Stat. Soc. Ser. B (Methodol.)*, 50(2), 157–224.

Madsen, A. L. (2008). "Belief update in CLG Bayesian networks with lazy propagation." *Int. J. Approximate Reasoning*, 49(2), 503–521.

Neil, M., Tailor, M., and Marquez, D. (2007). "Inference in hybrid Bayesian networks using dynamic discretization." *Stat. Comput.*, 17(3), 219–233.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*, Morgan Kaufmann, San Francisco.

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.

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." *Seismol. Res. Lett.*, 68(1), 199–222.

Straub, D., and Der Kiureghian, A. (2010a). "Bayesian network enhanced with structural reliability methods: Application." *J. Eng. Mech.*, 10.1061/(ASCE)EM.1943-7889.0000170, 1259–1270.

Straub, D., and Der Kiureghian, A. (2010b). "Bayesian network enhanced with structural reliability methods: Methodology." *J. Eng. Mech.*, 10.1061/(ASCE)EM.1943-7889.0000173, 1248–1258.

von Neumann, J., and Morgenstern, O. (1944). *Theory of games and economic behavior*, Princeton University Press, Princeton, NJ.

Wells, D., and Coppersmith, K. (1994). "New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement." *Bull. Seismol. Soc. Am.*, 84(4), 974–1002.

Yuan, C., and Druzdzel, M. (2003). "An importance sampling algorithm based on evidence pre-propagation." *Proc., 19th Annual Conf. on Uncertainty in Artificial Intelligence (UAI03)*, Morgan Kaufmann, San Francisco, 624–631.

Yuan, C., and Druzdzel, M. (2006). "Importance sampling algorithms for Bayesian networks: Principles and performance." *Math. Comput. Modell.*, 43(9), 1189–1207.

ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.