4th International Workshop on Structural Control, Columbia University, New York, June 2004.

Networked Sensing for Structural Health Monitoring

John Caffrey, Ramesh Govindan, Erik Johnson, Bhaskar Krishnamachari, Sami Masri, Gaurav Sukhatme, Krishna Chintalapudi, Karthik Dantu, Sumit Rangwala, Avinash Sridharan, Ning Xu, Marco Zuniga

ABSTRACT

This paper describes an ongoing project investigating embedded networked sensing for structural health monitoring applications. The vision is of many low-power sensor “motes” embedded throughout the structure with a smaller number of nodes that can provide local excitation. The challenge is to develop both the networking algorithms to reliably communicate within the network, and distributed algorithms to monitor the state of the structure. A wireless data acquisition network is described, including the methods of storing and transmitting the data. A damage detection scheme is described that uses extremely low transmission bandwidth, and is shown to be effective in detecting damage in a simulated structure. Finally, a large-scale structural testbed that is being used for this project is described. The outcome of this work-in-progress is expected to be strong recommendations and algorithms for distributed wireless sensor/actuator structural health monitoring networks.

INTRODUCTION

The emergence of wireless, battery-operated, small form-factor computing devices containing onboard MEMS sensors has enabled *embedded networked sensing*. This technology permits the placement of a relatively large number of such devices at fine spatial scales (on the order of a few meters). These devices can measure physical phenomena, and locally store or process readings. In addition, the devices self-organize into a computer network that can be used to collaboratively infer characteristics of a measured phenomenon. This technology has the potential to provide fine-grain measurements in a variety of natural settings such as terrestrial and marine ecosystems.

To what extent is such technology useful for structural health monitoring (SHM)? Most SHM research to date has focused on limited independent local damage detection mechanisms or on global damage assessment techniques using low resolution measurements of a structure’s vibration response to ambient excitation. To some extent, local SHM methodologies have focused on devices that have limited

Viterbi School of Engineering, University of Southern California, Los Angeles CA 90089.

range, detecting very local changes generally at ultrahigh frequency levels. While useful, such approaches do not use multiple sensor devices to collaboratively provide superior assessments of structural health. Motivated by the emergence of embedded networked sensing, we envision an approach to SHM that uses decentralized local excitation (using smaller actuators capable of exciting parts of a structure) and high resolution measurements of response to these excitations, detected and analyzed through a wireless network of devices.

This approach is attractive because a wireless network can significantly reduce installation and maintenance cost of a structural monitoring system; wiring large structures for high resolution measurements can be prohibitively expensive. Furthermore, since embedded networked sensing enables highly dense measurements, it promises potentially simpler and more accurate techniques to identify and even localize damage within the structure. Further, this approach promises a future where, for example, buildings are constructed using concrete mixed with several tens of thousands of embedded sensor devices as well as (a possibly smaller number of) low power local exciters. This network of sensors will be able to continuously monitor a structure, trigger alarms that identify the onset of damage, precisely pinpoint the location of damage and also provide a long-term history of ambient stresses imposed on the structure.

Of course, much work is needed before this vision becomes reality. Advances must be made in understanding how networked sensing applies to structural engineering. In addition, given that this is a nascent research area, research infrastructure must be developed to support this interdisciplinary endeavor. This paper summarize advances that have been made on these two fronts. Specifically, towards the former goal of understanding how and in what way networked sensing can be used for structural monitoring, a structural data acquisition system, denoted Wisden, has been prototyped using a network of wireless sensors to collect structural measurements; this prototype system is being used to investigate the feasibility of structural damage detection using wireless sensor devices together with an excitation at multiple locations. Progress has also been made towards developing a research infrastructure that will support the development of networked sensing for structures. A network simulator has been integrated with a structural model, enabling more rapid design and evaluation of SHM techniques; *mobile* exciters are currently being designed for a large testbed structure. The following sections describe these pieces of work.

A WIRELESS DATA ACQUISITION SYSTEM FOR STRUCTURAL MONITORING

The first use of networked sensing for structural monitoring is likely to be as a *data acquisition system* that collects structural measurements from multiple locations at a single node for centralized processing. A wireless sensor network system, called Wisden (short for Wireless Structural Data Extraction Network), has been developed for structural data acquisition. Wisden continuously collects structural response data from a multihop network of sensor nodes, and displays and stores the data at a base station. While the architecture of Wisden is simple—a base station centrally collecting data—its design is challenging given the requirements of structural sensing: relatively high data rates, loss-intolerant data transmission, and time-synchronized readings from sensors. The relatively low radio bandwidths, the high wireless packet loss rates observed in many environments, and the resource constraints of existing sensor platforms add significant challenges to the design of Wisden.

Wisden uses mostly off-the-shelf hardware. Specifically, the sensor nodes are each a Mica2 “mote”. Existing sensor platforms do not, however, have hardware support for high quality vibration sensing. So, Wisden uses a 16-bit *vibration card* originally designed for high frequency (up to 20 kHz), sampling at 16 bits per sample. The authors modified the card’s firmware to sample up to three channels of acceleration data at 100 Hz.

The first challenge in Wisden is to reliably transmit data from each sensor to a base station. Each source stores generated vibration data in its EEPROM, and transmits the data to its “parent” (the network node closer to the base station than itself).

Parents keep track of sequence numbers of packets that they receive, on a per source basis. A gap in the sequence number of sent packets indicates packet loss. Each node maintains a list of missing packets. When a loss is detected, a tuple containing a source ID and sequence number of the lost packet is inserted into this list. Entries in the “missing packets” list are piggybacked in outgoing transmissions, and children infer losses by overhearing this transmission. Nodes keep a small cache of recently transmitted packets, from which a child can repair losses reported by its parent. The base station also recovers packets using this mechanism, which is necessary to deal with situations where network routing paths may change due to node failure.

The data rate requirements for structural monitoring can be a significant fraction of radio bandwidth. Thus, *data compression* is a second crucial component of Wisden. While most prior research ([2]) has focused on data aggregation in order to increase network lifetime, the primary motivation for considering compression is to scale Wisden to many nodes. Of course, using higher bandwidth 802.11 radios represent a possible solution to this problem. However, platforms employing such radios typically consume an order of magnitude more power.

One possible approach to data compression is *event detection*. This approach is based on the observation that, if samples within a small window have a low value *and* are comparable in value, the structure is quiescent. Such quiescent periods are compressed using run length encoding; samples in non-quiescent periods are transmitted without compression. Event detection suppresses data transmission when events do not occur. Thus, the overall data rate required to transmit the samples is a function of the duty cycle of the vibrations. The current implementation of Wisden uses this approach.

A more general strategy is to use *progressive storage and transmission* that stores vibration data locally and transmits a lossy version (using wavelet compression) of the data to the base station. Such an approach enables low latency but lossy data acquisition. The stored data allows detailed views of the vibration data to be retrieved on demand. This technique will be useful in platforms that have significant local storage, a trend that is likely given the falling prices of flash memory. The use of such approaches is currently being examined for the Wisden prototype.

A third challenge in the design of Wisden was accurate time-stamping of sensor measurements (the use of GPS at each sensor node is precluded within structures, of course). In Wisden, each node calculates the amount of time spent by a sample at that particular node using its local clock. This amount is added to an *residence time* field attached to a packet (for simplicity, Wisden associates offsets with the first sample in a packet), as the packet leaves the node. Thus, the delay from the time of generation of the sample to the time it is received by the base station (or any node) is stored in the packet as the sample travels through different nodes in the network. This is the time the packet resides in the network. The base station (or any node) can thus calculate the time of generation of the sample by subtracting the residence time from its local time. If the base station is GPS synchronized, this approach gives a good approximation. If the residence time field is updated as close to the radio and the accelerometers as possible, then, assuming packet propagation times are negligible in dense sensor deployments, this approach can successfully timestamp the sample.

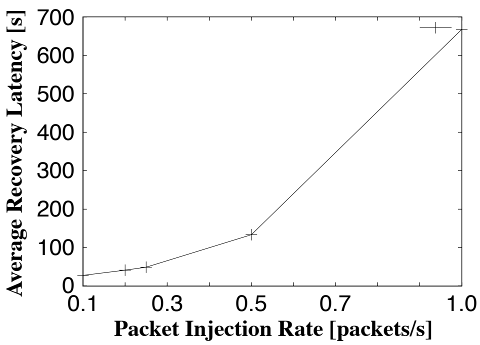

Figure 1. Average Recovery Latency

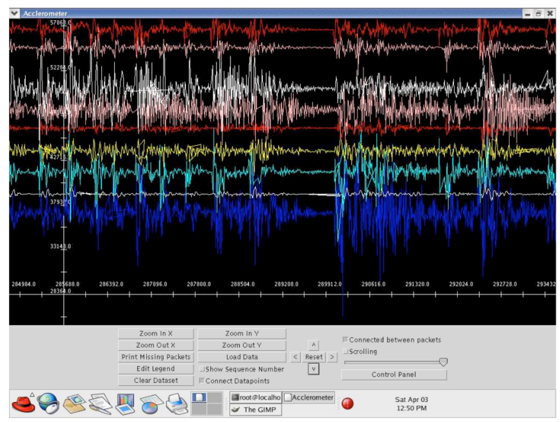

Figure 2. Data collected from the ceiling structure.

A 10-node Wisden system was deployed on the test bed described in a later section. The structure was instrumented by affixing accelerometers with heavy-duty double-sided tape, and wrapping the rest of the assembly with gaffer tape. The structure was then repeatedly hit with a 2-by-4 for 20 seconds. Figure 2 is a screenshot of the collected sample data, aligned at the base station. The 10 motes formed a multihop network and transmitted all of the recorded vibration data back to base station within 5 minutes. The average residence time incurred by a packet in this experiment was 142 seconds; some of the delay can be attributed to the sustained excitation, and some to packet loss. Finally, the *onset* time of the forced vibrations was within one sample time (actually 8ms) across all accelerometers, which is an indication that the time synchronization scheme was performing well.

DAMAGE DETECTION USING SENSOR NETWORKS

Systems that detect and locate damage in large structures such as buildings, bridges, ships and aircraft can improve safety and reduce their maintenance costs. This work takes a preliminary step in this direction. The goal is to get some understanding of the following questions: Can SHM techniques be implemented on a low-power wireless sensor network? Specifically, does there exist a solution to an SHM problem that admits of an implementation on a long-lived network of wireless sensor nodes? As a first step, the focus is on the most basic SHM problem, that of damage *detection*, *i.e.*, merely determining whether a structure is damaged.

Many existing SHM techniques attempt to detect damage by continuously recording structural response to ambient vibration. Such an approach might be infeasible for a wireless sensor network designed for long-lived operation, because the energy consumption of even low-power accelerometers is significant. SHM techniques that rely on forced vibrations, such as those delivered by electrodynamic shakers, are better suited to implementation on wireless sensor networks. Wirelessly controlled shakers can precisely deliver forced excitations at predetermined locations in coordination with the sensors to improve the accuracy of structural damage detection. Given that the forced vibrations can be scheduled, nodes can be duty cycled in order to increase network lifetime. Thus, the sensor-actuator network will consist of a number of wireless sensor nodes together with perhaps a smaller number of exciters.

How can such a system detect damage in structures? Frequency shift methods are a well-known class of damage detection techniques that infer damage by analyzing the frequency response of a structure. In our system, sensor network nodes analyze the structure's frequency response and locally (*i.e.,* without exchanging voluminous sensor readings with other nodes) extract the modes. They then collaboratively decide whether there exists damage in the structure by detecting frequency shifts. Given that the number of (detectable) modes is small, the communication overhead of determining damage is minimal. The system uses standard signal processing techniques to filter noise and extract the modes from the frequency response. Such techniques can be implemented on an ARM or XScale based platform.

In this work, it is assumed that a collection of sensors and actuators are dispersed throughout a structure. Given the focus on trying to understand how networked sensing can be used for structural monitoring, two questions should be posed: (*a*) Can multiple sensors overcome some of the challenges faced by damage detection and, if so, how?; and (*b*) Can exciting the structure at different places fundamentally lead to an enhanced detection capability? The answers to these questions are both affirmative, and follows from the basics of structural dynamics, of course. Even though real structures have very complex mode shapes, *structural response spectra are location dependent*, and different locations will typically have a different set of modes as dominant and others recessive, depending on their mode shapes and the chosen location. In other words, detection of certain modes may be difficult in certain locations and easy in others. To answer the second question, the structural response spectrum at any location also depends on the location of excitation, since the latter can determine which modes are excited.

The proposed algorithm for damage detection uses these ideas and is divided into two phases: local and collaborative. In the local phase, each node gathers data and performs spectral analysis to create a local list of tuples $\left\langle f_i^n, e_i^n \right\rangle$. Here, f_i^n is the modal frequency of the i^{th} mode as discovered at the n^{th} node and e_i^n is the fraction of signal energy contained in that mode. In the collaborative phase, nodes collaboratively compute the global list of tuple $\left\langle f_i, \max_k e_i^k \right\rangle$. This can be computed by sending all the tuples to a designated node in the network that “aggregates” the readings and determines mode shifts.

For brevity, the details of the aggregation algorithm are omitted here, but it should be emphasized that this algorithm is amenable to implementation on energy-constrained devices. All devices need to periodically wakeup (say once a day), take measures, and perform the collaborative phase. Thus, the network can operate a very low duty cycles and, given today’s technology, network lifetimes on the order of a few months are feasible.

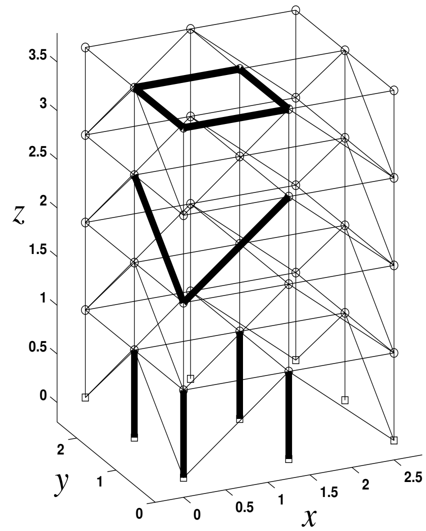

But does this algorithm really improve the efficacy of damage detection? In order to demonstrate the efficacy of the algorithm, it is evaluated using the structural model from the IASC-ASCE Task Group’s SHM benchmark model [3]. The benchmark structure is a 4-story, 2-bay by 2-bay steel-frame quarter-scale model structure in the Earthquake Engineering Research Laboratory at the University of British Columbia (UBC) [1]. 360 damage patterns were generated for the structure as follows. Each damage pattern corresponds to a partial loss of stiffness in a single element of the structure (a beam, a brace or a column). There are 9 columns, 8 braces and 12 beams per floor in the structure. Of these, 4 columns, 2 braces and 4 beams per floor were chosen (Figure). The fraction of loss of stiffness in an element was varied from 10-90% at intervals of 10%, leading to a total of 360 damage patterns.

Given the focus on understanding how sensor/actuator networks can help enhance damage detection, three different scenarios are evaluated:

- **s1** uses data from *best* sensor/actuator pair only,
- **s2** uses all the sensors and the *best* actuator, and
- **s3** uses the entire sensor/actuator network.

Comparing **s3** to **s1** helps quantify gains over using a single sensor/actuator pair, while comparing **s3** to **s2** provides insight into how much one can gain by using multiple actuators as opposed to a single one. To determine the best sensor/actuator pair for scheme **s1**, we find the pair that has the maximum number of spectral peaks having more than 5% of the signal energy. Ties were broken using the peak energies. For this structure, the best sensor/actuator pair is the actuator on floor 2 southeast column and the sensor on 4th floor southeast column. Similarly, to determine the best actuator position for scheme **s2**, the spectral peaks were aggregated over all 36 sensors for each of the four actuator locations, and the best actuator location is found to be the one on the third floor.

Figure 4 depicts the number of successful detections with percentage loss in stiffness of the members. As seen from Figure 4*a*, scheme **s3** significantly outperforms **s1** in the detection of braces. Even at only 30% loss of stiffness, all damage in braces is detected by **s2** and **s3** while **s1** is unable to detect any of them. For braces, there is no significant performance improvement when more than one actuator is used.

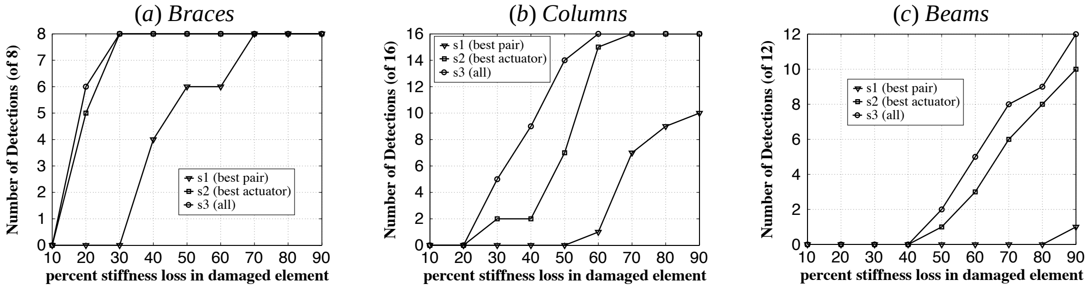

Figure 4. Performance comparison for actuator/sensor schemes as function of damage stiffness loss.

The need for multiple actuators becomes evident when detecting damage in columns. In Figure 4*b*, both schemes **s3** and **s2** significantly outperform **s1** at low values of loss of stiffness. For example, at 50% loss of stiffness, **s3** detected 14/16 column damage patterns while **s2** detected only 7/16 damage patterns. To summarize, having multiple sensors can dramatically improve detection capabilities. Having multiple actuators can further improve performance, depending on the type of element being considered.

A HYBRID SIMULATION FRAMEWORK

Within the structural engineering community, finite element models, as well as linearized and nonlinear control system models have been developed to provide rapid, yet accurate, structural analysis. Analogously, in the networking community, a variety of event-driven simulators have been designed that can recreate the execution of the actual distributed code in component nodes as well as the communication amongst them. For sensor networks, such existing simulators include NS2, TOSSIM, SensorSim, and Emstar. To make rapid progress, a critical piece of functionality is a *hybrid* simulator that integrates structural simulators with network simulators.

As a first step towards this goal, MATLAB® code, which uses the Control Toolbox®, is integrated with the TOSSIM network simulator. A state-space (*sys*) object model of the SHM Task Group benchmark structure [3], with input forces in *x*, *y* and *z* directions at each node and at the base (111 total inputs), and output absolute acceleration measurements in three dimensions at each node (108 total outputs), is used as the structure model. The output responses are computed from the inputs using the *lsim* function using a 0.01 s sampling time.

TOSSIM, the TinyOS simulator, provides an event driven simulation of the execution of code written for embedded devices, such as motes, as well as inter-device wireless communication. TOSSIM can simultaneously recreate the execution of a specific code on multiple motes. TOSSIM simulates the network at bit level granularity, capturing collisions and giving the developer an approximate model of the wireless channel. Tinyviz is the visualization tool for TOSSIM and it provides a generic framework to interact with TOSSIM and to control its execution. Plugins can be built upon this framework in order to exchange application-specific information with TOSSIM.

Two approaches are investigated for integrating these simulators. In the *singleshot* mode, the network devices (motes) do not generate any actuation. The MATLAB structural model is given an initial input sequence (possibly an external actuation), and the time series response is generated once, *a priori*, for all the motes and then input into TOSSIM. The *multishot* (or *closed-loop*) mode is more complex, since in this mode some motes may also generate actuation. Initially, the MATLAB model receives an input sequence, generates a response and feeds it into TOSSIM. Once a mote generates an actuation, the new input, along with the original input, is sent to MATLAB to get the new response. This is then repeated in some fashion. The key challenge in the multishot mode is the synchronization between the two simulators.

There are some features provided by TOSSIM and Tinyviz that facilitate the integration of the simulators. TOSSIM has the ability to simulate the sensor ADC channels present in the motes, and Tinyviz has the ability to provide values into these ADC channels. For this integration, a new plugin was developed for Tinyviz to act as a mediator between MATLAB and TOSSIM. The communication path between MATLAB and Tinyviz is shown in Figure 5. The client was incorporated in Tinyviz and the server in MATLAB. The details of the implementation are omitted here for brevity.

To operate this integrated simulation, a user must have installed both the MATLAB Control Toolbox (to use the state-space simulation model of the structure) and TOSSIM, and must provide an input file. The input file is loaded directly from the plugin developed in Tinyviz. The input file should indicate: the desired time of simulations, the time step for the convolution, the description of the structure to be simulated, and the input actuation times.

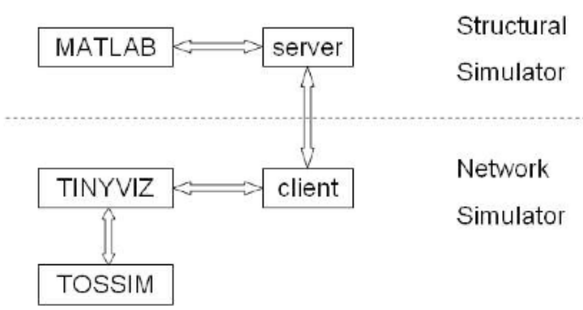

Figure 5. Block Diagram of the client/server architecture used to implement the TOSSIM/MATLAB integration.

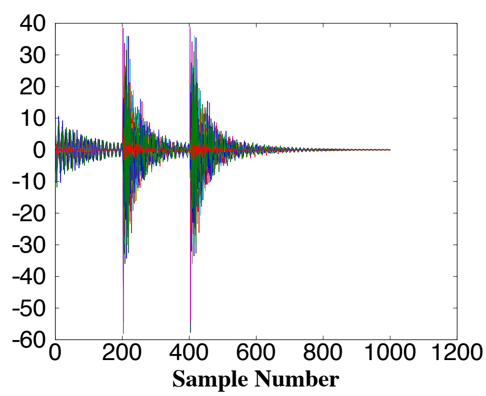

Figure 6. Multi-shot sensor responses.

Figure 6 illustrates sample observed sensor data plots obtained for the multishot case from a simple star topology network of more than a dozen nodes deployed on the structure being monitored. For the multishot case, several nodes inside the structure were instructed to generate actuation after some sensor inputs. Testing and validation of this integrated simulation mechanism is ongoing. In the near future, it will be used for experiments to develop and evaluate SHM protocols. Further, the integration technique will be extended and adapted to other network/structural simulators. Particular extensions in consideration are to the EmStar platform for more sophisticated network simulation, and to NASTRAN for more sophisticated structural simulation.

A LABORATORY TESTBED WITH A MOBILE EXCITER

An important piece of research infrastructure for this project is a test structure that supports local excitations. A novel idea that the authors are investigating is the development of an actuator with constrained mobility (*e.g.*, one that can be moved along rails). Such an actuator can flexibly deliver excitation at different points in a structure, and may be more cost effective than a collection of exciters that provide the same functionality. In this section, the laboratory test structure is described and the preliminary design of a mobile exciter to be placed on this structure is discussed.

The test structure (Figure 7) has a gridwork of massive structural members of various sizes and configuration configured in a form quite similar to a full-scale segment of a large bridge structure. The large size of this test structure will allow more attention on selection, design, and use of some realistic electromagnetic actuators that are capable of generating forces large enough to satisfy reasonable similitude relationships, thus making the range of experimental test parameters realistic. Currently, the structure is still undergoing assembly, instrumenting and preliminary testing. It is planned in the next year to perform simulation (using NASTRAN to generate synthetic data) as well as physical tests to perform structural health monitoring and damage detection experiments on this structure.

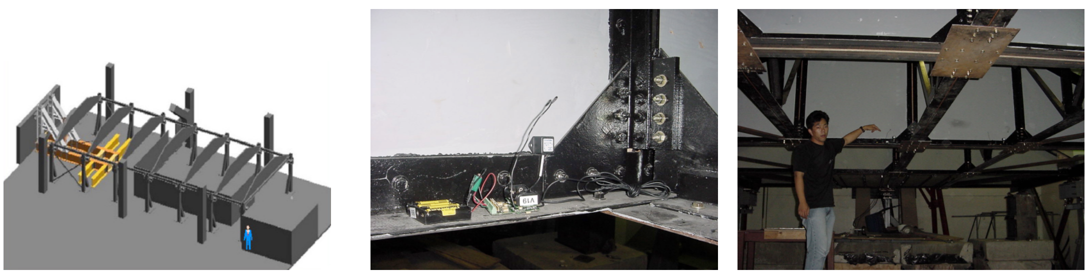

Figure 7. Test apparatus (*a*) sketch (scale indicated by person in lower right), (*b*) instrumentation, and (*c*) structural assembly (indicates high-degree of redundancy in structure)

CONCLUSIONS

An overview is provided of a multi-faceted network-based investigation for structural health monitoring, in which a network of wireless sensors and distributed actuators are used to develop optimum strategies for detecting damage in large structures.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 0325875 (ITR: Structural Health Monitoring Using Local Excitations and Dense Sensing). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

REFERENCES

1. Black, C. J., and C. E. Ventura (1998). "Blind Test on Damage Detection of a Steel Frame Structure." *Proc. 16th Annual Modal Analysis Conference (IMAC XVI)* (Santa Barbara, CA, 1998), pp. 623–629.
2. Intanagonwiwat, C., R. Govindan and D. Estrin (2000). "Directed diffusion: A scalable and robust communication paradigm for sensor networks." *Proc. 6th Annual Int'l Conference on Mobile Computing and Networking* (Boston, MA, August 2000), ACM Press, pp. 56–67.
3. Johnson, E. A., H. F. Lam, L. S. Katafygiotis and J. L. Beck (2004). "The Phase I IASC-ASCE Structural Health Monitoring Benchmark Problem using Simulated Data." ASCE *Journal of Engineering Mechanics*, **130**(1) (January 2004), 5–15.