Special issue on International Workshop on Networked Sensing Systems

Trans. of the Society of Instrument and Control Engineers
Vol.E-S-1, No.1, 109/114 (2006)

High-Frequency Distributed Sensing for Structure Monitoring†

Kirill MECHITOV *, Wooyoung KIM *,
Gul AGHA * and Tomonori NAGAYAMA **

Structural health monitoring (SHM) involves continuous monitoring of a structure's condition with near real-time analysis of sensed data. The emerging wireless sensor network (WSN) technology enables distributed data acquisition and processing for structure monitoring applications, which currently adopt a centralized solution, using a network of intelligent sensors. SHM poses a challenging problem for sensor networks with its dual requirements of high-frequency sensing and structure-wide coordination. We design a robust, modular WSN-based distributed sensing system for high-frequency data acquisition in structure monitoring. Our contributions are twofold. The system emulates the functionality of a centralized solution employing analog sensors, achieving comparable performance. The resulting compatibility with existing SHM algorithms will facilitate the adoption of sensor network technology for structure monitoring. Additionally, the data aggregation, distributed sensing and time synchronization services comprising the system enable *in situ* data processing and aggregation in WSNs. This functionality will enable a seamless transition from centralized to fully distributed SHM applications, with no changes to the infrastructure.

Key Words: sensor networks, structural health monitoring, distributed sensing, data aggregation

1. Introduction

The goal of structural health monitoring (SHM) is to determine the condition of the monitored structure, such as a building, bridge or an airplane, and to identify potential problems at an early stage, by examining the outputs of sensors attached to the structure. This process typically involves measuring strain values or vibration characteristics at different points in a building, or measuring the load on bridge supports. With few exceptions, existing SHM applications have adopted a centralized approach, employing a small number of analog sensors wired to a central controller [3]. In these systems, data acquisition and processing capacity of the central node and the wiring complexity place limits on scalability, usually supporting 10-20 sensors at the most. Wiring sensors to a central node in a large structure is expensive and cumbersome (the wires may cost more than the sensors!); it may also be detrimental to system reliability as wires may be damaged or severed.

Recent advances in sensing and networking technologies have led to the emergence of wireless sensor networks (WSNs) as a new computing platform. Composed of a large number of small, intelligent sensor nodes, WSNs have started supplanting centralized sensing and control systems with a low-cost distributed alternative [12]. WSNs are attractive for structure monitoring applications for a number of reasons. They offer reduced cost, increased robustness through decentralization, and a degree of resiliency to node and link failures. In particular, a distributed sensor network could continue to function, at diminished capacity, even after sustaining the loss of a large fraction of the sensors. Moreover, wireless communication means WSNs do not suffer from wire breaks in catastrophic events such as earthquakes. An additional benefit of "smart" sensors is that some of the data processing may occur locally at the sensors, reducing the control turnaround time and improving overall system responsiveness.

We develop a high-frequency distributed sensing system for structure monitoring applications on a WSN platform consisting of Mica-2 motes [2]. The nodes used in our experiments are equipped with a 433MHz RF transceiver for low data rate wireless communication. While several alternative radios are available, including Bluetooth and 802.15.4 (ZigBee), low power and low bandwidth are an inherent characteristic common to all wireless sensor plat-

† Presented at INSS2004 (2004.6)

* Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA

** Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA

(Received November 14, 2004)
(Revised June 27, 2005)

forms. The amount of available memory is also severely limited: 4KB SRAM and 512KB serial flash. We design the robust data acquisition system for this resource-constrained environment, bringing data processing and control into the network. The system is modular and highly customizable; options are available for *in situ* filtering and data processing, reliable or best-effort data aggregation, and multiple, dynamically-selected sensing modalities.

In this paper, we describe our experience with the distributed data acquisition system for centralized data processing on a wireless sensor platform. The objective of building such a system is to facilitate the transition to fully distributed SHM applications in the future, while supporting ongoing research on structure monitoring utilizing the existing centralized techniques. We describe the design of the system, present empirical evaluation results and discuss the limitations of current WSN platforms in SHM and propose directions for future research.

2. Structure Monitoring with Wireless Sensor Networks

In traditional centralized data acquisition systems, sensors continuously generate data that are periodically sampled by a central data processing unit. SHM computations are performed after the desired amount of data is collected. We develop a WSN-based system that emulates the functionality of centralized sensing systems, with added benefits of increased robustness, lower cost and seamless transition to distributed sensing, which we consider the future of SHM.

2.1 Requirements and Design

Buildings and bridges employing SHM need to be monitored for extended durations, periodically checking the state of the entire structure. Unlike applications like habitat monitoring, which continuously sample the sensors at a very low frequency, structure monitoring algorithms need to detect vibrations with frequencies up to 100Hz. Thus the Nyquist limit for the sampling frequency is at least 200Hz. With the current WSN technology, it is not possible to sample sensors continuously at this rate while exfiltrating the data over the wireless network.

In order to provide accurate sensing and reduce sample jitter, we must disable other sources of high-priority interrupts in the sensor node, *e.g.*, the radio. For this reason, we first store sensor data in flash memory to be retrieved later. Thus, the size of available flash memory limits the maximum duration of uninterrupted sensing. On Mica-2 motes, this limitation translates into recording approximately 90 seconds of continuous data sampled at 250Hz. Thus the data acquisition process operates via intermittent periods of high-frequency sensing and communication or idle time.

Secondly, the system mush have a way to deliver sensor data to the processing station reliably, since SHM algorithms developed for wired systems require that no sensor readings be lost in order to produce correct results. We design such a transport protocol, although it is costly to implement over an unreliable multi-hop wireless network. We also provide alternative protocols for best-effort and probabilistic data transport, *e.g.*, guarantee delivery of at least 99% of the packets with high probability.

Thirdly, overall system robustness is a critical requirement for a monitoring system operating for a period of several years. Individual sensor nodes are prone to failures, and changes in the environment may alter the topology of the wireless network. We develop an adaptive self-healing tree routing service for establishing a mesh network among the sensors to transport the data efficiently and reliably. We use *path length* and *link quality* to determine the best routes. Originally, we employed path length as the primary criterion for establishing the tree structure; however, we found that emphasizing link quality is far more important for robust packet delivery in volatile wireless networks. Connectivity within the network is periodically checked and maintained, as each node periodically sends out heartbeat messages to its neighbors. The advantages of this method include:

(1) optimal bandwidth utilization: the tree is built for sensor-to-sink communication, which is dominant in centralized data acquisition

(2) memory efficiency: the routing table consists of only one node

(3) simple and fast fault recovery: the aggregation tree can quickly adapt to node failure without causing global topology changes.

Finally, it is crucial that the sensors' clocks are synchronized within a tight error bound. This problem does not arise in a centralized system, where there is only one clock; however, distributed sensor readings are meaningless from the application standpoint unless they can be correlated on a consistent global time scale. We adapt the FTSP time synchronization protocol [7] to maintain clock synchrony. We make this protocol more efficient by piggybacking time synchronization beacon messages on the routing tree heartbeat messages. Thus we are able to establish network-wide synchronization with no significant overhead. The time synchronization service can maintain

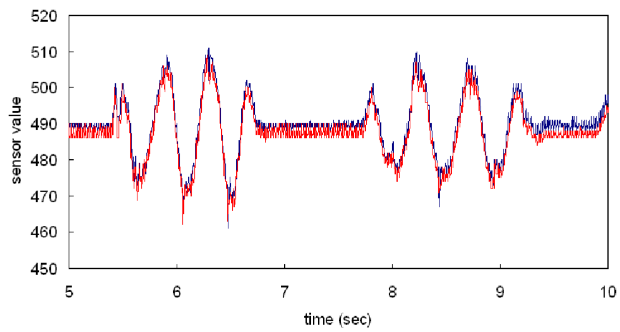

Fig. 1 Accelerometer data from two sensors showing tight synchronization.

better than 1ms synchronization on a 10-hop network for extended periods (**Fig. 1**). This is sufficient to ensure proper synchronization for our target application, where continuous sensing is limited to 90-second intervals, and the sampling period is 4ms.

2.2 Distributed Sensing and Data Aggregation

When put together, the sensing, routing and time synchronization services form the basis for building a distributed sensing and control platform. Distributed sensing enables the application of SHM techniques to large-scale structures, and local actuation is useful for control applications as it decreases reaction time and improves overall system responsiveness. We use a composition of these services to emulate a centralized sensing system that features sensing frequencies and synchronization precision similar to its centralized counterparts, albeit with higher reliability. Its purpose is to facilitate transition to truly distributed sensing and control.

The centralized nature of the current application puts significant stress on the low-bit rate wireless communications of the micro sensor devices. As data flows from the sensor nodes to the sink, congestion becomes a significant issue, exacerbated by the MAC layer based on a CSMA-CA (carrier sense multiple access–collision avoidance) scheduling algorithm [11]. To address this issue, we implement an opportunistic data aggregation service for the high-frequency sensing application.

The high sensing frequency places a bound on the amount of data processing that can occur within the network. In particular, the Fast Fourier Transform (FFT) algorithm, a central part of the vibration analysis-based structure monitoring algorithms, has a run time of several minutes on the Mica motes [1]. We are thus constrained primarily to aggregators based on linear functions on the time-domain data: average, median, extrema, zero-crossing and peak detection.

To be meaningful from the application standpoint, only data from immediate physical neighbors can be aggregated in this manner. Aggregation operations are performed on timestamped blocks of contiguous samples if the timestamps of two neighbors match. Each application data packet carries up to three such blocks. The data aggregation service is *opportunistic*, meaning that combinable blocks are only aggregated if they happen to be stored on an intermediate node at the same time while being transported to the sink. This method is less efficient than the more aggressive time-based aggregation [16], where messages are not forwarded along the tree until a later timestamp from the source node arrives; however, our choice is dictated by the very limited buffer space available on the nodes.

Our experimental results, discussed in the next section, demonstrate that our approach combines a large fraction of packets close to the source without the additional delays and buffer space requirements incurred by time-based aggregation methods. This is primarily due to the fact that our the aggregation tree closely matches the physical topology, given the emphasis on short, high-quality links noted above.

3. Evaluation

We study the performance of the distributed sensing system using two sets of experiments: accelerometer and strain gage sensor measurements to verify precision, and networking tests to examine time synchronization, scalability and data aggregation performance. The objective of the experiments is to demonstrate the viability of this system for research and industrial SHM applications.

We use Mica-2 motes from Crossbow Technology, Inc. [2], equipped with a mix of the MTS310 sensor boards, an improved accelerometer sensor board, and a strain gage sensor board. Due to the limited hardware availability of the custom sensor boards, strain gage and accelerometer experiments are conducted on a variable-resistance 3-story building model, while the networking experiments take place on a larger rigid 18-story building model (**Fig. 2**) placed on a shaking table. The table produces white noises at frequencies of 1 to 100Hz, the frequency range used in most vibration-based SHM applications. Sensing is performed at 250Hz for 60 second intervals. Sensors are located on each floor of the building model, spaced approximately 30cm apart and separated by a metal "floor." Radio transmit power of the CC1000 chip is reduced to the lowest setting (−20dB) and no external antenna is used in order to induce multi-hop communication. In experiments, we have observed

Fig. 2 Sensors are attached to an 18-story building model on a shaking table.

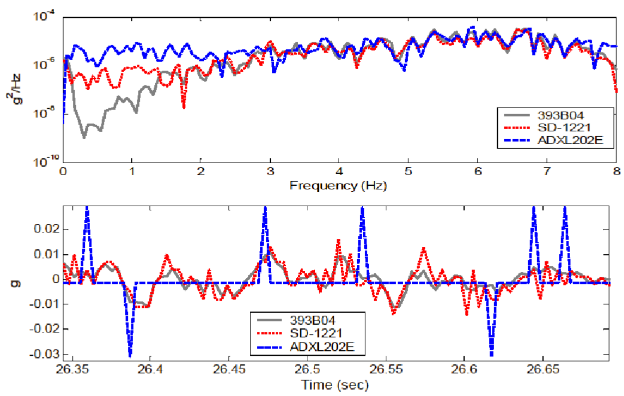

Fig. 3 Frequency- and time-domain performance of wired (SD-1221), Mica (ADXL202E) and custom (393B04) accelerometers[13].

path lengths of up to nine hops.

For comparison, a typical tower building where an SHM application may be deployed is 100 to 300m tall and contain one sensor per one to five floors. More advanced SHM applications may require multiple sensors per floor.

3.1 Sensor Precision

The Mica sensor board's standard accelerometer possesses insufficient precision to measure vibrations of the magnitude required for SHM. Instead, we employ sensor boards customized with an accelerometer (Silicon Designs model 1221[14]), which is suitable for low noise applications[13]. Its performance is comparable to that of analog accelerometers (*e.g.*, PCB Piezotronics model 393B04[10]) routinely used in wired SHM systems (**Fig. 3**). We also utilize a strain sensor developed specifically for use in SHM applications with Mica-2 motes[9], capable of matching the resolution and sensitivity of industrial wired strain gages (**Fig. 4**).

3.2 Data Aggregation

For centralized sensing in a Mica-2 sensor network, the nodes must store their measurements in local memory and later exfiltrate them to the processing node. Data from different sensors are aggregated along the way to the sink node, reducing the overall amount of data transmitted. The efficiency of data aggregation is critical to the scalability of the distributed sensing system. Even a fully-distributed system would still rely on data aggregation to reduce the amount of localized communication, though on a neighborhood, rather than network-wide, scale. We now evaluate the performance of the data aggregation component of our system.

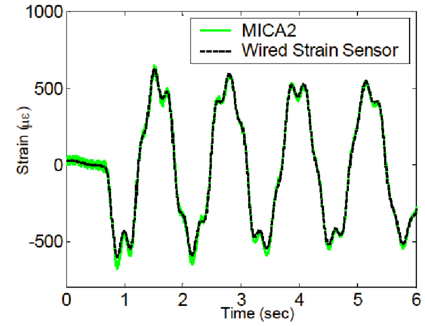

Fig. 4 Performance comparison of a wired strain sensor and the custom Mica-2 strain gage[9].

Table 1 Aggregation chance vs. distance from source.

Distance (hops)	Blocks aggregated (%)
1	46
2	13
3	2
≥ 4	0

As discussed earlier, compatible data blocks (physical neighbors, matching timestamps) are only combined if they happen to reside on the same intermediate node in transit to the sink. The total bandwidth consumed by the exfiltration of compatible data blocks can be computed as $B = \sum_i h_i + h_c$, where h_i is the number of hops the ith block is forwarded prior to aggregation and h_c is the number of hops remaining after aggregation. It is clear from this formula that the most important metric to gauge the efficiency of this strategy is not just the number of blocks actually combined, but also the *location* in the tree where the combination process takes place.

Table 1 summarizes this information. About 62% of the total number of data packets generated are actually combined by our data aggregation service. This is significantly less than the close to 100% combination rate that can be achieved with time-based data aggregation; however, this is primarily due to the very small five-block aggregation buffers available on the Mica-2 as a consequence of the limited RAM space.

Since near-real time processing of the data is desirable for long-term structure monitoring, latencies incurred in data transport and aggregation are also an important performance metric. **Table 2** shows the time it takes to transfer all of the data recorded at 250Hz for 60 seconds and the ratio of aggregation time to sensing time, for vary-

Table 2 Aggregation time and aggregation-to-sensing ratio.

Sensors	Time (s)	Ratio
1	288	4.80
2	491	8.18
4	818	13.63
8	2136	35.60
16	6287	104.78

ing numbers of sensors. The data aggregation phase is quite long even with a relatively small number of sensors. This is to be expected, since even two sensors can fully saturate the bandwidth of the sink node. (Note that the theoretical maximum bandwidth of a Mica-2 mote is limited to only 19.2Kbps.) Additional sensors competing for the bandwidth only increase congestion and exacerbate the sensing/aggregation imbalance.

There are several causes for the long aggregation times, particularly as the number of sensors is increased. First and most important, the bandwidth becomes saturated at the top levels of the routing tree. Second, as the wireless channel utilization grows, the frequency of collisions resulting in packet loss grows sharply since the collision avoidance-based MAC protocol is inefficient in high traffic volume conditions. Lastly, limited buffer space and acknowledgment-based reliable communication result in a large number of retransmissions as packets and acknowledgments are dropped at intermediate nodes. Retransmitted packets are also much less likely to be aggregated, as the chances of encountering packets from neighboring nodes with the same timestamp are quite small.

To demonstrate that the performance degradation inherent in the configuration of the nodes and not an artifact of or design or implementation, we simulate the performance of the system with the same network topology, but with a collision-free MAC layer with 100% channel capacity utilization and lossless channels. This is the absolute best-case scenario for this problem. Note however that even in this case packet drops and retransmissions may still occur due to insufficient buffer space at intermediate nodes. **Fig. 5** presents a performance comparison of the best-case simulation and the experimental implementation. While these numbers are somewhat better than the experimental results, the best-case scenario exhibits similar scalability behavior. Bandwidth and buffer space considerations continue to limit the scalability of the data aggregation service. We discuss the implications of this finding in the next section.

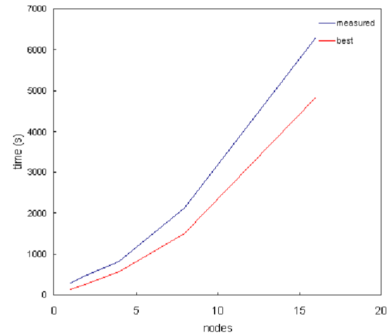

Fig. 5 Performance comparison of the best-case simulation and experimental measurements.

4. Discussion

The core problem for high-frequency sensing applications is that they simply generate too much data to be efficiently transported over the network. Data compression may help alleviate the sensing/aggregation imbalance to some extent, although only by a constant factor. More aggressive data aggregation approaches may also provide some benefit [4),6),16)]. Neither is likely to solve the problem completely, however.

Another way to tackle the bandwidth and buffer space limitations is to employ *in situ* data processing: letting sensors perform local processing on the raw data and transfer only the summary results would keep the network from being saturated. In the Mica-2 platform, one microprocessor controls both sensing and communication on a node. Having separate controllers, a capability offered by more powerful Intel motes [8)], would enable concurrent sensing and communication, further reducing the extent of the imbalance. Such an architecture would also open up the possibility of new communication controller designs optimized for power usage [5)].

In principle, the problem stems from the design of low-rate communication protocols in emerging WSNs (*e.g.*, the Mica-2 protocol and ZigBee/IEEE 802.15.4 [15)]). Given the limited power supply, the choice of "low-rate" is uncompromisable. Thus, these low-bandwidth networks cannot perform well in high-volume data transfers, which are unavoidable in centralized SHM applications. Our experiments have confirmed this, and lead us to conclude that SHM applications built on WSNs should be fully distributed, limiting the scope over which the data is to be aggregated. Our communication and synchronization services will facilitate the transition.

5. Conclusion

Through the experiments using realistic building models we have shown that our data acquisition system is capable of achieving time synchronization precision and sensing resolution comparable to those observed in wired

SHM systems employing analog sensors. The data collection phase, while significantly longer than in wired sensing systems, can still accommodate SHM applications where sensing periods are interspersed with long intervals of inactivity. This system is also more robust than its centralized, wired predecessors. Further, the high adaptability of our modular design makes it suitable for use in fully distributed sensing and control applications, which we view as a key future direction for structural health monitoring and control research.

6. Acknowledgements

This material is based upon work supported in part by the Defense Advanced Research Projects Agency (DARPA) under Award No. F33615-01-C-1907, by the National Center for Advanced Secure Systems Research (NCASSR) Award No. N0014-04-1-0562 and by ONR MURI under Award No. N0014-02-1-0715.

References

1) T. Canli, M. Terwilliger, A. Gupta, and A. Khokhar. Power-time optimal algorithm for computing FFT over sensor networks. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, pages 293–294, 2004.
2) Crossbow Technology, Inc. http://www.xbow.com/.
3) S. Doebling, C. Farrar, M. Prime, and D. Shevitz. Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: A literature review. Technical Report LA-13070-MS, Los Alamos National Laboratory, 1996.
4) T. He, B. M. Blum, J. A. Stankovic, and T. Abdelzaher. AIDA: Adaptive application-independent data aggregation in wireless sensor networks. *Transactions on Embedded Computing Systems*, 3(2), 2004.
5) J. L. Hill. *System Architecture for Wireless Sensor Networks*. PhD thesis, University of California, Berkeley, 2003.
6) C. Intanagonwiwat, D. Estrin, R. Govindan, and J. Heidemann. Impact of network density on data aggregation in wireless sensor networks. In *Proceedings of the 22nd International Conference on Distributed Computing Systems*, 2002.
7) M. Maroti, B. Kusy, G. Simon, and A. Ledeczi. The flooding time synchronization protocol. In *Proceedings of the 2nd Conference on Embedded Networked Sensor Systems*, 2004.
8) L. Nachman, R. Kling, R. Adler, J. Huang, and V. Hummel. The Intel Mote Platform: A Bluetooth-based sensor network for industrial monitoring. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks*, 2005.
9) T. Nagayama, M. Ruiz-Sandoval, B. F. Spencer, K. Mechitov, and G. Agha. Wireless strain sensor development for civil infrastructure. In *Proceedings of the 1st International Workshop on Networked Sensing Systems*, pages 97–100, 2004.
10) PCB Piezotronics, Inc. http://www.pcb.com/.
11) J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd Conference on Embedded Networked Sensor Systems*, 2004.
12) G. J. Pottie and W. J. Kaiser. Wireless integrated network sensors. *Communications of the ACM*, 45(5):51–58, 2000.
13) M. Ruiz-Sandoval, B. F. Spencer, and N. Kurata. Development of a high sensitivity accelerometer for the mica platform. In *Proceedings of the 4th International Workshop on Structural Health Monitoring*, 2003.
14) Silicon Designs, Inc. http://www.silicondesigns.com.
15) ZigBee^{TM} Alliance. http://www.zigbee.org/.
16) S. Zou, I. Nikolaidis, and J. J. Harms. Efficient data collection trees in sensor networks with redundancy removal. In *Proceedings of the 3rd International Conference on Ad-Hoc Networks and Wireless*, 2004.

Kirill Mechitov

Kirill Mechitov received his B.S. in computer science from the Georgia Institute of Technology in 2001. He is currently a Ph.D. candidate in computer science at the University of Illinois at Urbana-Champaign. His research interests include wireless sensor networks, distributed algorithms and middleware for network embedded systems.

Wooyoung Kim

Wooyoung Kim is a research assistant professor at the University of Illinois at Urbana-Champaign. He earned MS and BS degrees in Computer Engineering at the Seoul National University in 1987 and 1989, respectively. He went on studying to get his Ph.D. in Computer Science at the University of Illinois at Urbana-Champaign in 1997.

Gul Agha

Gul Agha is Professor of Computer Science at the University of Illinois at Urbana-Champaign. He received his schooling in Sindh where he grew up, his BS with honors from Caltech, and AM (Psychology), M.S. and PhD (Computer and Communication Science) from the University of Michigan. Dr. Agha is a Fellow of the IEEE and serves as the Editor-in-Chief of ACM Computing Surveys.

Tomonori Nagayama

Tomonori Nagayama obtained his B.S. (2000) and his M.S. (2002) in civil engineering from the University of Tokyo. He is currently a Ph.D. candidate at department of civil and environmental engineering, University of Illinois at Urbana-Champaign, USA. His research interests include structural health monitoring and intelligent structure.

Reprinted from Trans. of the SICE
Vol. E-S-1 No. 1 109/114 2006