Sensor Network Optimization using Bayesian Networks, Decision Graphs, and Value of Information

Carl Malings
Graduate Student, Dept. of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, USA

Matteo Pozzi
Assistant Professor, Dept. of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, USA

ABSTRACT: Bayesian Networks (BNs) and decision graphs provide a useful framework for modeling the uncertain behavior of civil engineering infrastructures subjected to various risks, as well as the potential outcomes of risk mitigation actions undertaken by managing agents. These graphs can also guide optimal sensing and inspection of infrastructure by maximizing the value of information of sensing efforts. This paper presents a general framework for modeling infrastructure systems using BNs and for evaluating sensor placement metrics within this model. An example application of the use of the value of information metric in guiding optimal sensing in a system of infrastructure assets in the San Francisco Bay area subjected to seismic risk is then presented. A parametric study also investigates the sensitivity of the value of information metric to various parameters of the BN system model.

1. INTRODUCTION

Effective management of infrastructure requires information about the status of the infrastructure system so that managers may make informed decisions to minimize potential losses, in terms of lost revenues or potential harm to the public. In order to cost-effectively collect this information, an optimization strategy should be used, where the benefits of additional information should be weighed against their costs. This paper presents such a strategy, in which a probabilistic model of an infrastructure system is used to optimize the value of information of a sensor network monitoring the system.

Analyzing the benefits of a sensor network before it is put in place is a form of pre-posterior analysis. Such analysis requires a model of the sensed system, such that potential sensor measurements may be predicted and their consequences assessed. We propose a Bayesian Network model for this purpose. Bayesian Networks (BNs) are a type of probabilistic graphical model (PGM). PGMs represent physical systems using random variables with a joint probabilistic distribution. Relationships between variables are encoded graphically. The reader is referred to Koller and Friedman (2009) for further background on PGMs. BNs have applications in the modeling of infrastructure systems; Bensi et al. (2014) present a BN model of a transportation system subjected to seismic risk. By modeling outputs of sensors as observable variables within a BN, pre-posterior analysis can also be performed. In an infrastructure system BN, similarities between components are encoded in the model, allowing observations of one sensor to provide information about multiple components.

Within this model for pre-posterior analysis, the usefulness of potential measurements must be quantified, such that optimization can be performed. A common metric for this purpose is the conditional entropy metric, which quantifies uncertainty in one set of random variables conditioned on observations of another set. The reader is referred to Cover and Thomas (2006) for further information on conditional entropy. If the goal of a sensor network is characterized as reducing uncertainty in the system, this metric can be used to quantify this goal and perform sensor network optimization, as discussed by Krause (2008).

In the case of infrastructure management, however, uncertainty reduction is not necessarily the goal of sensing. Infrastructure managers use information from sensors to guide decision-making about which actions to take to reduce long-term management costs for the system. Potential actions might include closing an unsafe component to prevent injury to the public or leaving the component in operation until it fails. Each action comes with certain costs or losses, such as lost revenues or potential injuries. To assess the benefits of information in such decision-making problems, we use the value of information metric, which quantifies the reduction in expected loss in a decision-making problem due to the availability of the information. An introduction to value of information is presented in this paper, and a more complete background is provided by Raiffa and Schlaifer (1961).

The decisions of managers, as well as the potential costs of different outcomes, can be included explicitly into a BN model of an infrastructure system. Such a model is referred to as a decision graph or influence diagram. Pre-posterior analysis can then be conducted in the decision graph using the value of information metric, and an optimal sensor network can be designed for the monitoring of the modeled system. In this paper, we present a general method for creating such a model and performing this optimization. We demonstrate this method on an example infrastructure system in the San Fransisco Bay area subjected to seismic risk, as well as on a simple system to perform a parametric analysis on the value of information metric.

2. PROBLEM FORMULATION

We now present a BN model of a civil infrastructure system, extend this model to a decision graph, and use this graph to define a metric for optimal sensor placement for management of the system. Consider an infrastructure system made of n binary components, which can either be operational or not. The system is subjected to a random risk scenario parameterized by S. Although this scenario is uncertain a priori, we assume that, after the occurrence of the scenario, information about the scenario is available to infrastructure managers. The functioning of component i is governed by (potentially multi-dimensional) variable W_i. From this variable, a binary random variable $X_i \in \{0,1\}$ can be defined as the state variable of component i, i.e., component i is operational if $X_i = 1$ and has failed if $X_i = 0$. Together, the set of variables $X = \{X_1,\dots,X_n\}$ describes the state of every component of the system.

For component i, the manager of this infrastructure system must select an action $a_i \in A_i$ for the management of this component. Depending on the chosen action a_i and the state of the component x_i, the manager incurs some loss (or reward) $L_i(x_i,a_i)$, a deterministic outcome of the action and state. Finally, the total loss incurred by the manager for the entire system, $L^{\mathrm{G}}(X,A)$, is a function of the joint state of all system components, as well as all selected actions. Here, we assume that this system loss is the sum of the losses incurred in the management of each component:

$$L^{\mathrm{G}}(X,A) = \sum_{i=1}^{n} L_i(x_i,a_i) \tag{1}$$

In selecting which actions to take, an infrastructure system manager may have the ability to observe certain variables, from which he or she may infer the states of components within the system. We denote with $Z = \{Z_1,\dots,Z_n\}$ the set of all observations which might be made by the manager, where Z_i is the set of observations related to component i, that is, $Z_i \subseteq W_i$. Due to limited time and resources, the manager may not be able to make all possible observations; therefore, he or she must select a subset $Y = \{Y_1,\dots,Y_n\}$ of these variables to measure, where $Y_i \subseteq Z_i$, and base their decision on the specific outcome $Y = y$ of these measurements. Note that Y_i may be an empty set if no observations relating to component i are selected.

A graphical representation of the infrastructure model described above is shown in Figure 1. The symbols used follow a common convention, described by Koller and Friedman (2009); circles represent random variables, squares indicate agent actions, and rhombi denote utilities or losses. Arrows or lines indicate relationships between variables. Dashed arrows indicate temporal precedence. Shaded circles indicate observed variables.

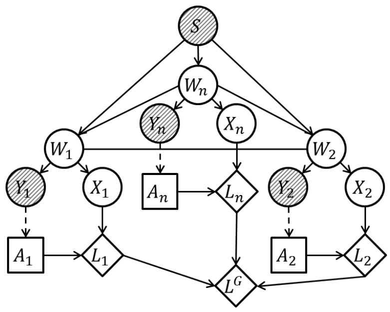

Figure 1: A decision graph for an infrastructure system with n components. Variable S describes the risk scenario, variable W_i describes component i, X_i describes its state, and observation Y_i is made on certain observable features of the component $Z_i \subseteq W_i$. Action A_i is chosen by the manager of the infrastructure system, who incurs loss L_i for component i, and global loss L^{G} for the whole system.

Within the decision graph of Figure 1, we define an optimal sensor network as follows:

$$Y^* = \mathrm{argmax}_{Y \subseteq Z} m_X(Y) \text{ subject to } \mathrm{C}(Y) \leq B \quad (2)$$

where Y^* is the optimal set of sensed variables, $C(Y)$ is the cost of measuring Y, B is a fixed budget constraint, and $m_X(Y)$ is a metric which quantifies how observing Y will improve the managing agent's ability to effectively and efficiently manage the infrastructure system. We denote this metric as $m_X(Y)$ since X describes the functionality of each component, and is therefore of primary interest for the management of the system. Our metric should therefore depend on how well sensor placement Y improves the manager's knowledge about X.

It is necessary that the optimal sensor placement Y^* be robust under uncertainty in the risk scenario S. Since S is not known a priori, we must compute our metric for specific values of $s \in S$; we denote these scenario-specific metric values with $m_{X|s}(Y)$. We then compute $m_X(Y)$ by taking the expected value over potential scenarios:

$$m_X(Y) = \mathbb{E}_S[m_{X|s}(Y)] = \int_s m_{X|s}(Y)\mathrm{p}(s)\mathrm{d}s \quad (3)$$

where $\mathbb{E}_S[\cdot]$ represents the statistical expectation under the distribution $\mathrm{p}(S)$ of S. To evaluate (3), we adopt a Monte Carlo sampling approach:

$$m_X(Y) \approx \frac{1}{n_s} \sum_{i=1}^{n_s} m_{X|s_i}(Y) \quad (4)$$

where n_s is the number of scenarios $s_1, \ldots, s_{n_s}$ sampled independently from $\mathrm{p}(S)$.

3. VALUE OF INFORMATION FOR SENSOR PLACEMENT

Value of information quantifies the explicit benefit, in terms of a reduction in expected losses, that an infrastructure manager would see after implementing a sensor network to measure Y. For component i, the expected loss for this component under scenario $s \in S$, without any observations of Y, is:

$$\mathbb{E}L_{i,s}(\emptyset) = \min_{A_i} \mathbb{E}_{X_i|s}\left[L_i(x_i, a_i)\right] \quad (5)$$

That is, a manager should choose an action $a_i \in A_i$ which minimizes his or her expected loss under the possible outcomes of X_i in scenario s. Note that, with binary components, there are only two outcomes for X_i: 0 (a failure) or 1 (no failure).

Given an observation y of Y, the managing agent can update the probability distribution of the state of X_i based on the conditional distribution $\mathrm{p}(X_i|y,s)$. Taking this into account, the expected loss given that an observation of Y will be available before an action is chosen is given in Eq. (6):

$$\mathbb{E}L_{i,s}(Y) = \mathbb{E}_{Y|s}\left[\min_{A_i} \mathbb{E}_{X_i|y,s}\left[L_i(x_i, a_i)\right]\right] \quad (6)$$

where the inner expectation is taken over the conditional distribution of X_i given a specific observation of $Y = y$, and the outer expectation is taken over the distribution for these observations in scenario $S = s$. The outer expectation accounts for the fact that the observation y is not known a priori, but the optimal action will depend on this observation (thus, the minimization over A_i is within this expectation).

For component i, the value of information of observing Y before choosing an action is the difference between the expected loss for this component given this information and the expected loss without information:

$$\mathrm{VoI}_{i,s}(Y) = \mathbb{E}L_{i,s}(\emptyset) - \mathbb{E}L_{i,s}(Y) \quad (7)$$

That is, the value of information of Y is the decrease in expected loss due to the availability of an observation of Y prior to making a decision.

For the entire system, since the system loss is equivalent to the sum of individual component losses, by Eq. (1), and since the expectation is linear, the value of information of Y for the system is:

$$\mathrm{VoI}_s(Y) = \mathbb{E}L_s^{\mathrm{G}}(\emptyset) - \mathbb{E}L_s^{\mathrm{G}}(Y) = \sum_{i=1}^{n} \mathrm{VoI}_{i,s}(Y) \quad (8)$$

To maximize the net benefit of a sensor network, the difference of the value provided by this sensor network, quantified as $\mathrm{VoI}_s(Y)$, and the cost of implementing this network, $\mathrm{C}(Y)$, should be maximized. We can use this difference as a metric to assess sensor network Y under scenario s:

$$m_{X|s}(Y) = \mathrm{VoI}_s(Y) - \mathrm{C}(Y) \quad (9)$$

where the dependence of the metric on X is implicit in the formulation of the value of information.

4. METHODOLOGY

In section 4.1, a general method for computing $m_X(Y)$ in a PGM as outlined in Figure 1 is described. In section 4.2, details for implementing this method in a Gaussian graphical model are briefly presented. The reader is referred to Malings and Pozzi (2014) for details. In section 4.3, an algorithm is given for efficiently approximating Eq. (2).

4.1. General metric evaluation

The evaluation of a generic metric $m_X(Y)$ for some proposed sensor network measuring Y in an infrastructure system PGM as outlined in section 2 can be performed by the following steps:

1. Begin with the distributions $\mathrm{p}(S)$, $\mathrm{p}(W|s)$, $\mathrm{p}(X|w,s)$, and $\mathrm{p}(Y|w,s)$ which parameterize the PGM.
2. Marginalize over W to obtain distributions $\mathrm{p}(X|s)$ and $\mathrm{p}(Y|s)$.
3. For any observation $Y = y$, perform inference within the PGM to obtain the updated distribution $\mathrm{p}(W|y,s)$.
4. Use this distribution to obtain an updated distribution $\mathrm{p}(X|y,s)$.
5. Define a mapping $m_{X|y,s}(Y)$ from the probability distributions $\mathrm{p}(X|s)$ and $\mathrm{p}(X|y,s)$ to the metric to be evaluated. Note that we have assumed in section 2 that the value of metric $m_{X|s}(Y)$ depends on how an observation y of Y allows for updating knowledge about X under scenario s, from $\mathrm{p}(X|s)$ to $\mathrm{p}(X|y,s)$.
6. Take the expectation of $m_{X|y,s}(Y)$ over $\mathrm{p}(Y|s)$ to compute $m_{X|s}(Y)$. This accounts for the fact that observation y is not known a priori.
7. Take the expectation of $m_{X|s}(Y)$ over $\mathrm{p}(S)$ to compute $m_X(Y)$, as in Eq. (3), following the approach of Eq. (4).

Details on marginalization and inference in PGMs are given by Koller and Friedman (2009).

Note that for the value of information metric, $m_{X|y,s}(Y) = \sum_{i=1}^{n} \{\nu_{i,s}(\emptyset) - \nu_{i,s}(y)\} - \mathrm{C}(Y)$, where $\nu_{i,s}(\emptyset) = \min_{A_i} \mathbb{E}_{X_i|s}[L_i(x_i,a_i)]$ and $\nu_{i,s}(y) = \min_{A_i} \mathbb{E}_{X_i|y,s}[L_i(x_i,a_i)]$. Evaluation of the value of information metric therefore only involves the marginal distributions $\mathrm{p}(X_i|y,s)$ and $\mathrm{p}(X_i|s)$, rather than the joint distributions $\mathrm{p}(X|y,s)$ and $\mathrm{p}(X|s)$. The ability to evaluate value of information in this way is due to the decomposition of system-level loss into the sum of component-level losses in Eq. (1). Computation of the value of information is therefore relatively efficient, since $\mathrm{p}(X_1|y,s),\ldots,\mathrm{p}(X_n|y,s)$ only require the computation of n values, i.e., the probability of failure for each component, to parameterize them while $\mathrm{p}(X|y,s)$ requires $2^n - 1$ values to parameterize.

4.2. Gaussian linear models

Following the approach of Ditlevsen and Madsen (1996), we assume that component i is fully characterized by the demand (d_i) and the capacity to resist this demand (c_i), i.e., $W_i = \{c_i, d_i\}$. Furthermore, we assume that under scenario s, these variables, via an appropriate transformation, can be represented by joint Gaussian distributions:

$$\mathbf{w} = \begin{bmatrix} \mathbf{c} \\ \mathbf{d} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_C \\ \mu_D \end{bmatrix}, \begin{bmatrix} \Sigma_C & \mathbf{0} \\ \mathbf{0} & \Sigma_D \end{bmatrix} \right) \quad (10)$$

where $\mathbf{d} = [d_1,\ldots,d_n]^{\mathrm{T}}$ and $\mathbf{c} = [c_1,\ldots,c_n]^{\mathrm{T}}$, μ_D and μ_C are the corresponding means, Σ_D and Σ_C the covariance matrices, and $\mathbf{0}$ is an n by n matrix of zeros. A multivariate Gaussian distribution is denoted by $\mathcal{N}(\cdot,\cdot)$, with the mean vector as the first argument and the covariance matrix as the second. This model assumes capacities are marginally independent of demands. We also assume that candidate observation variables Z are noisy measurements of the capacity and demand for each component, i.e. $\mathbf{z} = \mathbf{w} + \varepsilon$, where ε represents the random error of these observations, assumed to be a zero mean Gaussian random vector independent of $\mathbf{w}$ with a diagonal covariance matrix.

To encode the proposed sensor network Y, we use the matrix $\mathbf{A}$. For m observed variables selected from $2n$ potentially observed variables, $\mathbf{A}$ is an m by $2n$ matrix. Each row of $\mathbf{A}$ has all entries zero, except for an entry of one in the position corresponding to the selected observable variable. We model the observations of sensors as:

$$\mathbf{y} = \mathbf{A}\mathbf{z} = \mathbf{A}(\mathbf{w}+\varepsilon) \tag{11}$$

The relationship between X and W is given by:

$$\mathbf{x} = \mathbb{I}(\mathbf{c}-\mathbf{d} \geq 0) = \mathbb{I}\left(\begin{bmatrix} \mathbf{I} & -\mathbf{I} \end{bmatrix}\mathbf{w} \geq 0\right) \tag{12}$$

where $\mathbb{I}[\cdot]$ is an indicator function, taking on value 1 when its argument is true and 0 otherwise, and $\mathbf{I}$ represents and n by n identity matrix. Under this model, a component fails when the demand placed on it exceeds its capacity.

Equations (10), (11), and (12) define the distributions for $\mathrm{p}(W|s)$, $\mathrm{p}(Y|w,s)$, and $\mathrm{p}(X|w,s)$ respectively, where the latter is a deterministic function of $\mathbf{w}$. These distributions, together with $\mathrm{p}(S)$, which gives a distribution on the parameters μ_D, μ_C, Σ_D, and Σ_C (which may differ for different scenarios), are all that are needed for the general procedure outlined in section 4.1. Inference and marginalization for Gaussian random variables are outlined by Koller and Friedman (2009).

4.3. Greedy sensor placement

An exact solution of the optimal sensor placement problem of Eq. (2) would require the enumeration of all subsets Y of Z, which is computationally prohibitive in all but the smallest problems. An alternative, approximate solution approach involves a heuristic known as the greedy algorithm, as discussed by Krause (2008). In this method, single elements of Z are iteratively added to the set Y, where these elements most improve the objective to be optimized. Pseudo-code for implementation of the greedy algorithm is given in Algorithm 1.

```
Input: Z; m_X(·); C(·); B
Y ← ∅ ;
while C(Y) < B, |Z| > 0 do
    y* ← argmax_{y∈Z} m_X(y ∪ Y) ;
    Y ← Y ∪ y* ;
    Z ← Z\y* ;
    foreach z ∈ Z do
        if C(Y ∪ z) > B then
            Z ← Z\z ;
        end
    end
end
return Y ;
```

Algorithm 1: Pseudo-code for the greedy algorithm. $m_X(\cdot)$ is evaluated as outlined in section 4.1. Based on algorithms of Krause (2008).

5. EXAMPLE APPLICATION TO SEISMIC RISK IN SAN FRANCISCO

As an illustrative application of sensor placement optimization to a practical infrastructure management problem, we examine an infrastructure system consisting of 18 bridges and 9 tunnels in the San Francisco Bay area subjected to seismic risk. It should be noted that this system will merely serve to illustrate the application of the metric and techniques discussed above, and is not meant to be a practical recommendation. Seismic risk scenarios are modeled using a homogeneous Poisson process for earthquake occurrence, based on the model outlined by Anagnos and Kiremidjian (1988), with empirical data for the San Francisco Bay area presented by Field et al. (2009) and USGS (2008). Using this generative model, for the evaluation of (4), $n_s = 1000$ sample seismic scenarios are generated, with each scenario s consisting of an earthquake with magnitude M and epicenter location E.

Earthquake demands are defined in terms of peak ground acceleration, using attenuation equations presented by Douglas (2011). Ground accelerations are modeled as lognormal random variables; under the logarithm transformation, these become Gaussian random variables, as in section 4.2. Correlations between demands are modeled by a squared-exponential kernel function, as in Bensi et al. (2014). Capacities of components are also modeled as lognormal random variables using fragility curves presented in Hazus (2012). Correlations between capacity variables are assumed to be higher for components with a similar overall typology. These lognormal variables are collected in the vector $\mathbf{w}'$, such that $\log(\mathbf{w}') = \mathbf{w}$ is Gaussian.

It is assumed that a binary decision must be undertaken as to whether or not to close each potentially damaged bridge or tunnel in the wake of the event. The option to close down the component comes with a certain cost in terms of lost toll revenues and service loss, which is roughly estimated for each component considered. If the component is not closed, no costs would be incurred, but if the component is severely damaged and fails while in use, a high cost of failure is incurred.

Detailed measurements of the capacities of each component and/or the ground acceleration at each location would potentially be available, at a certain cost. These costs are assumed to consist of an installation cost for sensors as well as ongoing maintenance costs for the sensor network which are discounted to their present value using a discount rate of 5%. Sensor noise is modeled as a multiplicative lognormal error ε' with median 1, such that under a logarithm transformation, this noise would be zero mean additive Gaussian noise ε, as in Eq. (11):

$$\log(\mathbf{w}'\varepsilon') = \log(\mathbf{w}') + \log(\varepsilon') = \mathbf{w} + \varepsilon \qquad (13)$$

Further details of the model outlined above are provided in Malings and Pozzi (2014).

Figure 2 shows the optimal measurement selections for the management of this infrastructure system based on the value of information metric of Eq. (9). Capacity measurements for the Golden Gate Bridge and the Caldecott tunnel, as well as demand measurements at both locations and at the San Francisco-Oakland Bay Bridge, are indicated as the optimal measurement set for the management of this example infrastructure system.

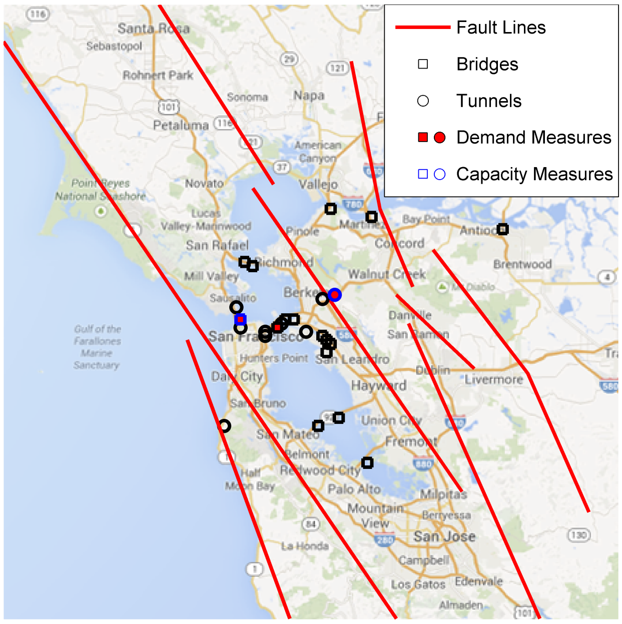

Figure 2: Optimal measure selections based on the value of information metric. Background image from www.maps.google.com.

6. PARAMETRIC ANALYSIS

A parametric study is presented to provide insight into the sensitivity of the value of information metric to some of the parameters of the BN used to represent an infrastructure system. This study is performed on an example system with 12 components subjected to a single magnitude 7 earthquake scenario, as shown in Figure 3. Demands on all components are computed for this scenario as discussed in section 5. Component capacities are defined using lognormal distributions, with median capacities for each type as listed in Table 1 and coefficient of variation of 0.6 for all types. Components of the same type have correlated capacities, with the correlation coefficient listed in Table 1 between any pair of components of the same type. The binary decision described in section 5 is again used, with closure cost set at $10 million for all components and failure cost set as shown in Table 1. Sensors are modeled as having multiplicative lognormal noise

Table 1: BN model parameters for the system of Figure 3. The parameters to be varied are C_f, μ_C^{II}, and ρ.

Component Type	Cost of Failure [$M]	Median Capacity [g]	Correlation Coefficient
I	C_f	1	0.1
II	100	μ_C^{II}	0.1
III	100	1	ρ

with median 1 and coefficient of variation 0.2. All sensors are assigned a cost of $1 million, and a budget constraint of $1 million is used, such that only 1 sensor will be selected. For this parametric study,

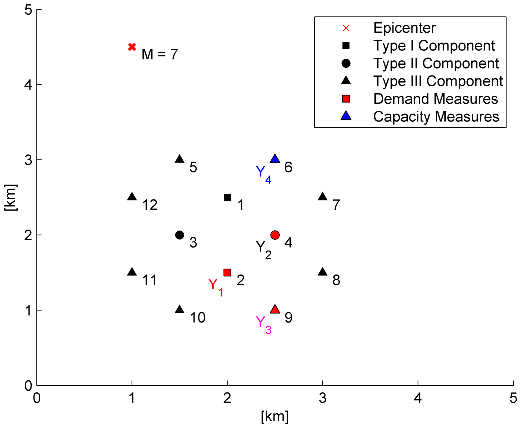

Figure 3: Example system for this parametric study. Example measures Y_1, Y_2, Y_3, and Y_4 are shown.

we track the value of information for four proposed measurements, Y_1, Y_2, Y_3, and Y_4, each of which are selected as the most optimal measurement for the system under a specific setting of the parameters C_f, $\mu_{C(\mathrm{II})}$, and ρ.

We begin by studying the effect of varying failure cost parameter C_f from 100 to 10 for components of type I. At $C_f = 100$, demand measurement Y_1 on component 2 has the highest value of information. As C_f is lowered, the relative value of information for Y_2 compared to Y_1 grows as C_f is decreased. This is due to the lower expected costs for components of type I compared to components of other types as C_f decreases. It eventually becomes more worthwhile to directly monitor a component with

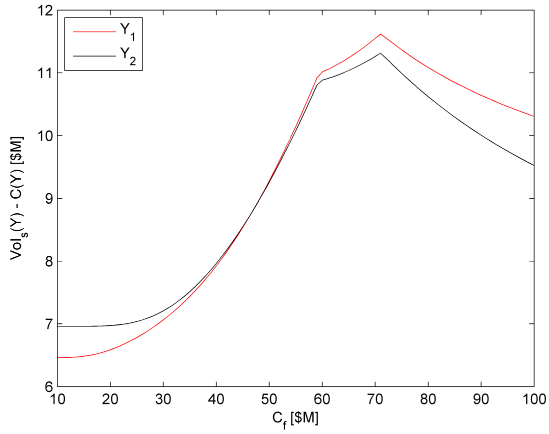

Figure 4: Parametric study on C_f, with $\mu_C^{\mathrm{II}} = 1$ and $\rho = 0.1$.

higher expected costs. The two sharp bends in each curve of Figure 4 are due to the choice of optimal management action for components 1 and 2 changing as the relative costs of these actions change.

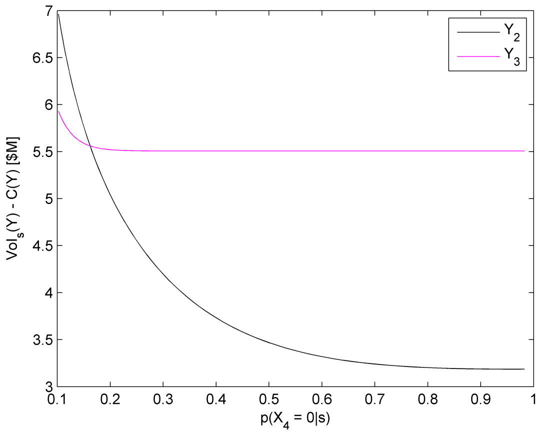

Figure 5: Parametric study on μ_C^{II}, with $C_f = 10$ and $\rho = 0.1$. Probability of failure for component 4 is plotted on the horizontal axis.

Next, the median capacity of type II components, μ_C^{II}, is varied from 1 to 0.1, increasing the probability of failure for these components. The probability of failure of component 4, where measure Y_2 is taken, is used as the axis of Figure 5. As this probability of failure increases, the value of measuring Y_2 declines, as this component is more likely to fail, and therefore the optimal action, to close the component, is no longer in question. It soon becomes more valuable to take measure Y_3 instead, as the optimal management action for the component where this measure is taken is still unclear.

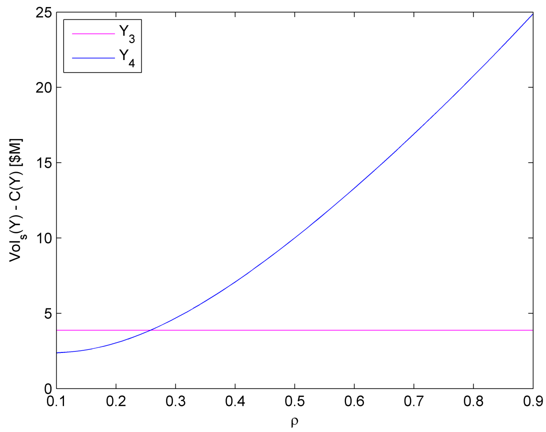

Figure 6: Parametric study on ρ, with $C_f = 10$ and $\mu_C^{\mathrm{II}} = 0.1$.

Finally, the correlation coefficient for components of type III, ρ, is varied from 0.1 to 0.9, as shown in Figure 6. As this coefficient increases, the value of information for measuring Y_4, the capacity of a component of type III, also increases, since information about the capacities of many components in the system will be gained from this single measurement. Note that $\mathrm{VoI}_s(Y_3)$ remains constant.

7. CONCLUSIONS

This paper presents a general framework for modeling systems of infrastructure using BNs and decision graphs. Within such a model, methods for computing sensor placement metrics in general, and the value of information metric in particular, are presented. These methods are then demonstrated using an example system of infrastructure assets in the San Francisco Bay area subjected to seismic risk. A parametric analysis is also conducted to demonstrate the sensitivity of the value of information metric to various BN parameters. The ability to directly trade off value of information and sensor system cost using the metric demonstrated above can allow for senor network selections which maximize the net benefit of the information provided by these networks in terms of minimizing expected losses for decision-making problems in infrastructure management.

8. REFERENCES

Anagnos, T. and Kiremidjian, A. (1988). "A review of earthquake occurrence models for seismic hazard analysis." *Probabilistic Engineering Mechanics*, 3(1), 3–11.

Bensi, M., Der Kiureghian, A., and Straub, D. (2014). "Framework for post-earthquake risk assessment and decision making for infrastructure systems." *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 9, 1–17.

Cover, T. and Thomas, J. (2006). *Elements of Information Theory*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 2 edition.

Ditlevsen, O. and Madsen, H. O. (1996). *Structural Reliability Methods*. John Wiley & Sons Ltd, Chichester, UK, 1 edition.

Douglas, J. (2011). "Ground-motion prediction equations, 1964-2010." *PEER PEER Report 2011/102*, Pacific Earthquake Engineering Research Center, College of Engineering, University of California, Berkeley, Berkeley, California, USA (April).

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

Hazus (2012). "Hazus-MH 2.1 technical manual: Earthquake model – multi-hazard loss estimation methodology." *Report no.*, Department of Homeland Security, Federal Emergency Management Agency, Washington, DC, USA.

Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

Krause, A. (2008). "Optimizing sensing: Theory and applications." Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA 15213 (December).

Malings, C. and Pozzi, M. (2014). "Conditional entropy and value of information based metrics for sensor placement in infrastructure systems." *Proceedings of the 6th World Conference on Strucutral Control and Monitoring*, Barcelona, Spain (July).

Raiffa, H. and Schlaifer, R. (1961). *Applied statistical decision theory*. Harvard University Press, Cambridge, Mass., USA.

USGS (2008). "Uniform california earthquake rupture forecast." *Report No. 3027*, United States Geological Survey (April).