The 6th International Conference on
Structural Health Monitoring of Intelligent Infrastructure
Hong Kong | 9-11 December 2013

OPTIMAL TOPOLOGY OF SENSOR NETWORKS FOR MANAGEMENT OF INFRASTRUCTURE SYSTEMS

C. Malings [1], M. Memarzadeh [1] and M. Pozzi [1]
[1] Department of Civil and Environmental Engineering,
Carnegie Mellon University, Pittsburgh, PA, USA. Email: mpozzi@cmu.edu

ABSTRACT

Many spatially distributed phenomena relevant to the management of civil infrastructure, from temperature in buildings to earthquake ground excitation, can be modeled as Gaussian Processes (GPs). GPs generalize the concept of multivariate Gaussian distributions in the continuous domain. They can be used for obtaining consistent predictions by modeling spatial correlation. Furthermore, exact inference and learning are possible using GPs when data are collected in the field. Observations from a sensor network can be processed by a GP-model to reduce the uncertainty and guide decision making. In this paper, we investigate the applicability of a metric based on mutual information to assess the benefits of alternative topologies for sensor networks. We develop a methodology to optimize the sensor placement for the management of infrastructure. We apply the proposed methodology to a network of assets under seismic risk, considering sensor systems measuring ground acceleration or providing information about the structural capacity, for assessing the reliability of the infrastructure after an earthquake.

KEYWORDS

civil infrastructure systems, gaussian processes, sensor networks, structural health monitoring.

INTRODUCTION

Structural Health Monitoring (SHM) systems can provide real-time information about the condition of both new and existing infrastructure assets, in terms of structural capacities and demands placed on them. An interesting problem for civil infrastructure planners and managers is how best to deploy such systems, with a view to optimally determining not only the status of individual assets, but the status of an interconnected and interdependent system.

In this paper, we make use of the specific context of a seismic event impacting an infrastructure network to evaluate several methods for optimally designing a sensor network to monitor this system. To model the relevant variables, a Gaussian Process (GP) framework is used. A GP is the generalization of a multivariate Gaussian distribution to a continuous domain. By providing both mean and covariance functions, the GP can be used to define a multivariate normal distribution for any set of points over the problem domain. For this reason, GPs have been used to model spatially distributed physical quantities, such as temperature or soil permeability. GPs can also be updated based on exact or uncertain observations of these physical quantities. Further background on GPs and their properties can be found in the textbooks of Rasmussen and Williams (2006) or of Barber (2010).

An investigation into optimal sensor placement using GPs has been conducted by Krause (2008) and Krause *et al.* (2008). They propose the maximization of mutual information between sensed and un-sensed locations within the domain of the process to be the criterion to guide sensor placement. To perform this optimization, they employ a simple greedy algorithm which iteratively adds the sensor locations with highest mutual information. This algorithm allows for rapid computation and provides a provably near-optimal solution, even for very complicated problems. These techniques were applied to problems of positioning temperature and precipitation sensors.

The problem of optimal sensing can also be viewed from a data mining and machine learning perspective as the problem of optimal feature selection for classification or regression (Hastie *et al.*, 2009). Given a large number of potential explanatory features for a phenomenon, it is often of interest to select a small number of features which are highly correlated with the phenomenon and provide the best basis for a predictive model. The optimal locations for the sensors are those corresponding to the selected features. If a model is available for system simulations, it can be used to populate a synthetic data-set, upon which feature selection techniques can be applied.

For the purposes of this paper, the specific domain of seismic risk is chosen as a setting for the demonstration and validation of sensor positioning techniques. Seismic demand variables, such as peak ground acceleration (PGA), are often modelled as log-normally distributed random variables. By applying a logarithm transformation to these, a set of Gaussian-distributed variables is obtained, which is amenable to modelling through GPs. Seismic demand, e.g. in the form of PGA, is also easily measured with acceleration sensors. Seismic capacity is more difficult to model and measure; however, assumptions have been made for the purpose of this research, as outlined below, which allow them to be modelled as log-normally distributed variables as well.

PROBLEM STATEMENT

The basic problem to be addressed can be stated as: "Given a network of distributed infrastructure assets, which are to be subjected to some uncertain global demand and have intrinsic capacities to resist this demand, how might sensors be best distributed over this network to provide the best assessment of the condition of this system as a whole?" To address this problem, a model of the system in question must first be developed. As stated above, we focus our attention on seismic risk. To define the corresponding scenarios, a model for seismic demand over a spatial domain, a model for the seismic capacity of infrastructure assets, and a model of the performance of the network of infrastructure assets are developed.

Demand Model

We assume that the seismic demand can be described by the PGA. Among many models available in the literature, depending on the seismic zone, we use that calibrated by Mandal *et al.* (2009), as reported by Douglas (2011). The model takes into account magnitude, distance from epicenter, and certain soil type characteristics, and includes additive Gaussian error. The basic equation of this model is:

$$\ln(PGA) = a + bM - \ln\left(\sqrt{r^2 + c^2}\right) + dS + \varepsilon\,; \quad \varepsilon \sim \mathcal{N}(0, \sigma) \tag{1}$$

where the PGA is evaluated at a specific location, M represents the moment magnitude of the earthquake, r is distance to the epicentre (in kilometres), and S is a soil type factor. The natural logarithm is denoted by $\ln(\bullet)$, and a, b, c, and d are empirical constants calibrated for this equation. The error factor ε describes a normally distributed random error in the prediction, with zero mean value and standard deviation σ, which has also been calibrated for this equation. Together, these define a log-normal distribution for the PGA at any distance from the epicentre. This ground motion attenuation equation is typical of many similar models, in that it combines a deterministic median prediction for ground motion with a stochastic residual error term. In many equations, this residual is divided into inter-event and intra-event residuals, to capture the variability from one seismic event to the next or the variability in motions within an event, respectively. In Eq. 1, the residuals are not distinguished, and the combined term is treated as an intra-event residual.

To define a GP model for PGA, it was assumed that a positive correlation exists between the values of PGA at nearby locations. As in the work by Bensi *et al.* (2011), the value of this correlation was determined by a squared exponential kernel function:

$$\rho_{ij} = \exp\left(-\frac{\left\|\Delta x_{ij}\right\|^2}{\lambda^2}\right) \tag{2}$$

where ρ_{ij} represents the correlation factor between locations i and j, while the distance between these locations is designated Δx_{ij}. The scale-length λ designates how quickly this correlation decays, and has been set to 10 km for this model. Overall, this model indicates a strong positive correlation for nearby points, and practically no correlation for distant points ($\Delta x_{ij} \gg 10$ km). Together, Eqs. 1 and 2 define a GP model for PGA over a domain, assuming basic parameters of the earthquake and soil type are known.

Despite the fact that PGA values may be a poor predictor of seismic demand for low-frequency components (Grigoriu and Kafali, 2004), we have selected this parameter for the sake of model simplicity, and we assume in the following that the natural frequency of the components are sufficiently high that the PGA can be taken as a relevant engineering parameter for the demand.

Capacity Model

We also define seismic capacity of infrastructure components in terms of acceleration. This model is based on seismic fragility curves presented in the manual for Hazus-MH software developed by the US Federal Emergency Management Agency (FEMA). These curves relate the probability of various damage levels on generic infrastructure components (i.e. pumping stations, electrical substations, storage tanks, etc.) to the applied seismic acceleration. For this model a "moderate" damage level is considered to represent a damaged infrastructure component, since this level of damage is consistent with the component being offline for a about week or more, meaning it would be unavailable during immediate post-event recovery efforts. The fragility curves presented are based on log-normal cumulative distribution functions, and so these curves can be easily adapted to describe log-normal probability distributions over the capacity of these assets in terms of their resistance to PGA loadings.

We also define correlations between capacity variables for different assets. The underlying concept is that assets with similar construction materials, designs, and functionalities would have positively correlated capacities. No prior investigations into values for these correlations were found, and so, for the purposes of this model, the correlations between capacities of assets of different types have to be assigned based on expert knowledge. In principle, the correlation should reflect the similarity of structural capacities, as perceived by the model designer. In the absence of such a systematic approach, however, we define a multivariate log-normal distribution for the set of assets using the capacities derived from the Hazus-MH software manual as marginal distributions, combined with assigned correlation factors.

System Model

With log-normal models for capacity and demand, the natural logarithm transformation can be used to convert these to Gaussian distributions which can be modelled as GPs. Next, a basic technique of risk assessment is used to define a limit state function as difference between a capacity and demand variable. When demand exceeds capacity, the difference will be negative, and the limit state is said to be exceeded, indicating that the asset is in a damaged or failed state. The probability of the asset being in this state can then be found by evaluating the cumulative probability distribution of the limit state function up to the boundary of the failure domain. This function can also be mapped to a binary indicator of the state of an asset (i.e. 0 for a failed system, 1 otherwise).

The functionality of an infrastructure system can be defined in terms of its ability to provide some basic public service (i.e. water supply, power supply, transportation connectivity). If the system fails in providing this service, the system can be said to have failed. For such systems, a fault tree analysis can be performed, indicating which asset failures or combinations of asset failures might lead to a breakdown of the system as a whole. The results of this analysis can be used to define a network model in which assets are arranged in parallel or series combination, as they relate functionally to each other and make up the system. A mathematical model can be defined for the state of the system in terms of the states of each asset. This can be presented as a function of the binary indicators for the state of each asset, which determines a binary indicator of the system state (S). The minimum cut-sets representation can be used to define the failure domain of the system. The probability of system failure can then be computed by integrating over this failure domain. The reader is referred to Bensi *et al.* (2011) for a more complete overview of the subject.

PROPOSED METHODOLOGY

We propose to model a distributed network of infrastructure assets using GPs for the logarithms of both capacity and demand variables, as outlined above. Optimal positioning of sensors on this network can then be carried out by maximizing the mutual information between sensor measurements and variables of interest. A data-driven approach, where informative features are selected from simulated data sets provided for potential sensor locations, is also investigated.

Approach 1: Model Based Approach for Optimal Sensor Placements to Predict Demand or Capacity

The properties of GPs allow for the effect of observations made upon the field to be used in updating the joint probability of all variables. Following the work of Krause (2008), we use a greedy optimization algorithm which sequentially selects sensors providing the highest mutual information between sensed and un-sensed locations in the system. The goal for this optimization procedure is to find A^* so that:

$$A^* = \text{argmax}_{A:|A|=k}\, I(X_A; X_{V \setminus A}) \tag{3}$$

where A denotes the set of sensed locations (that is, locations at which sensors can measure some variable), and X_A denotes the random variables which these sensors measure. The entire domain of the problem is denoted as V; thus, $V\backslash A$ is the set of locations in the domain where sensors are not positioned. For practical implementation, the domain V is discretized to a finite number of potential sensor locations, making $X_{V\backslash A}$ the joint set of variables which are not measured. The operators $|\bullet|$ and $I(\bullet;\bullet)$ denote the size of a set and the mutual information operators, respectively. The reader can refer to Cover and Thomas (2006) for more information on mutual information and entropy calculations. Therefore, Eq. 3 expresses the goal of choosing the set of locations A of size k which maximizes the mutual information between sets of measured variables X_A and unmeasured variables $X_{V\backslash A}$.

To implement the greedy optimization algorithm, the potential gain in mutual information for each proposed sensor location is computed, and the location resulting in the largest gain is included into the set of sensor locations. This process is described as:

$$y^{*}=\operatorname{argmax}_{y\in V\backslash A}\, Gain(y) \tag{4}$$

where y is the proposed sensor location selected from the set of un-sensed locations $V\backslash A$, and $Gain(y)$ is presented in Eq. 5. The proposed location which maximizes the gain in mutual information will be added to the set of proposed sensor locations A before the next iteration of the greedy algorithm optimization procedure. $Gain(y)$ is computed as:

$$Gain(y) = I\left(X_{A\cup y};X_{V\backslash A\cup y}\right) - I\left(X_A;X_{V\backslash A}\right) = H\left(X_y\middle|X_A\right) - H\left(X_y\middle|X_{V\backslash A\cup y}\right) \tag{5}$$

where $H(\bullet)$ denotes the entropy operator, $H(\bullet\mid\bullet)$ denotes conditional entropy, $A\cup y$ denotes the set of proposed sensor locations including the new proposed location y, and $V\backslash A\cup y$ denotes the domain of the problem excluding all proposed sensor locations including the newly proposed location y.

In the context of the problem outlined in the previous section, where GPs are defined over the logarithms of log-normally distributed variables, this gain can be computed explicitly. The formula for the entropy of log-normally distributed random variables is given as:

$$H\left(X_y\right)=\frac{1}{2}+\frac{1}{2}\ln\left(2\pi\zeta_{X_y}^2\right)+\lambda_{X_y} \tag{6}$$

where $\zeta_{X_y}^2$ and λ_{X_y} represent the variance and mean of the associated normal distribution for $\ln(X_y)$, respectively. The conditional entropy can therefore be expressed as:

$$H\left(X_y|X_A\right)=\frac{1}{2}+\frac{1}{2}\ln\left(2\pi\zeta_{X_y|X_A}^2\right)+\mathbb{E}_{X_A}\left[\lambda_{X_y|X_A}\right] \tag{7}$$

The parameters in Eq. 7 can be determined via Bayesian updating, as:

$$\zeta_{X_y|X_A}^2=\zeta_{X_y}^2-\boldsymbol{\Sigma}_{X_y,X_A}\boldsymbol{\Sigma}_{X_A,X_A}^{-1}\boldsymbol{\Sigma}_{X_A,X_y}$$

$$\lambda_{X_y|X_A}=\lambda_{X_y}+\boldsymbol{\Sigma}_{X_y,X_A}\boldsymbol{\Sigma}_{X_A,X_A}^{-1}\left(\ln(X_A)-\lambda_{X_A}\right) \tag{8}$$

where $\boldsymbol{\Sigma}_{X_A,X_A}$ represents a covariance matrix for X_A and $\boldsymbol{\Sigma}_{X_y,X_A}$ represents a covariance vector between variable X_y and variables X_A. Using the definition for $\lambda_{X_y|X_A}$, we can compute the expected value over possible values of observations X_A. It can be shown that:

$$\mathbb{E}_{X_A}\left[\lambda_{X_y|X_A}\right]=\lambda_{X_y} \tag{9}$$

Using this result, Eq. 7 can be rewritten as:

$$H\left(X_y|X_A\right)=\frac{1}{2}+\frac{1}{2}\ln\left(2\pi\zeta_{X_y|X_A}^2\right)+\lambda_{X_y} \tag{10}$$

The derivation for $H\left(X_y\middle|X_{V\backslash A\cup y}\right)$ is carried out in a similar fashion. Substituting these expressions into Eq. 5, we can obtain an expression for the gain in mutual information for a proposed sensed variable y:

$$Gain(y)=\frac{1}{2}\ln\left(\frac{\zeta_{X_y|X_A}^2}{\zeta_{X_y|X_{V\backslash A\cup y}}^2}\right) \tag{11}$$

Since the logarithm is monotonic, we can merely evaluate the argument for the logarithm in Eq. 11 to compute and rank the gains in mutual information for each potentially observed variable y.

This algorithm for optimal sensor placement in GPs can be employed in a number of ways. To optimally position sensors for demand measurement over the entire domain of a model, this domain can be discretized to a grid of potential sensor locations. A given number of acceleration sensors can be positioned over the grid to maximize mutual information between the acceleration demand variables at sensed locations (the chosen locations for the sensors) and the rest of the domain (i.e. every potential sensor location over the discretized domain where no sensor has yet been positioned). In the context of infrastructure system assessment, only the accelerations at infrastructure component locations will have an effect on the system. In this case, the domain of interest (V) can be limited to include only the asset locations, while sensors might still be placed at any point on the grid discretization of the domain. The sensor locations are optimized in terms of gain in mutual information between sensor locations and asset demand variables.

This approach can also be used on the capacity model used for the components. We assume we can refine the assessment of the structural capacity, for example by using a structural health monitoring effort underway upon an asset (such efforts are labelled as "Capacity Measures" in the figures in the following section). Using the correlation between capacities, an optimal scheme for conducting these capacity measurements might be developed, as for demand variables. In this case, the domain is limited to the set of assets included in the system, and mutual information is measured between the monitored assets and the unmonitored assets in the system. This has the effect of prompting the algorithm to measure components whose capacities are most correlated with those of other unmeasured components.

Finally, since the limit state variables for each component are themselves defined by multivariate Gaussian distributions, this same approach can be applied to these variables as well. Considering both capacity and demand measurements together, the algorithm can be used to select the measurements with the greatest mutual information with the limit state variables.

Approach 2: Model Based Approach for Optimal Sensor Placements to Predict System Status

From a system performance perspective, however, the true variable of interest is the functionality of the system, S. This variable has only two possible states; either the system can be functioning after the earthquake or it can be failed. The probability of being in one state or the other is defined by the system probability of failure. S is therefore in the form of a Bernoulli binary random variable. We can seek a method to optimally place sensors throughout the system so as to gain information about S in an efficient manner.

For S, the mutual information gain criterion described in Eq.5 collapses to a simple entropy gain condition:

$$Gain(y) = I\left(X_{A\cup y};S\right) - I(X_A;S) = H(S|X_A) - H(S|X_{A\cup y}) \quad (12)$$

In this case, the approach simply selects the observations which most reduce the conditional entropy of S. Being a Bernoulli variable, its entropy is:

$$H(S) = -P(S=0)\ln[P(S=0)] - P(S=1)\ln[P(S=1)] \quad (13)$$

The conditional entropy of S given a set of observations X_A is:

$$H(S|X_A) = \int_{X_A}\{-P(S=0|X_A)\log[P(S=0|X_A)] - P(S=1|X_A)\log[P(S=1|X_A)]\}p(X_A)dX_A \quad (14)$$

where $p(X_A)$ describes the multivariate probability density. Evaluation of Eq.14 is computationally demanding. In this paper, the integral is approximated through Monte Carlo simulations.

Approach 3: Data Driven Approach for Optimal Sensor Placements to Determine System Status

We also propose a brute-force method for sensor placement, based on simulations. By running multiple simulations for all variables in the problem, a set of sensor readings for each proposed location is generated. Then, sets of these features (i.e. simulated sensor readings) are used to train classifiers for predicting the system state. The set of features resulting in the lowest misclassification rate is selected and cross-validated over different data sets. In this paper, a simple approach to feature selection is used based on a Support Vector Machine (SVM) classifier. We adopt SVM for its robustness in binary classification and its computational efficiency. The reader is referred to Hastie *et al.* (2009) for more information on SVM.

NUMERICAL VALIDATION

For implementation of the methods described in the previous section, we make use of the network of infrastructure elements presented in Figure 1. This network for a water supply system is arranged spatially over a 40×40 km^2 area. Assets in this system include an electrical generator station (Gen), a power substation (Sub), two power distribution grids (Gd1& Gd2, represented as point assets), a medium and small pumping station (MPP & SPP), a water treatment plant (WTP), and a water storage tank (Tnk). A system failure is considered to have occurred if either the substation and generator elements fail, both discretized grid elements fail, both pumping plants fail, or the water treatment plant and storage tank assets fail.

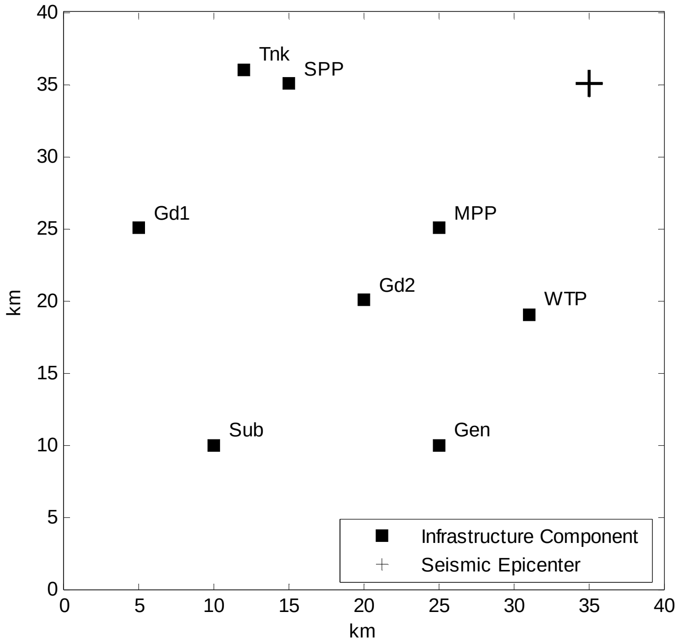

Figure 1. Arrangement of assets in the infrastructure network

Table 1. Mean values and coefficients of variation for asset capacity variables

Asset name and label	Mean value of ln(capacity/g)	Coefficient of variation of ln(capacity/g)
Medium Pumping Plant (MPP)	-1.273	0.25
Small Pumping Plant (SPP)	-1.273	0.25
Electric Substation (Sub)	-1.6094	0.25
Electric Generator (Gen)	-1.7720	0.25
Power Distribution Grid 1 (Gd1)	-1.1087	0.04
Power Distribution Grid 2 (Gd2)	-1.1087	0.04
Water Storage Tank (Tnk)	-0.8675	0.49
Water Treatment Plant (WTP)	-1.0498	0.16

Demands on the system are defined in terms of the PGA model presented in Eqs.1 and 2, based on an earthquake with Magnitude 6 and known epicentre location. Capacities are defined for each component, as shown in Table 1, with non-zero correlations defined between the capacity variables of the two grid elements (ρ=0.25) and those of the two pumping plant assets (ρ=0.4).

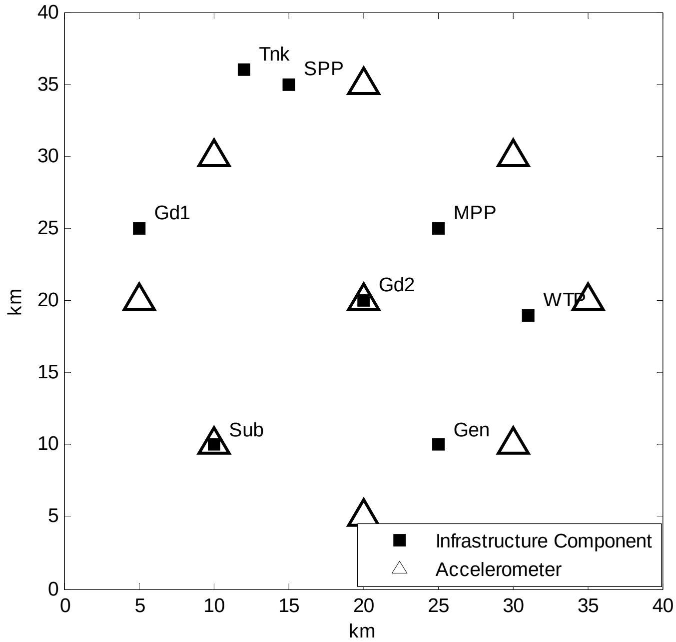

Figure 2. Results of acceleration sensor positioning for demand determination over the entire domain.

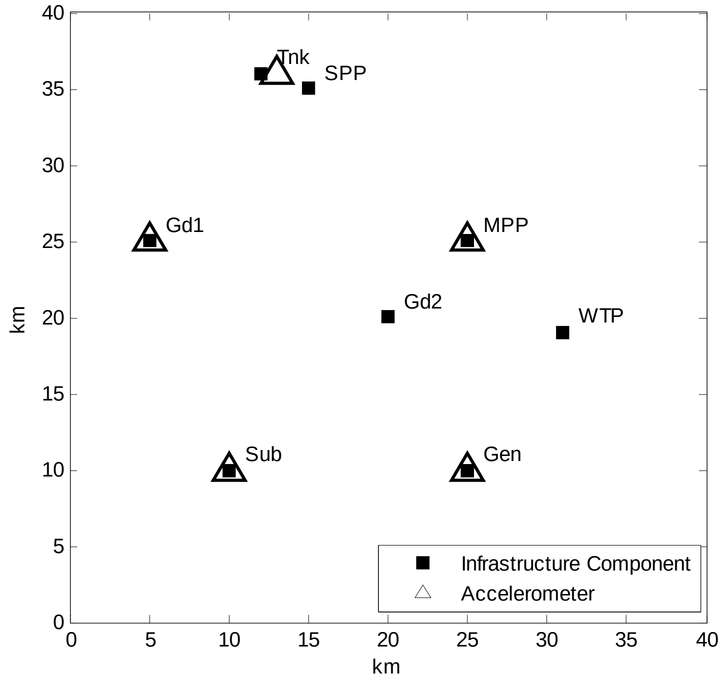

Figure 3. Results of acceleration sensor positioning for demand determination at asset locations only

Results for optimal sensor positioning to determine demand are shown in Figures 2 and 3, with acceleration sensors denoted by triangular symbols. Results for sensor positioning of nine sensors over the entire domain, shown in Figure 2, are fairly intuitive, exhibiting a symmetrical distribution over the domain so as to optimally "cover" the entire area. When the domain of interest is restricted to the locations of infrastructure assets only, as shown in Figure 3 for five sensors, the algorithm tends to locate sensors at asset locations, as these are clearly the most directly informative about the acceleration demand on assets. Only when assets are fairly close spatially is there a benefit to placing a sensor between assets, so as to gain information about both. Similar results were observed when the algorithm was applied to optimally place sensors to measure capacity, although in this case, such sensors can only be placed at asset locations. For this problem, the placed sensors took advantage of the correlations between capacity variables, monitoring only one element of each pair of positively correlated grid or pump type assets in the network.

Combined placement of acceleration and capacity measurements for system reliability assessment is approached in two ways, and the results are presented below in Figures 4 and 5. The domain was discretized to a 5×5 km^2 grid over which acceleration sensors might be placed.

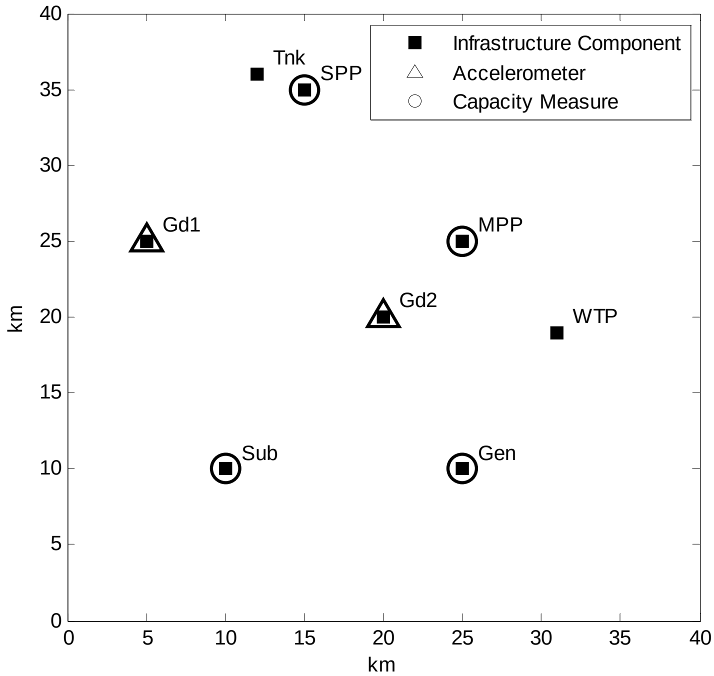

Figure 4. Positioning of sensors with highest mutual information with limit state variables

In Figure 4, the greedy algorithm is applied to maximize the mutual information between observed variables and the joint set of component limit state variables. While the uncertainties in these variables themselves are not direct indicators of the uncertainty in the state of each asset, that state is a function of these variables. Furthermore, these variables have a joint Gaussian distribution, allowing for the easy application of the previous approaches for sensor placement to this problem. As shown in Figure 4, six measurements are placed over the network, with capacity measurements placed on the small and medium pump assets and the substation and generator assets, and demand measurements on both grid assets. The two acceleration sensors are placed first by the algorithm. With these two sensors in place, the uncertainty in acceleration over the asset locations is reduced significantly, such that uncertainty in capacity is the dominant factor in the uncertainty of the limit state variables. The algorithm next places the four capacity sensors at four of the five assets with the most uncertain capacities. This placement of sensors most reduced the uncertainty in the limit state variables. However, whether these variables are closer to or farther from the failure domain of the asset had no bearing on this sensor selection, so it is not optimized in terms of focusing on the most at-risk assets first.

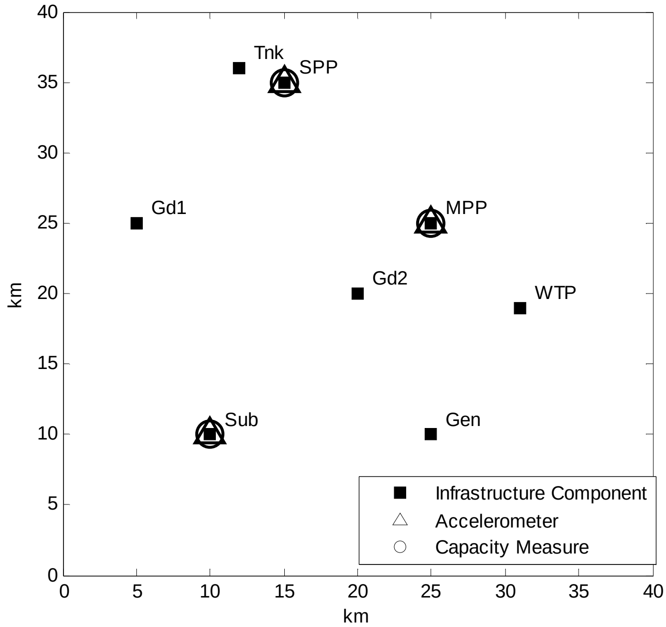

Figure 5. Sensor placement for system reliability assessment via system state variable entropy reduction

In Figure 5, sensors are placed in order to minimize the conditional entropy of the system state variable S. As discussed previously, this is a more complicated problem, as this variable is no longer Gaussian, and so an approximation method is used to compute posterior entropies in order to produce these results. In Figure 5, the six demand and capacity measurements are placed on only three assets (the small and medium pumping plants and the substation). The two pumping plants might be considered the most at-risk assets, as they are closest to the epicentre, and are redundant, so at least one of the two must remain intact for the system to operate. The water treatment plant and tank are slightly farther away, and have slightly greater safety margins from their failure domains.

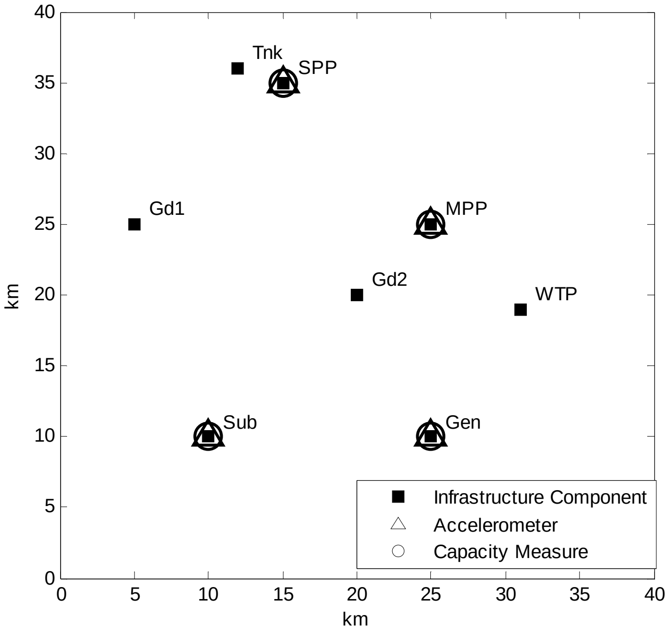

Figure 6. Results of feature selection method (SVM)

For comparative purposes, the results of a feature selection approach to sensor positioning are presented in Figure 6, above. For this approach, a series of twenty thousand simulations of the system were conducted based on the system model, with values simulated for every potentially sensed variable, as well as the state of the system. A SVM classifier is then trained on these simulations to predict the state of the system using potential sensor observations as features. The eight most significant features for system state classification are selected, and the sensors corresponding to these features (that is, the sensors which would be used to collect them) are selected as the optimal sensor positions for the network, based on this approach. According to Figure 6, the most relevant features to the system status are the accelerations and capacities of the pumping plants, the electrical generator, and substation. This agrees with the previous results, with an overlapping of the features selected. The two methods do not necessarily rank the features in the same order, since it required the selection of eight features via the SVM method to obtain, as a subset of this selection, the six locations selected in Figure 5.

CONCLUSIONS

In this paper, we propose methods for optimal sensor placement over a network of interrelated and interdependent infrastructure components. The method involving minimizing the entropy of S is computationally expensive because the dimension of the domain of integration in Eq. 14 grows linearly with k, that is the number of sensors we are placing. The Monte Carlo method we use to solve Eq. 14 requires more and more samples to obtain a constant accuracy with increasing k. In the numerical example of the previous section, we use one hundred simulations, and the order of selection for the sensors is not always the same, but is sensitive to the randomness within the simulations. A major advantage of this method is that it exploits the particular structure of the GP, and it can easily support changes in the model of the system resulting from new information.

The running time of the SVM method, including generating the necessary training data, is comparable with the previous one but, as noted above, it is a more general method, as it does not rely on the GP structure. Any change to the model, furthermore, will require that an entirely new data set be generated.

It should be kept in mind that all presented results apply only to a specific scenario, i.e. an earthquake with given magnitude and location. In future implementations, this approach will be extended to the case where these variables are given by a probabilistic distribution over multiple scenarios. Furthermore, future implementations will seek to quantify the value of information provided by sensor placements through the inclusion of a well-defined decision making problem on the post-event management of the infrastructure system. In this context, where decisions must be made based on available information, the monetary or utilitarian value of the information provided by sensors to decision makers can be computed.

REFERENCES

Barber. D. (2010). *Bayesian Reasoning and Machine Learning*, Cambridge University Press, UK.

Bensi, M., Der Kiureghian, A. and Straub, D. (2011). *A Bayesian Network Methodology for Infrastructure Seismic Risk Assessment and Decision Support*, PEER Report 2011/02.

Cover, T.M., and Thomas, J.A. (2006). *Elements of Information Theory*, John Wiley & Sons, Chichester, UK

Pozzi, M., Der Kiureghian, A., Yue, Y., and Zonta, D. (2012). "Bayesian network with Gaussian variables for post-earthquake emergency management", *Proceedings of the 16th IFIP WG 7.5 Working Conference on Reliability and Optimization of Structural Systems*, Yerava, Armenia, June.

Douglas, J. (2011). *Ground-Motion Prediction Equations, 1964-2010,* Pacific Earthquake Engineering Research Center, College of Engineering, University of California, Berkeley, USA.

Grigoriu, M. and Kafali, C. (2004). "Seismic fragility analysis", *Proceedings of the Ninth ASCE Joint Specialty Conference on Probabilistic Mechanics and Structural Reliability*. Albuquerque, New Mexico, July.

Hastie, T., Tibshirani, R., and Friedman, J. (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer Science and Business Media, New York, USA.

Hazus-MH 2.1 Technical Manual: Earthquake Model – Multi-hazard Loss Estimation Methodology, Department of Homeland Security, Federal Emergency Management Agency, Washington DC, USA.

Krause, A. (2008). *Optimizing Sensing: Theory and Applications*, PhD Thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, USA.

Krause, A., Singh, A. and Guestrin, C. (2008). "Near-optimal sensor placements in Gaussian processes: theory, efficient algorithms and empirical studies", *Journal of Machine Learning Research,* 9, 235-284.

Rasmussen, C.E. and Williams C.K.I. (2006). *Gaussian Processes for Machine Learning*, M.I.T. Press.

Zografos, K. and Nadarajah, S. (2005). "Expressions for Rényi and Shannon Entropies for multivariate distributions." *Statistics and Probability Letters* 71, 71-84.