

Monitoring Civil Structures with a Wireless Sensor Network

Structural health monitoring (SHM) is an active area of research devoted to systems that can autonomously and proactively assess the structural integrity of bridges, buildings, and aerospace vehicles. Recent technological advances promise the eventual ability to cover a large civil structure with low-cost wireless sensors that can continuously monitor a building's structural health, but researchers face several obstacles to reaching this goal, including high data-rate, data-fidelity, and time-synchronization requirements. This article describes two systems the authors recently deployed in real-world structures.

**Krishna Chintalapudi,
Tat Fu, Jeongyeup Paek,
Nupur Kothari,
Sumit Rangwala,
John Caffrey,
Ramesh Govindan,
Erik Johnson, and Sami Masri**
University of Southern California

Large civil structures such as buildings and bridges form the backbone of our society and are critical to its daily operation. Inspectors typically assess them manually, but a networked computer system that could automatically assess structural integrity – and pinpoint the existence and location of any damage – could measurably lengthen a structure's lifetime, reduce its operational cost, and improve overall public safety.

Structural health monitoring (SHM) is a highly active area of research devoted to developing the tools and techniques needed for automatic structural-integrity assessments. Broadly speaking, SHM research attempts to use sensors to localize damage and detect its extent through *structural response* (via the spatio-

temporal patterns of vibrations induced throughout the structure).

Given the challenges involved, the key question is whether it's possible to design large-scale sensor networks for SHM. What principles govern their design? What kinds of system abstractions can help structural engineers develop SHM applications? Our research group at the University of Southern California recently examined these questions, prototyped two software systems, and experimented with them on realistic structures. *Wisden*,¹ a wireless sensor network-based data acquisition system, delivers time-synchronized structural-response data reliably from several locations over multiple hops to a base station. It supports flexible self-organizing sensor-network deployments of up to several tens of

untethered wireless nodes and avoids the high cabling, installation, and maintenance costs incurred by traditional wired data-acquisition systems. *netSHM* provides a programmable sensor-actuator-network system that SHM engineers can use to implement algorithms in a higher-level language such as Matlab or C. It uses a two-tier hierarchy, with resource-constrained sensor nodes in the lower tier and more endowed gateway nodes in the upper tier; theoretically, it can scale to hundreds of nodes. This article summarizes the design of these two systems, discusses their performance, and describes our experiences with them.

The Benefits of Going Wireless

Automatically acquiring and processing data from the several hundred sensors needed to monitor, say, a moderately sized office building requires us to deploy a network. Naturally, the high cost of cabling required to set up a wired network of this size is a serious impediment to developing large-scale SHM systems, but tiny wireless sensors are an easily deployable, low-cost alternative that could bring SHM systems within the realm of practicality. Today's battery-powered, coin-sized devices can contain a processor, significant flash memory, and a low-power radio, together with micro-electromechanical systems (MEMS) sensors capable of measuring vibration. Moreover, these wireless sensor nodes are relatively easy to mount within a few meters of each other. Dense placement greatly increases the spatial resolution of data collection and improves the quality of damage assessment.

The real challenge, though, is how to develop software subsystems that can support a large-scale network of wireless sensor nodes. The batteries on today's wireless sensor platforms barely last a few days, and nodes typically expend a lot of energy in sensing and wireless communication. To conserve energy, sensor-node platforms can operate in various low-power modes, with sensor-node software intelligently duty-cycling the hardware components. In addition, this software can process data locally to reduce the amount of data transmitted wirelessly, a procedure called *in-network processing*.

Although these techniques are generally useful for a variety of networked sensing applications, SHM places additional stringent requirements on wireless sensor node software. SHM sensors generate data at extremely high rates – a single sensor, for example, can generate several hundred

samples per second. Data from sensors or a suitably processed representation thereof must be reliably delivered across the network, but reliable communication in very noisy wireless environments is a significant challenge.

SHM: An Overview

Researchers broadly classify SHM techniques into categories: local and global. Local techniques detect tiny incipient cracks or defects in structures via sophisticated ultrasound, thermal, X-ray, magnetic, or optical imaging techniques. Such imaging equipment is usually expensive, power hungry, and bulky; as such, local SHM techniques are beyond the reach of dense wireless sensing.

Global techniques, on the other hand, can discover damage large enough to influence the properties of the entire structure or large sections of it – for example, significant damage to an entire cable on a bridge or an entire column of a building. Most existing global schemes infer damage from changes in the modes of structural response due to external excitations, either ambient (such as heavy winds and passing vehicles) or forced (such as from shakers or impact hammers). Damage in some of a structure's members results in a change in the mode shapes and frequencies induced in the structure. Responses in larger structures usually comprise frequencies in the tens of Hz and can be sensed with relatively inexpensive low-energy MEMS-based accelerometers. Global techniques therefore fit within the constraints that existing sensor-network platforms impose and are a viable application space for sensor-network-based SHM.

Although SHM techniques based on ambient excitation are easier to deploy, the inherent unpredictability in the excitation's nature and duration makes their design very difficult. A significant advantage of forced excitation systems over ambient ones is that sensors and actuators can be scheduled to perform tests at prespecified times (say, once a day). Forced excitation-based implementations let sensor-actuator networks operate on an extremely low-duty cycle, in which sensors and actuators sleep most of the time and wake up once a day to test the structure. These techniques promise viable long-lifetime sensor-actuator networks. Many structural exciters are commercially available today, producing a variety of excitations such as impulsive impact, white noise, sine sweeps, and so on.

The research on damage detection and localization is vast, as Table 1 illustrates. In general,

Table 1. Taxonomy of structural health monitoring (SHM) techniques.

Time series	Damage detection			Damage localization	
	Modal frequency	Mode shape	Neural network	Time domain	Frequency domain
Inferred from changes in autoregressive/autoregressive moving average (AR/ARMA) coefficients	Inferred by detecting shifts in modal frequencies	Inferred via changes in mode shapes	Networks trained via simulated data	Stiffness of members estimated by constructing a state-space model	Stiffness of members estimated by using mode shapes

SHM techniques infer the existence of damage by detecting changes in overall structural response and then locate the damage itself by estimating the change in stiffness of the structure’s individual members or sections. Many of these techniques use intricate signal-processing methods to detect such changes; we don’t cover such methods in this article.

The SHM application domain places very stringent requirements on wireless sensing software. The complex signal-processing operations that many SHM applications perform can’t tolerate data loss, thus reliable transmission of SHM sensor data is an absolute requirement. This poses significant challenges for wireless sensor networks: their low-power radios, often deployed in harsh radio environments, sometimes exhibit message loss rates of 30 percent. Moreover, SHM signal-processing techniques temporally correlate data from multiple sensors. If the clocks on these sensors aren’t synchronized, phase errors will adversely affect the data analysis. This is particularly true of damage-localization algorithms that attempt to infer differences in mode shapes, which require clock synchronization within tens of microseconds. Finally, SHM sensors generate data at rates comparable to the nominal data rate of today’s low-power radios. A typical structure’s modal frequencies fall within a few tens of Hz, but structural engineers suggest sampling at several hundred Hz for robust estimation. A single sensor node generating 16-bit vibration data along three axes at 500 samples per second can easily consume one-fourth of the nominal data rate of the IEEE 802.15.4 low-power radio commonly used in many wireless applications. A large structure typically comprises hundreds of members and will require at least two tri-axial sensors measuring accelerations at each end of every member to be able to detect damaged members. Clearly, it’s infeasible to continuously collect raw sensor data at a central computer for analysis.

A novel approach to circumventing this data-

rate bottleneck is to process the sensor data within the network, before transmitting it to a central computer. The key challenge, then, is how to adapt existing SHM signal-processing techniques to perform as much data reduction within the network as possible. Time-series-based damage-detection techniques, for example, use autoregressive/autoregressive moving average (AR/ARMA) coefficients — each sensor node can locally compute these coefficients and transmit them to the central location, instead of transmitting the entire data. If each node transmits 40 complex ARMA coefficients instead of 5,000 raw sensor samples, we gain a more than 99 percent savings in communication overhead.

Wisden

The first software system we designed, implemented, and deployed was a wireless structural-data-acquisition system. Wisden accommodates short-term structural testing in situations similar to the following scenario. Let’s say a transportation agency is ready to declare a newly built bridge open, but it allows a team of structural engineers one or two days to measure structural properties. Engineers can deploy a wireless network in tens of minutes, but the challenge in such scenarios is in knowing exactly *where* to instrument the structure because the structural characteristics might not be precisely known. Wisden’s benefit is that it lets engineers iterate on sensor placement to determine appropriate locations.

Specifications

A Wisden deployment typically comprises tens of wireless nodes (placed at various locations on a large structure) that self-configure to form a tree topology and reliably send time-synchronized vibration data to the sink (the tree root), potentially over multiple hops. The sink forwards this data to a base station, usually a high-end PC. Engineers can seamlessly remove or add sensors in a working Wisden deployment by placing a

The Third Generation

In the next phase of our research, we intend to explore robotic excitation. Specifically, we're designing robotic impactors for our seismic test structure. Our eventual goal is to have several of these impactors mounted on the test structure so that we can examine coordinated excitation. The impactors will also carry imaging devices, which will let us explore automated robotic imaging of the structure. Work is now currently under way to equip the seismic test structure with a set of robotic exciters. The robotic actuators will impose a known impulse with sufficient force to excite the structure's various local and glob-

al modes. A total of four robotic devices will be made and mounted to the main support trusses shown in Figure 1 in the main text via an aluminum rail. The robotic devices traverse the structure's length on aluminum rails attached underneath the main trusses. We hope to have the initial prototype ready by August 2006.

Figure A shows each robotic device's final design. The timing of the various actuation operations must be well controlled, and all four robots must operate autonomously. We plan to accomplish this with an onboard computer that will act as an upper-tier node in a netSHM network.




Figure A. Impactor schematic. This robotic impactor will be mounted on rails and can be wirelessly commanded to move to a specific location on our seismic test structure and then impart an impulsive force.

new node and then turning it off or on. In our current implementation, each Wisden node is a Mica-2 or Mica-Z *mote* (sensor platforms currently available from Crossbow Technologies; www.xbow.com) that measures structural vibrations with the help of a *vibration card* specifically designed for the high-quality, low-power vibration sensing suitable for most SHM applications. Attached to this card is a highly sensitive tri-axial accelerometer that can sense up to three channels (three axes) of vibration data.

Wisden implements a negative acknowledgment (NACK)-based hybrid hop-by-hop/end-to-end reliability scheme. Hop-by-hop reliability recovers from message transmission losses between neighboring nodes by observing gaps in received sequence numbers and retransmitting lost packets from a message cache. This technique is necessary to overcome the high (up to 30 percent) wireless message loss rates often observed in real-world deployments. End-to-end reliability is required because hop-by-hop reliability alone can't recover losses when topology changes; a copy of every generated packet is also stored in the source node's flash memory for retransmission. Wisden's base station keeps track of missing packets from all nodes and notifies the appropriate sensor nodes of these packets so that they can retransmit them as needed. Retransmissions can increase network contention at high traffic loads. In our current implementation, nodes are rate-limited to avoid this situation, but we're currently examining adaptive rate-control mechanisms for higher network utilization.

Wisden uses a lightweight approach that time stamps the data consistently at the base station, rather than synchronizing clocks across the network. The Wisden sink synchronizes samples from all the nodes by estimating their generation times according to its own local time. For this, Wisden estimates the received sample's *residence time* (time elapsed between the generation of a sample and its receipt at the sink) and subtracts it from the sample's receipt time at the sink. Ignoring the radio waves' propagation delay (on the order of nanoseconds incurred over several hundred meters of path distance to the sink), the sink can calculate a packet's residence time by summing the times the packet spent at every intermediate node traversed. Each node calculates the time between the sample's receipt to its eventual successful transmission to the next hop and adds it to a residence time field in the packet. The base station (or any node) can thus calculate the time of the samples' generation by subtracting the residence time from its local time.

Finally, Wisden employs an onset-detection-based compression scheme that lets the nodes detect a significant event's occurrence and transmit only data corresponding to that event. As we mentioned earlier, continuous transmission of vibration data from several tens of sensors demands data rates far beyond the capabilities of existing sensor-network platforms. Fortunately, most SHM techniques can receive data corresponding to the occurrence of a "large" event. Wisden's onset detector maintains running estimates of both the noise floor and signal envelope's mean




Figure 1. Two real environments. We deployed Wisden in (a) a seismic test structure and (b) the Four Seasons building.




Figure 2. Wisden's topology. In the Four Seasons experiment, the blue lines show the dominant links Wisden used, and the red lines show some of the occasionally used links.

and standard deviation. When the signal envelope exceeds the noise floor by a few standard deviations, an event is said to have started; Wisden then collects and transmits all samples until the signal envelope subsides.

Deployment

We deployed Wisden in two real environments: a seismic test structure and the Four Seasons, an abandoned office building in Los Angeles. The seismic test structure, shown in Figure 1a, is a full-scale realistic imitation of a 28-foot $\times$ 48-foot hospital ceiling used as a platform for conducting seismic experiments. The ceiling is designed to support 10,000 pounds of weight and has functional electric lights, fire sprinklers, drop ceiling installations, and water pipes. The entire ceiling can be subjected to uni-axial motion with a peak-to-peak stroke of 10 inches, using a

55,000-pound hydraulic actuator with a ± 5 -inch stroke. The hydraulic pump delivers up to 40 gallons per minute at 3,000 pounds per square inch. The total weight of the moving portion of the test structure is approximately 12,000 pounds; currently, it requires a human in the loop to actuate the shaker.

The Four Seasons structure, shown in Figure 1b, is a four-storey office building that was significantly damaged during the 1994 Northridge earthquake and thereafter abandoned and yellow-tagged (meaning its use is restricted due to unsafe areas). This building style is fairly common, which makes the analysis of its earthquake damage especially compelling. The University of California, Los Angeles/Network for Earthquake Engineering Simulation (NEES) project received permission to perform forced vibration testing on the building, so we took this unique opportunity to field-test our system and collect vibration data from the building using our own accelerometers.

Results

Perhaps the single most significant lesson we learned from our deployment on the seismic test structure was that real structures are heavily damped. A typical structural response due to a sudden impact lasts less than a second, a very small time window in which to collect enough samples for analysis. This isn't all that surprising – real structures are designed specifically to discourage sustained oscillations.

However, this does imply that SHM deployments on real structures must support sampling rates high enough to let researchers capture “enough” samples for a reasonable analysis. A sampling rate of, say, 500 Hz provides approxi-

mately 300 samples of structural response. This motivated us to try as high a sampling rate as possible and explore the limits of existing mote platforms. Our explorations revealed that flash memory access latencies on Mica-Z (used to store packets for retransmission in case of losses) become the primary bottlenecks that limit sampling rates. Experiments further revealed that flash memory write latencies vary from a best case of 44 ms to a worst case of 103 ms, allowing a best case of 372 Hz and a worst case of 160 Hz for the sampling rate.²

The Four Seasons deployment provided a very clear example of wireless communication's lossy and highly variable nature in real environments. The communication environment in the building was noticeably worse than in the seismic test structure. The overall average message-reception rate for the links Wisden selected (the system's self-configuration technique estimates wireless message loss rates and picks neighbors with low loss) was 81.12 percent, but for some links, the reception rate was as low as 37.6 percent. Considering the fact that these values are for the chosen links, other candidate links must have been even worse. Figure 2 shows the topology during the experiment. Although the black lines represent the dominant links Wisden used, we observed frequent route changes throughout the experiment; the red lines show some of the sample paths. Some nodes had two or three hop paths to the base station, whereas others had even longer geographic distances with direct connections. The wireless links' variable nature explains the frequent route changes, thus proving that in very large structures, the quality of wireless communication can vary significantly and dynamically.

Ultimately, our experiments proved Wisden's success. Our civil engineering colleagues took several days to set up their wired data acquisition system, but Wisden required roughly 30 minutes for set up before each experiment. This demonstrates the fundamental rationale behind wireless sensor networks: flexibility and ease of deployment.

netSHM

Wisden was a handcrafted system we developed in consultation with structural engineers, but developing additional tools for structural engineers in this manner clearly doesn't scale. We currently classify Wisden as a first-generation system, whose development has yielded insights

```
function shift = getModalShiftsFromBuilding()
%create a group for sensors
idSensors = NetSHMCreateGroup([16,7,13,14,5,2,4,3])
%create a group for actuators
idActuator = NetSHMCreateGroup([6]);
%actuate after 22 seconds
NetSHMcmdActuate(gidActuator,22);
%task motes to sense and send data
% 4000 samples at 200Hz along x axis
    starting 20 secs from now
samples = NetSHMGetSamples(gidSensors,20,200,1,4) ;
% find the modal frequencies from all the samples
modes = findModes(samples);
%read the originally stored modes
load originalModes ;
%detect possible damage
shift = findModalFreqShifts(modes,originalModes);
```

Figure 3. Matlab code for damage detection. All netSHM API functions are prefixed with netSHM.

into the design of a more general programmable system called netSHM.

netSHM is a second-generation wireless sensor network-based SHM system that lets structural engineers implement and test algorithms in a higher-level language such as C or Matlab without having to understand the intricacies of wireless networking. netSHM provides a generic API in the higher-level language that lets users specifically task wireless sensor nodes.

Specifications

Figure 3 shows an example Matlab code for a modal frequency-shift-based damage detection scheme; all netSHM API functions are prefixed with netSHM. The function NetSHMgetSamples(gidSensors,20,200,1,4), for example, schedules eight wireless sensor devices with the specified identifiers in the array gidSensors (16, 7, 13, ..., 3) to collect 4,000 samples at 200 Hz starting from 20 seconds along the x-axis. (An alternative to this procedural programming model is an acquisitional query interface; we haven't explored this because our focus has been on netSHM's networking mechanisms. Our networking mechanisms differ significantly from those provided by a query system such as TinyDB,³ which isn't tiered and doesn't support reliable transmission).

netSHM uses a two-tier hierarchical architecture: resource-constrained wireless sensor nodes




Figure 4. netSHM's two-tier hierarchical structure. Upper-tier nodes have more processing power and higher-capacity radios, and contain application logic. Lower-tier nodes perform sensing, but can also locally reduce data.

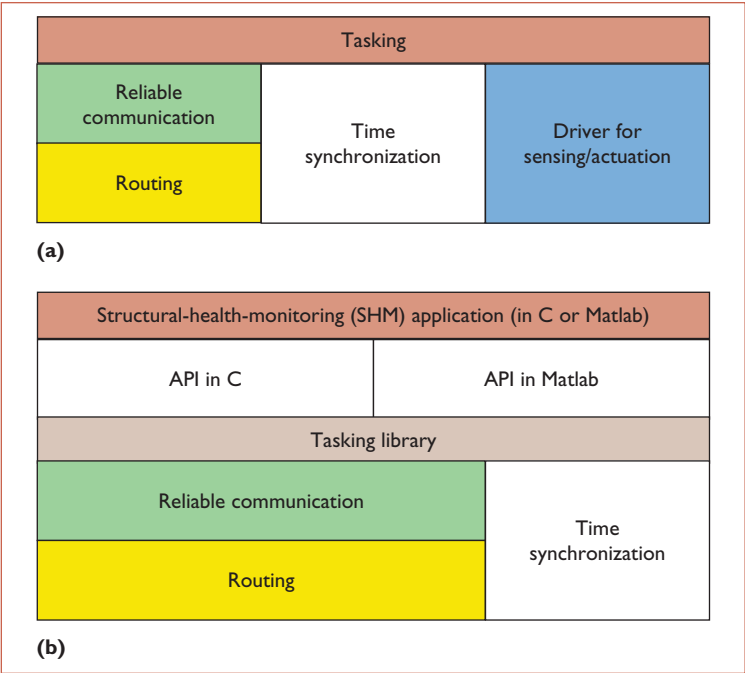


Figure 5. The netSHM stack. On (a) mote-class devices, the mote's tasking library provides a generic task interface, and on (b) gateway-class devices, the gateway node's tasking library translates higher-level functions into commands.

(such as Mica-Z) form a mote-class tier, whereas more endowed nodes (embedded platforms with 32-bit processors, 802.1x radios, and sev-

eral megabytes of flash memory) form the upper tier. In this design, the high-end nodes (for example, Stargates or PCs) self-configure to provide a high-capacity backbone that helps netSHM scale to larger numbers of sensor nodes. These high-end nodes also provide the computational power required for SHM applications, as Figure 4 illustrates.

netSHM's software design departs from current practice, in which energy constraints motivate sensor-network programmers to place optimized application-specific code in the sensor nodes themselves (thus severely restricting code reusability). To allow design of a reusable, programmable system, netSHM constrains application-specific code to reside only in the upper-tier nodes. To achieve this, the high-end nodes issue commands to mote-class nodes via a task interface provided by the mote-class devices themselves. These tasks might command mote-class devices to sample sensors at a specified rate or subject the sensed data to some generic signal processing. An upper-tier device could, for example, task a mote-class device to collect structural response, compute the fast Fourier transform (FFT), and transmit the FFT coefficients in response.

This constraint has another benefit: the networking code for the wireless sensor nodes becomes generic and can be implemented as reusable middleware. The basic netSHM middleware also provides self-configuration, routing, and reliable data transfer and time-synchronization services. Figure 5a shows the netSHM stack on a mote-class device; the mote's tasking library provides a generic task interface and builds on existing reliability and time-synchronization services. Figure 5b shows the netSHM stack on an upper-tier device; the gateway node's tasking library translates higher-level language API functions into mote-tasking commands and uses the reliability layer for command transmission.

Deployment

We deployed netSHM on two structures — a scaled model of a four-storey building and the seismic test structure described earlier. As Figure 6 shows, the scaled building model is 48 inches high, with 1/2-inch × 12-inch × 18-inch aluminum plates that serve as floors and are supported by 1/2-inch × 1/8-inch steel columns. Removable 5.5 lbs-per-inch springs serve as braces between the structure's floors, also aug-

menting the stiffness between the floors. Damage is induced by removing these springs. The building has four wirelessly controlled shakers built from off-the-shelf components. These can be tasked via an attached Mica-Z mote to deliver impulses to the top floor.

Results

We implemented two SHM techniques – one for damage detection and another for damage localization – as netSHM applications. The damage-detection technique is an instance of a modal-frequency-shift-based approach, and the localization algorithm is borrowed from Caicedo and colleagues.⁴ As Figure 7 shows, we successfully detected significant stiffness changes on the floor of the scaled model on which a single instance of damage was induced, as well as damage induced on multiple floors.

Our deployment in the seismic test structure comprised 15 Mica-Z motes and two Stargates. We confirmed netSHM's robustness to gateway-node failures by shutting down one of the Stargates during sample collections. Experiments indicated that netSHM could reconfigure the network topology and ensure reliable delivery.

Our admittedly small-scale sensing demonstrations have shown that sensor networks have a role to play in SHM. We believe they bring intrinsic value because they enable the flexible, low-cost, dense sensing necessary for global damage detection and localization. Moreover, as wireless technology matures and sensor-node platforms evolve, we believe SHM-imposed requirements can be realistically met within one generation of platform technology.

Over the next three to five years, we anticipate that networked sensing and actuation will replace wired instrumentation in structural testing. Many research institutions and companies test various kinds of structures or structural components in specialized laboratory settings. In these settings, the embedded sensor networks that incorporate local processing will replace wired instrumentation and enable dense sensing. These initial deployments will be specialized and will require significant operator intervention to replace batteries or recalibrate sensors. Such intervention is acceptable at small scales.

However, if large-scale structure sensing is to be feasible, it must be accompanied by the devel-




Figure 6. Scaled four-storey building model. To the right, four wirelessly controller actuators impart impulsive force to the structure. On each floor, two accelerometers at diametrically opposite corners measure the structural responses to excitation.

opment of actuator technology suitable for different kinds of structures. It still seems to us that forced-excitation-based methods will be necessary for long-lived networked sensing. Such methods enable sensor-network energy conservation and promise higher accuracy than techniques that rely on ambient sources.

Beyond these, future generations of SHM systems will include movable shakers that can impart structural excitations or might even contain wireless imaging sensors that can detect surface damage. Although we're five to 10 years away from a turnkey SHM system based on wireless sensors and mobile actuators, our research has started to lay the foundations. □

References

1. N. Xu et al., "A Wireless Sensor Network for Structural Monitoring," *Proc. ACM Symp. Networked Embedded Sensing (Sensys)*, ACM Press, 2004, pp. 13–24.
2. J. Paek et al., "A Wireless Sensor Network for Structural Health Monitoring: Performance and Experience," *Proc. 2nd IEEE Workshop on Embedded Networked Sensors*, IEEE CS Press, 2005; www.cse.unsw.edu.au/~emnet.
3. S. Madden et al., "TAG: A Tiny Aggregation Service for Ad Hoc Sensor Networks," *Proc. Usenix Symp. Operating Sys-*




Figure 7. Damage localization using netSHM. The changes in stiffness on each floor result from different damage patterns; the localization algorithm relies on these changes to identify which floors have been damaged.

tems Design and Implementation, Usenix Assoc., 2002; www.usenix.org/events/osdi02.

- J.M. Caicedo, S.J. Dyke, and E.A. Johnson, "Natural Excitation Technique and Eigen System Realization Algorithm for Phase I of the IASCASCE Benchmark Problem: Simulated Data," *J. Eng. Mechanics*, vol. 130, no. 1, 2004, pp. 49–61.

Krishna Chintalapudi is currently involved in research in emerging wireless technologies at the Bosch Research and Technology Center in Palo Alto, California. His research interests include delay-sensitive and fault-tolerant wireless sensor networks, high-data-rate wireless sensor network applications, hybrid sensor networks, and wireless control and automation. Chintalapudi has a PhD in computer science from the University of Southern California and an MS in electrical engineering from Drexel University. Contact him at krishna.chintalapudi@rtc.bosch.com.

Tat Fu is a graduate research assistant at the University of Southern California, where he is pursuing both a PhD in civil engineering and an M.Arch. from the School of Architecture. Fu has a BS and an MS in civil engineering from the University of Southern California. Contact him at tsf@usc.edu.

Jeongyeup Paek is a member of the Embedded Networks Laboratory at the University of Southern California and a graduate student in the university's PhD program. Paek has a BS from Seoul National University and an MS from the

University of Southern California, both in electrical engineering. His research focuses on topics in wireless sensor network systems such as reliable transport infrastructures, architecture for tiered embedded networks, and sensor-network applications such as structural health monitoring. Contact him at jpaek@enl.usc.edu.

Nupur Kothari is a PhD student in the University of Southern California's computer science department. Her research interests include programming techniques for sensor networks, data aggregation in sensor networks, and geographic routing. Contact her at nkothari@usc.edu.

Sumit Rangwala is a graduate student at the University of Southern California, Los Angeles. His research focuses on congestion control in wireless sensor networks. Rangwala has a BS in computer science from SGSITS, India, and an MS in computer science from the University of Southern California. Contact him at srangwal@usc.edu.

John Caffrey is an associate professor of mechanical engineering at the California State Polytechnic University, Pomona, and a research assistant professor at the University of Southern California. He has a BS in mechanical engineering from the University of Akron, an MS in mechanical engineering from the University of California, Santa Barbara, and a PhD in civil engineering from the University of Southern California. Contact him at jcaffrey@usc.edu.

Ramesh Govindan is an associate professor in the University of Southern California's computer science department. His research interests include scalable routing in internetworks and wireless sensor networks. Govindan has a B.Tech. from the Indian Institute of Technology at Madras, and an MS and a PhD from the University of California, Berkeley. Contact him at ramesh@usc.edu.

Erik Johnson is an associate professor of civil engineering at the University of Southern California, where his research interests include structural health monitoring, structural control, and computational stochastic dynamics. Johnson has a BS, an MS, and a PhD in aeronautical and astronautical engineering from the University of Illinois at Urbana-Champaign. Contact him at johnsone@usc.edu.

Sami Masri is a professor in the Departments of Civil and Environmental Engineering and Aerospace and Mechanical Engineering at the University of Southern California. His research interests include the modeling and control of nonlinear systems. Masri has a BS from the University of Texas, and an MS and a PhD from the California Institute of Technology. Contact him at masri@usc.edu.