Invited Paper

When Sensor and Actuator Networks Cover the World

John A. Stankovic

The technologies for wireless communication, sensing, and computation are each progressing at faster and faster rates. Notably, they are also being combined for an amazingly large multiplicative effect. It can be envisioned that the world will eventually be covered by networks of networks of smart sensors and actuators. This fact will give rise to revolutionary applications. However, to make this vision a reality, many research challenges must be overcome. This paper describes a representative set of new applications and identifies several key research challenges.

Keywords: Wireless, communications, sensors, networks, knowledge creation, robustness, openness, security, privacy.

I. Introduction

Wireless sensor and actuator networks (WSANs) constitute an important and exciting new technology with great potential for improving many current applications as well as creating new revolutionary systems in areas such as global-scale environmental monitoring, precision agriculture, home and assisted living medical care, smart buildings and cities, industrial automation, and numerous military applications (see Fig. 1). Typically, WSANs are composed of large numbers of minimal capacity sensing, computing, and communicating devices and various types of actuators (see Fig. 2). These devices operate in complex and noisy real world, real-time environments. Current and past research have produced many excellent low level mechanisms and protocols to collect, transport, and perform sensor fusion of this raw data and react with control actions. However, many challenges remain.

In this paper, a projection into the future is hypothesized when WSANs will cover the world. Section II briefly relates WSAN to other terms, such as ubiquitous computing and wireless sensor networks (WSNs). Section III presents a vision for several future applications. Section IV articulates some of the key challenges to achieve such applications. Section V describes related work. Section VI summarizes the main ideas and mentions additional research topics.

II. Terminology

Terminology is very important, but it can also be confusing. For example, embedded systems and real-time systems are highly related fields, but not identical. These two areas are represented by different research communities with different emphasis: embedded systems emphasize form factor, cost and other constraints, while real-time systems emphasize timing

Manuscript received May 22, 2008; revised June 26, 2008.

John A. Stankovic (phone: + 1 434 982 2275, email: stankovic@cs.virginia.edu) is with the Department of Computer Science, University of Virginia, Virginia, USA.

Fig. 1. Three important applications: volcano monitoring, bridge monitoring, and assisted living.

properties. Nevertheless, most real-time systems are embedded and vice versa.

Similarly, ubiquitous computing, pervasive computing, WSNs, and WSANs are highly related, but not identical. They have different communities of researchers. Ubiquitous and pervasive computing emphasize computation, while WSNs emphasize communications and sensing. In the future, these fields will increasingly overlap, taking properties from each other. This paper uses the term WSANs to refer to the future view of the progression of WSNs from mostly passive sensing to sensing, actuation, and increased computation.

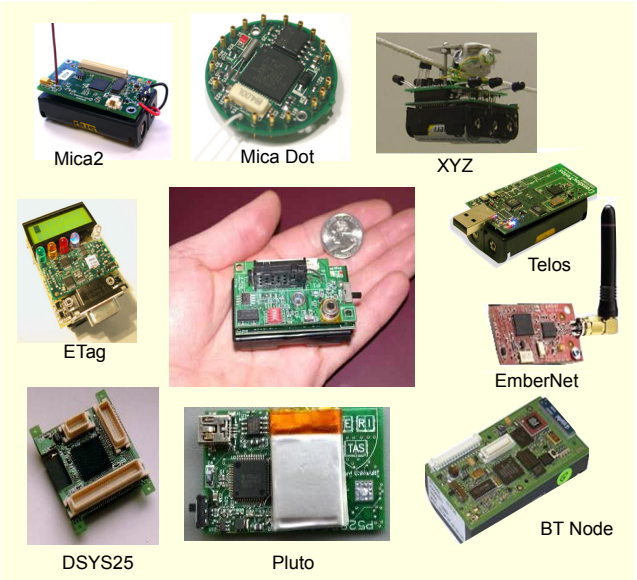

Fig. 2. Representative smart sensor nodes.

III. Future Applications

To illustrate the transformative nature of WSAN technology, three future applications are briefly discussed: global scale environmental monitoring and control, social participatory computing, and continuous birth-to-death health care.

1. Global Scale Environmental Monitoring and Control

Today, many sensors exist around the world collecting environmental data. In most cases, the systems focus on single problems, such as the effect of tides on barrier islands off the coast of Virginia or tornado watches in the central parts of the United States. Most of these systems measure a limited number of parameters at a large granularity (sensors separated by 100s or 1,000s of meters). WSANs have the potential of dense and flexible coverage and most importantly enabling correlation across many WSANs. Such capabilities will result in new understanding of environmental conditions. Dense coverage might include sensors placed within centimeters or meters of each other, enabling a precise understanding of certain phenomena. This will permit applications such as micro-agriculture where local conditions can dictate control of pesticide and fertilizer amounts. When unexpected environmental events occur or when natural disasters destroy the current infrastructure (earthquakes, cyclones), WSAN technology can be rapidly deployed for almost immediate collection of data. This can improve rescue efforts and increase our understanding and control of disease outbreaks. When WSANs cover the world, it will be possible to collect data in real time and produce correlations among systems. Instead of

studying the effect of pollution on a particular fish in one river system, it will be possible to correlate that phenomenon with populations of other fish, other river systems, weather conditions, social activities, and many other data collections. Some of these correlations may be discovered via data mining. For another example, it will also be possible to track the spread of human disease such as the flu and correlate that with outbreaks of animal disease as well as weather conditions. Global weather patterns based on air, land, and ocean interactions will be better understood.

2. Social Participatory Computing

The ubiquity of WSAN technology will include devices worn and carried by individuals as well as many emplaced systems in all our surroundings. The access to real-time data will transform how people work, socialize, and go about their daily routines. Individuals will be able to track commuting delays and minimize or avoid problems associated with these delays. Municipalities will be able to control traffic patterns to reduce congestion. People will also be able to link to groups of individuals with similar interests and interact with them from anywhere at anytime. This includes family groups, work groups, medical groups, and social groups. Preferences can be automatically incorporated into these activities. Automatic notifications for special social events such as a concert or sales on products currently of interest to an individual will be routine. Overall, individuals will become enmeshed with and rely on these WSANs for efficient and happier lifestyles.

3. Continuous Birth-to-Death Health Care

One use of the widespread availability of WSANs will be for medical care. It will be possible to create an "account" for each person when they are born and via physiological, activity, and environmental sensing keep a record of their health and activities that relate to health. Using such information, preliminary diagnosis can be achieved without a doctor's visit. This will also enable preventative care. As one ages or as one's health deteriorates, specific devices can be included in smart clothes or within the home to ameliorate the condition and/or treat it more effectively. In some cases, automatic medical treatment might be administered. Any such automatic control actions must be guaranteed to be "safe." Holistic and long term health information on individuals will enable dramatic improvements in their care as well as overall understanding of health problems and the effectiveness of treatment. It has been shown that elderly individuals are happier and healthier when participating in social activities. The fact that WSANs will create social participatory computing will further improve health care.

IV. Challenges

Many challenging research problems must be addressed before WSANs of the future are in widespread use. This paper concentrates on a few critical ones, specifically the following: from raw data to knowledge; robust system operation; as well as openness, security, and privacy.

1. From Raw Data to Knowledge

In a world covered by WSANs, a vast amount of raw data will be continuously collected. It will be necessary to develop techniques that convert this raw data into usable knowledge. For example, in the medical area, raw streams of sensor values must be converted into semantically meaningful activities performed by or about a person, such as eating, poor respiration, or exhibiting signs of depression. The main challenges for data interpretation and the formation of knowledge include addressing noisy physical world data and developing new inference techniques that do not suffer the limitations of Bayesian or Dempster-Shafer schemes. These limitations include the need to know *a priori* probabilities and the cost of computations. Rule based systems may also be too ad hoc for dealing with WSANs. Given that a very large number of WSANs will exist, with each providing many real-time sensor streams, it will be common for a given stream of data to be used in many different ways for many different inference purposes. Enabling streams to act as primitives for unexpected future inferences is an interesting research problem. In addition, the overall system solution must deal with the fact that no inference method is 100% correct. Consequently, uncertainty in interpreted data can easily cause users not to trust the system.

Trust is at the crux of next generation WSAN technology. Security and privacy are essential elements of trust, and these are discussed in their own sections. However, as a basis for trust, it is also necessary to develop new in-field sensor calibration techniques and reliable transport protocols [1], [2]. Without these basic underlying system-level capabilities, further inference might be operating with wrong or too much missing data, resulting in wrong conclusions. If these wrong conclusions drive actuators, then serious safety problems can occur. One approach is to ensure that all inferred information is accompanied by a confidence level in the form of a probability that the information is correct or incorrect [3] and to use that information to guarantee safe actuator operation. In many applications, informing users how information was derived is necessary. Another main challenge is making good (control)

decisions using the created knowledge. However, in making decisions, it is necessary to minimize the number of false negatives and false positives and guarantee safety; otherwise, the system will be dismissed as unreliable.

Due to the expected pervasiveness of WSANs, many individuals will often be in the sensing area of the same sensors. It is necessary to perform correct data association ensuring that the collected data and subsequent inferences are associated with the correct individual or individuals. This is a very challenging problem for many situations. When users are wearing RFIDs or when cameras with pattern recognition are used, the problem is solved. However, in many other situations, it will be necessary to combine a set of current sensor readings with a trace of the recent past readings and utilize a history of a given user's activities and personal characteristics to arrive at an accurate data assignment. Very little work has been done on this problem.

2. Robust System Operation

Many applications in WSNs typically initialize themselves by self-organizing after deployment [4]. At the conclusion of the self-organizing stage, it is common for the nodes of a WSN to know their locations, have synchronized clocks, know their neighbors, and have a coherent set of parameter settings, such as consistent sleep/wake-up schedules, appropriate power levels for communication, and pair-wise security keys [5]. However, over time, these conditions can deteriorate. The most common (and simple) example of this deterioration problem is with clock synchronization. Over time, clock drift causes nodes to have different enough times to result in application failures. While it is widely recognized that clock synchronization must re-occur, this principle is much more general. For example, even in static WSANs, some nodes may be physically moved unexpectedly. More and more nodes may become out of place over time. To make system-wide node locations coherent again, node re-localization needs to occur (albeit at a much slower rate than for clock sync). This issue can be considered a form of entropy where a WSAN will deteriorate (tend towards disorder) unless energy in the form of re-running protocols and other self-healing mechanisms is applied. Note that control of actuators can also deteriorate due to their controlling software and protocols, but also due to physical wear and tear.

These types of required coherence (entropy) services must be combined with many other approaches to produce robust system operation. This includes formal methods to develop reliable code, in-situ debugging techniques [6], on-line fault tolerance [7], in-field-maintenance [8], and general health monitoring services [9]. These problems are exacerbated due to the unattended operation of the system, the need for a long lifetime, the openness of the systems, and the realities of the physical world. The goal is for this collection of solutions to create a robust system [10] in spite of noisy, faulty, and non-deterministic underlying physical world realities.

Another problem barely addressed to date is that, in some cases, run time assurances must be given to authorities, e.g., to (re)certify that the system is operating as expected. Consider a fire fighting WSAN deployed in a sky scraper office building to detect fires, alert fire stations and aid in evacuation. Periodically, it is necessary to demonstrate to certification authorities that this system meets these requirements. WSANs will need services that can support run time certification.

3. Openness

Traditionally, the majority of sensor-based systems have been closed systems. For example, cars, airplanes, and ships have had networked sensor systems that operate largely within those vehicles. However, these systems and other WSAN systems are expanding rapidly. Cars automatically transmit maintenance information, and airplanes send real-time jet engine information to manufacturers. WSANs will enable an even greater cooperation and 2-way control on a wide scale: cars (and aircraft) talking to each other and controlling each other to avoid collisions, humans exchanging data automatically when they meet and this possibly affecting their next actions, and physiological data uploaded to doctors in real-time with real-time feedback from the doctor. WSAN technology requires openness to achieve these benefits. However, supporting openness creates many new research problems. All of our current composition techniques, analysis techniques and tools need to be re-thought and developed to account for this openness. New unified communications interfaces will be required to enable efficient information exchange across diverse systems. Of course, openness also causes difficulty with security and privacy, the topics for the next two subsections. Consequently, openness must provide a correct balance between access to functionality and security and privacy.

To better illustrate some of the complexities involved with openness, consider feedback control. Most sensor systems heavily utilize feedback control theory to provide robust performance. The classical methodology includes creating a model of the system and then deriving a controller using well known techniques to meet stability, overshoot, settling time, and accuracy requirements. A sensitivity analysis is also possible and strongly encouraged. However, openness and scale create many difficulties for this methodology. Openness means that the model of the system is constantly changing. Human interaction is an integral aspect of openness, and this

makes modeling extremely difficult. The scaling and interactions across systems also dynamically change the models and create a need for decentralized control. While some work has been performed in areas such as stochastic, robust, distributed, and adaptive control, these areas are not developed well enough to support the degree of openness and dynamics expected in WSAN. A new and richer set of techniques and theory is required. It is especially important to understand how large numbers of control loops might interact with each other. To date, there have already been examples of WSANs in which control loops have competed with each other, one loop indicating an increase in a control variable and antoher loop indicating a decrease in the same variable at the same time. Such dependencies must be addressed in real-time and in an adaptive manner to support the expected openness of WSAN.

4. Security

A fundamental problem that must be solved in WSAN is dealing with security attacks [11]-[13]. Security attacks are problematic for WSANs because of the minimal capacity devices being used in parts of the systems, the physical accessibility to sensor and actuator devices, and the openness of the systems including the fact that most devices will communicate wirelessly. The security problem is further exacerbated because transient and permanent random failures are commonplace in WSANs, and failures are vulnerabilities that can be exploited by attackers. However, the considerable redundancy in WSANs creates great potential for designing them to continue to provide their specified services even in the face of failures. To meet realistic system requirements that derive from long-lived and unattended operation, WSN must be able to continue to operate satisfactorily in the presence of security attacks and to recover effectively from them. The system must also be able to adapt to new attacks that were unanticipated when the system was first deployed. These problems are beginning to be addressed by works such as [14]. In this work, WSAN technology operates with a base level of support including strong detection capabilities. Once an attack is detected, reaction to it occurs. This reaction is supported by the self-healing features of the system. This strategy works better than trying to build a completely secure system.

To heal from security attacks, a WSAN system needs to detect the attack, diagnose the attack, and deploy countermeasures and repairs; however, all of this must be performed in a lightweight manner due to the types of low capacity devices involved. Most of today's mainframe security solutions require heavyweight computations and large memory requirements, so solutions for WSAN are major research challenges. Ideally, for a quick response, given the real-time nature of WSANs, detection, countermeasures, and repairs must run in real-time as part of a runtime self-healing architecture. Sometimes, healing requires re-programming parts of the WSAN remotely, such as when an unanticipated attack occurs. In these cases, healing instructions need to be securely (with authentication and attestation) delivered to the appropriate nodes, and then the node's running programs need to be amended by the runtime architecture. Regardless of where the changes originate, the required changes may vary across the WSAN and within the node. Only nodes in a certain location (in the area being attacked) may need to activate a certain self-healing protocol and only for a certain interval of time (the duration of the attack). Solutions such as these do not yet exist for WSAN. It is likely that significant hardware support [15] will be necessary for providing encryption, authentication, attestation, and tamper proof keys.

5. Privacy

The ubiquity and interactions of WSAN will provide many conveniences and useful services for individuals but will also create many opportunities to violate privacy [16]. To solve the privacy problem created by single and interacting WSANs of the future, the privacy policies for each (system) domain must be specified. Once specified, WSAN systems must enforce privacy. Consequently, a system must be able to express users' requests for data access and the system's policies such that the requests can be evaluated against the policies in order to decide if they should be granted or denied. A new language is required to express privacy policies because WSANs have the following requirements not easily expressed in current privacy languages:

- Need to express the different types of context information in the environment such as time, space, physiological sensing, environmental sensing, and stream-based noisy data. Most of the context information needs to be collected and evaluated in real-time.
- Need to represent different types of data owners and request subjects in the system as well as external users and their rights when domains interact. Unlike other privacy enforcing systems where the subjects and data owners are human individuals or groups, a WSAN privacy language might also support physical entities such as "room", "floor", and other system entities as request issuers and data owners.
- Need to represent high-level aggregating requests, such as querying the average, maximum, or minimum reading of specified sensing data. This capability must be supported by anonymizing aggregation functions.
- Need to support not only adherence to privacy for queries

of data (pulling data value from the system), but also privacy on requests to set a system's parameters (pushing new values to the system), such as a private use of an actuator.

- Need to allow dynamic changes to the policies and perform a myriad of analyses some of which are context dependent.

One of the more difficult privacy problems is that systems may interact with other systems, each having their own privacy policies. Consequently, inconsistencies may arise across systems. Once again, on-line consistency checking and notification and resolution schemes are required.

Because WSANs monitor and control a large variety of physical parameters in different contexts, it is necessary to permit a high degree of dynamics and possibly even temporary privacy violations in order to meet functional, safety, or performance requirements. For example, an individual wearing an EKG might experience heart problems, and the real-time reporting of this problem takes precedence over privacy. To send an alarm quickly, it may be necessary to skip privacy protections. If this occurs, at a minimum users need to be notified whenever such a privacy violation has occurred.

V. Related Works

This paper focuses on future applications and research. For readers interested in the individual open research questions discussed, those particular sections contain references to related works. This section presents several additional comprehensive papers related to WSAN.

In [17], there is an excellent survey of the issues relating to control and communications for WSAN. This includes wireless communication issues. A key application domain for WSAN is surveillance. An information based approach for sensor management in evolving environments is found in [18]. Another survey and projection of future research needs for WSAN is given in [19]. Finally, a new text, [20], presents a nice overview of the technologies, analysis, and design issues for WSAN.

VI. Conclusion

When WSANs cover the world, a new revolution similar to the Industrial and Internet revolutions will occur. Applications such as those discussed in this paper and many others not even conceived of today will appear. New jobs, industries, and economic models will emerge. Daily life will be changed in profound ways. Unanticipated legislation will also affect the progress and direction of the impact of this technology.

Optimistically, new research results will enable openness, robustness, security, and privacy to co-exist—not an easy task! Knowledge will be produced by a vast collection of real-time sensors and inference software. This knowledge will occur at an ever increasing rate and with proper focusing and filtering will enable beneficial consequences for individuals, corporations, and governments.

Note that, in this short paper, only a few very important research areas were discussed. Many other important topics targeting WSAN must also be addressed including the following: heterogeneity, standards, programming abstractions and languages, real-time stream databases, middleware, operating systems, scaling, composition theory and analysis, formal methods, the wireless spectrum, wireless realities including interference, real-time, system safety, design, analysis and debugging tools, energy scavenging and power control, mobility, time synchronization, location services, decentralized algorithms, swarm computing, and signal processing. Simultaneously addressing several of these issues in the context of WSAN will produce many interesting research problems.

References

[1] A. Srinivasan, J. Teitelbaum, J. Wu, "DRBTS: Distributed Reputation-Based Beacon Trust System," *2nd IEEE Int'l Symp. Dependable, Autonomic, and Secure Computing* (DASC), 2006, pp. 277-283.

[2] A. Woo, T. Tong, and D. Culler, "Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks," *Proc. the 1st ACM International Conf. Embedded Networked Sensor Systems* (Sensys), 2003, pp. 14-27.

[3] B. Przydatek, D. Song, and A. Perrig, "SIA: Secure Information Aggregation in Sensor Networks," *Proc. the 1st ACM Int'l Conf. Embedded Networked Sensor Systems* (Sensys), 2003, pp. 255-265.

[4] A. Cerpa and D. Estrin, "ASCENT: Adaptive Self-Configuring Sensor Networks Topologies," *Proc. the IEEE Infocom*, 2002, pp. 1278-1287.

[5] H. Chan, A. Perrig, and D. Song, "Random Key Predistribution Schemes for Sensor Networks," *IEEE Symp. Research in Security and Privacy*, 2003, pp. 197-213.

[6] N. Ramanathan et al., "Sympathy for the Sensor Network Debugger," *Proc. the 3rd ACM Int'l Conf. Embedded Networked Sensor Systems (Sensys)*, 2005, pp. 255-267.

[7] L. Paradis and Q. Han, "A Survey of Fault Management in Wireless Sensor Networks," *Journal of Network and Systems Management*, vol. 15, no. 2, June 2007, pp. 171-190.

[8] Q. Cao and J. Stankovic, "An In-Field Maintenance Framework for Wireless Sensor Networks," *DCOSS*, June 2008, pp. 457-468.

[9] S. Rost and H. Balakrishnan, "Memento: A Health Monitoring

System for Wireless Sensor Networks," *Proc. IEEE SECON*, 2006, pp. 575-584.

[10] L. Gu and J.A. Stankovic, "t-kernel: Providing Reliable OS Support for Wireless Sensor Networks," *Proc. ACM Conf. on Embedded Networked Sensor Systems (Sensys)*, 2006, pp. 1-14.

[11] S. Capkun and J.-P. Hubaux, "Secure Positioning of Wireless Devices with Application to Sensor Networks," *Proc. IEEE Infocom*, 2005, pp. 1917-1928.

[12] A. Perrig, D. Wagner, and J. Stankovic, "Security in Wireless Sensor Networks," *CACM*, vol. 47, no. 6, June 2004, pp. 53-57.

[13] A. Wood and J. Stankovic, "Denial of Service in Sensor Networks," *IEEE Computer*, vol. 35, no. 10, Oct. 2002, pp. 54-62.

[14] A. Wood et al., "SIGF: A Family of Configurable, Secure Routing Protocols for Wireless Sensor Networks," *ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, 2006, pp. 35-48.

[15] D. Arora et al., "Secure Embedded Processing through Hardware-Assisted Run-Time Monitoring," *Proc. the Design, Automation and Test in Europe Conference and Exhibition (DATE)*, 2005, pp. 178-183.

[16] P. Kamat et al., "Enhancing Source Location Privacy in Sensor Network Routing." *Proc. Int'l Conf. Distributed Computing Systems*, 2005, pp. 559-608.

[17] J. Baillieul and P. Antsaklis, "Control and Communication Challenges in Networked Control Systems," *Proc. the IEEE*, vol. 95, no. 1, Jan. 2007, pp. 9-28.

[18] C. Kreucher et al., "An Information-Based Approach to Sensor Management in Large Dynamic Networks," *Proc. the IEEE*, vol. 95, no. 5, May 2007, pp. 978-999.

[19] J. Stankovic et al., "Opportunities and Obligations for Physical Computing Systems," *IEEE Computer*, vol. 38, no. 11, Nov. 2005, pp. 23-31.

[20] R. Verdone et al., *Wireless Sensor and Actuator Networks: Technologies, Analysis and Design*, Academic Press, Jan. 2008.

Professor John A. Stankovic is the BP America Professor in the Computer Science Department at the University of Virginia. He served as Chair of the department for 8 years. He is a Fellow of both the IEEE and the ACM. He won the IEEE Real-Time Systems Technical Committee's Award for Outstanding Technical Contributions and Leadership (inaugural winner). He also won the IEEE Technical Committee on Distributed Processing's Distinguished Achievement Award (inaugural winner). He has won three Best Paper awards, including one for ACM SenSys 2006. He is ranked among the top 250 highly cited authors in CS by Thomson Scientific Institute. He has given many Keynote talks at conferences and in Distinguished Lecture series at major Universities. Professor Stankovic also served on the Board of Directors of the Computer Research Association for 9 years. Before joining the University of Virginia, Professor Stankovic taught at the University of Massachusetts where he won an outstanding scholar award. He has also held visiting positions in the Computer Science Department at Carnegie-Mellon University, at INRIA in France, and Scuola Superiore S. Anna in Pisa, Italy. He was the Editor-in-Chief for the IEEE Transactions on Distributed and Parallel Systems and was founder and co-editor-in-chief for the Real-Time Systems Journal. He was General and Program Chair for many conferences including ACM Sensys 2004 and ACM/IEEE IPSN 2006. His research interests are in distributed computing, real-time systems, operating systems, and wireless sensor networks. He has built three sensor networks: VigilNet, a military surveillance system funded by Darpa and now being constructed by Northrup-Gruman, Luster, an environmental science system for measuring the effect of sunlight on plant growth, and AlarmNet, an emulation of as assisted living facility. Prof. Stankovic received his PhD from Brown University.