Research Directions for the Internet of Things

John A. Stankovic, *Life Fellow, IEEE*

***Abstract*—Many technical communities are vigorously pursuing research topics that contribute to the Internet of Things (IoT). Today, as sensing, actuation, communication, and control become ever more sophisticated and ubiquitous, there is significant overlap in these communities, sometimes from slightly different perspectives. More cooperation between communities is encouraged. To provide a basis for discussing open research problems in IoT, a vision for how IoT could change the world in the distant future is first presented. Then, eight key research topics are enumerated and research problems within those topics are discussed.**

***Index Terms*—Cyber Physical Systems, Internet of Things, Mobile Computing, Pervasive Computing, Wireless Sensor Networks.**

I. Introduction

Smart devices. Smartphones. Smart cars. Smart homes. Smart cities. A smart world. These notions have been espoused for many years. Achieving these goals has been investigated, to date, by many diverse and often disjoint research communities. Five such prominent research communities are: Internet of Things (IoT), Mobile Computing (MC), Pervasive Computing (PC), Wireless Sensor Networks (WSN), and most recently, Cyber Physical Systems (CPS). However, as technology and solutions progress in each of these fields there is an increasing overlap and merger of principles and research questions. Narrow definitions of each of these fields are no longer appropriate. Further, research in IoT, PC, MC, WSN and CPS often relies on underlying technologies such as real-time computing, machine learning, security, privacy, signal processing, big data, and others. Consequently, the *smart vision of the world* involves much of computer science, computer engineering, and electrical engineering. Greater interactions among these communities will speed progress.

In this paper, as a backdrop to identifying research questions, I briefly highlight a vision for a smart world (Section II). I then discuss open research questions categorized into 8 topics (Section III). The research discussed is representative rather than complete. Two goals of the paper are: (i) to highlight a number of significant research needs for future IoT systems, and (ii) to raise awareness of work being performed across various research communities.

This work was supported in part by the National Science Foundation under grants CNS-1239483, CNS-1017363, and CNS-1319302.

John A. Stankovic is the BP America Professor in the Computer Science Department at the University of Virginia, Charlottesville, Virginia, 22904. USA. (e-mail: stankovic@ cs.virginia.edu).

II. Vision and IoT Scope

Many people [8], including myself [28][29], hold the view that cities and the world itself will be overlaid with sensing and actuation, many embedded in "things" creating what is referred to as a *smart world*. But it is important to note that one key issue is the degree of the density of sensing and actuation coverage. I believe that there will be a transition point when the degree of coverage triples or quadruples from what we have today. At that time there will be a *qualitative* change. For example, today many buildings already have sensors for attempting to save energy [7][38]; home automation is occurring [3]; cars, taxis, and traffic lights have devices to try and improve safety and transportation [9]; people have smartphones with sensors for running many useful apps [2]; industrial plants are connecting to the Internet [1]; and healthcare services are relying on increased home sensing to support remote medicine and wellness [11]. However, all of these are just the tip of the iceberg. They are all still at early stages of development. The steady increasing density of sensing and the sophistication of the associated processing will make for a significant *qualitative change* in how we work and live. We will truly have systems-of-systems that synergistically interact to form totally new and unpredictable services.

What will be the platform or platforms that support such a vision? One possibility is a global sensing and actuation utility connected to the Internet. Electricity and water are two utilities that can be used for a myriad of purposes. Sensing and actuation in the form of an IoT platform will become a utility. IoT will not be seen as individual systems, but as a critical, integrated infrastructure upon which many applications and services can run. Some applications will be personalized such as digitizing daily life activities, others will be city-wide such as efficient, delay-free transportation, and others will be worldwide such as global delivery systems. In cities perhaps there will be no traffic lights and even 3D transportation vehicles. Smart buildings will not only control energy or security, but integrate personal comfort, energy savings, security and health and wellness aspects into convenient and effective spaces. Individuals may have patches of bionic skin with sensing of physiological parameters being transmitted to the cloud which houses his digital health, and to the

surrounding smart spaces for improved comfort, health, efficiency, and safety. In fact, smart watches, phones, body nodes, and clothes will act as personalized input to optimize city-wide services benefiting both the individual and society. Consequently, we will often (perhaps 24/7) be implicitly linked into the new utility. Some examples of new services include immediate and continuous access to the right information for the task at hand, be it, traveling to work or a meeting, exercising, shopping, socializing, or visiting a doctor. Sometimes these activities will be virtual activities, or even include the use of avatars or robots. Many outputs and displays for users may be holographic. Credit cards should disappear and biometrics like voice or retinas will provide safe access to buildings, ATMs, and transportation systems.

A sensing and actuation utility will not only exist in public spaces, but also extend into the home, apartments, and condominiums. Here people will be able to run health, energy, security, and entertainment apps on the infrastructure. Installing and running new apps will be as easy as plugging in a new toaster into the electric utility. One app may help monitor and control heart rate, another perform financial and investments services, another automatically ordering food and wine, or even predicting a impending medical problem that should be addressed early to mitigate or even avoid the problem. Humans will often be integral parts of the IoT system. The Industrial Internet is also a form of IoT where the devices (things) are objects in manufacturing plants, dispatch centers, process control industries, etc. Consequently, in the future the scope of IoT is enormous and will affect every aspect of all our lives.

III. Research

The spectrum of research required to achieve IoT at the scale envisioned above requires significant research along many directions. In this section problems and required research are highlighted in 8 topic areas: massive scaling, architecture and dependencies, creating knowledge and big data, robustness, openness, security, privacy, and human-in-the-loop. Each of the topic discussions primarily focuses on new problems that arise for future IoT systems of the type described in Section II. The research topics presented in each case are representative and not complete.

Many important topics such as the development of standards, the impact of privacy laws, and the cultural impact on use of these technologies are outside the scope of the paper.

A. Massive Scaling

The current trajectory of the numbers of smart devices being deployed implies that eventually trillions of things will be on the Internet. How to name, authenticate access, maintain, protect, use, and support such a large scale of things are major problems. Will IPv6 suffice? Will protocols such a 6LowPAN play a role? Will entirely new standards and protocols emerge? Since many of the things on the Internet will require their own energy source, will energy scavenging and enormously low power circuits eliminate the need for batteries? How will the massive amounts of data be collected, used, and stored? What longitudinal studies will be performed? How will the real-time and reliability aspects be supported [5][13]? How will devices including mobile devices be discovered? Will the emergence of a utility model, if it occurs, mean entirely new standards? How will such a utility be achieved? It is unlikely that any solution immediately becomes the norm. Many protocols and variations will co-exist. What will be the architectural model that can support the expected heterogeneity of devices and applications?

B. Architecture and Dependencies

As trillions of things (objects) are connected to the Internet it is necessary to have an adequate architecture that permits easy connectivity, control, communications, and useful applications. How will these objects interact in and across applications [37]? Many times, things or sets of things must be disjoint and protected from other devices. At other times it makes sense to share devices and information. One possible architectural approach for IoT is to borrow from the smartphone world [2][4]. Smartphones employ an approach where applications are implemented and made available from an app store. This has many advantages including an unbounded development of novel applications that can execute on the smartphones. Various standards and automatic checks are made to ensure that an app can execute on a given platform. For example, the correct version of the underlying OS and the required sensors and actuators can be checked when the app is installed [12]. A similar architectural approach for IoT would also have similar advantages. However, the underlying platform for IoT is much more complicated than for smartphones. Nevertheless, if IoT is based on an underlying sensor and actuator network that acts as a utility similar to electricity and water, then, different IoT applications can be installed on this utility. While each application must solve its own problems, the sharing of a sensing and actuation utility across multiple simultaneously running applications can result in many systems-of-systems interference problems, especially with the actuators. Interferences arise from many issues, but primarily when the cyber depends on assumptions about the environment, the hardware platform, requirements, naming, control and various device semantics. Previous work, in general, has considered relatively simple dependencies related to numbers and types of parameters, versions of underlying operating systems, and availability of correct underlying hardware. Research is needed to develop a comprehensive approach to specifying, detecting, and resolving dependencies across applications. This is especially important for safety critical applications or when actuators can cause harm.

Let's consider a few examples of dependencies [21][31][32]. Assume that we integrate several systems responsible for energy management (controlling thermostats [17], windows, doors, and shades) and home health care (controlling lights, TVs, body nodes measuring heart rate and temperature, and sleep apnea machines [33]). If information can be shared, this would allow the energy management system to adjust room temperature depending on the physiological status of the

residents as detected by the home health care system. Also, integration will allow avoiding negative consequences. For example, the integrated system will not turn off medical appliances to save energy while they are being used as suggested by the home health care system. In addition to these advantages, all the systems can share sensors and actuators, which will reduce cost of deployment, improve aesthetics of the rooms, and reduce channel contention. However, integrating multiple systems is very challenging as each individual system has its own assumptions and strategy to control the physical world variables without much knowledge of the other systems, which leads to conflicts when these systems are integrated without careful consideration. For example, a home health care application may detect depression and decide to turn on all the lights. On the other hand, the energy management application may decide to turn off lights when no motion is detected. Detecting and resolving such dependency problems is important for correctness of operation of interacting IoT systems.

C. *Creating Knowledge and Big Data*

In an IoT world there will exist a vast amount of raw data being 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. 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 be used, but may also be too ad hoc for some applications.

The amount of collected data will be enormous. It can be expected that a very large number of real-time sensor data streams will exist, that it will be common for a given stream of data to be used in many different ways for many different inference purposes, that the data provenance and how it was processed must be known, and that privacy and security must be applied. Data mining techniques are expected to provide the creation of important knowledge from all this data. 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 one important aspect of the usefulness of big data. 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. 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 and 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.

Many IoT applications will be designed to work for a particular person. 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 then the problem is solved (except for the privacy issues). 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. More research is necessary on this problem.

D. *Robustness*

If our vision is correct, many IoT applications will be based on a deployed sensing, actuation, and communication platform (connecting a network of things). In these deployments it is common for the devices to know their locations, have synchronized clocks, know their neighbor devices when cooperating, 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. However, over time these conditions can deteriorate. The most common (and simple) example of this deterioration problem is with clock synchronization [18]. 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, 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 system will deteriorate (tend towards disorder) unless energy in the form of re-running protocols and other self-healing mechanisms is applied [35]. Note that control of actuators can also deteriorate due to their controlling software and protocols, but also due to physical wear and tear. In other words, how can a long-lived, dynamic, and mobile IoT be maintained?

The required coherence (entropy) services must combine with many other approaches to produce robust system operation. This includes formal methods to develop reliable code, in-situ debugging techniques, on-line fault tolerance, in-field-

maintenance, and general health monitoring services [23][24][25]. 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 in spite of noisy, faulty and non-deterministic underlying physical world realities.

Another problem barely addressed to date is that in some IoT applications, especially safety critical ones, run time assurances must be given to authorities, e.g., to (re)certify that the system is operating as expected. Consider a fire fighting system 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. Such IoT applications will need services that can support run time certification.

E. 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 that vehicle. However, these systems' capabilities are expanding rapidly. Cars are automatically transmitting maintenance information and airplanes are sending real-time jet engine information to manufacturers. There is or will be 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. These systems require 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. Many sensor and actuator 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. The openness means that the model of the system is constantly changing. The human interaction is an integral aspect of openness (see Section III H.) and this makes modeling extremely difficult, and the scaling and interactions across systems also dynamically change the models and creates a need for decentralized control. While some work has been performed in topics such as stochastic control, robust control, distributed control and adaptive control, these areas are not developed well enough to support the degree of openness and dynamics expected in some IoT sytems. 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 where control loops have competed with each other, one indicating an increase in a control variable while the other loop indicating a decrease in the same variable at the same time. Such dependencies (see Section III B.) must be addressed in real-time and in an adaptive manner to support the expected openness of IoT.

Openness is also playing a major role in industrial things on the Internet. Remote access across factories or to individual products is often very beneficial to Industry. However, security concerns arise, especially if there is any safety issue involved.

F. Security

A fundamental problem that is pervasive in the Internet today that must be solved is dealing with security attacks [22] [36]. Security attacks are problematic for the IoT because of the minimal capacity "things" (devices) being used, the physical accessibility to sensors, actuators and objects, 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 and failures are vulnerabilities that can be exploited by attackers. However, the considerable redundancy that is available creates potential for designing applications 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, IoT applications must be able to continue to operate satisfactorily in the presence of, and to recover effectively from security attacks. Solutions may require downloading new code [10] and this itself is open to security attacks. The system must also be able to adapt to new attacks unanticipated when the system was first deployed. These problems are beginning to be addressed by work such as that found in [34]. In [34], the system operates with a base level of support including strong attack detection capabilities. Once an attack is detected then reaction to it occurs, by self-healing.

To heal from security attacks, a system needs to detect the attack, diagnose the attack, and deploy countermeasures and repairs, but perform all of this 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 IoT are major research challenges. Ideally, for a quick response, given the real-time nature of many IoTs, the detection, countermeasures and repairs must run in real-time as part of a runtime self-healing architecture. Sometimes, healing requires re-programming, e.g., 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. It is likely that significant hardware support [22] will be necessary for providing encryption, authentication, attestation, and tamper proof keys. Even if new devices are security-aware, dealing with legacy devices will prove difficult.

G. Privacy

The ubiquity and interactions involved in IoT will provide many conveniences and useful services for individuals, but also create many opportunities to violate privacy. To solve the privacy problem created by IoT applications of the future, the privacy policies for each (system) domain must be specified. Once specified either the individual IoT application or the IoT infrastructure (e.g., the utility capability) must enforce privacy. Consequently, the IoT paradigm must be able to express users' requests for data access and the 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 the following requirements not *easily* expressed in current privacy languages:

- The need to express the different types of context in the environment such as time, space, physiological sensing, environmental sensing, and stream based noisy data. Most of the context needs to be collected and evaluated in real-time. But what will collect policies and data and support privacy? Is it the utility infrastructure, an individual application, both, or some new approach?
- The 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, an IoT privacy language might also support physical entities such as "refrigerator", "room", "floor", and other system entities (things) as request issuers and data owners.
- The 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. This capability needs to exist for real-time streams and across the big data repositories. Note that inference is very powerful and having access to vast amounts of data and inference techniques it is often easy to violate privacy in spite of anonymization.
- The 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), e.g., a private use of an actuator.
- The 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 interact with other systems, each having their own privacy policies. Consequently, inconsistencies may arise across systems in the IoT world. On-line consistency checking and notification and resolution schemes are required.

H. Humans in the Loop

As IoT applications proliferate they will become more sophisticated. Many of these new applications will intimately involve humans, i.e., humans and things will operate synergistically. Human in-the-loop systems offer exciting opportunities to a broad range of applications including energy management [17], health care [15], and automobile systems [9][16]. For example, it is hypothesized that explicitly incorporating human-in-the-loop models for driving can improve safety, and using models of activities of daily living in home health care can improve medical conditions of the elderly and keep them safe. Although having humans in the loop has its advantage, modeling human behaviors is extremely challenging due to the complex physiological, psychological and behavioral aspect of human beings. New research is necessary to raise human-in-the-loop control to a central principle in system design and to solve three main challenges [20].

Challenge 1: The need for a comprehensive understanding of the complete spectrum of types of human-in-the-loop controls.

There are many variations for human-in-the-loop controls. We need to understand the complete spectrum to determine the underlying principles and subtleties that separate them. Human-in-the-loop applications can be classified into four categories: (i) applications where humans directly control the system, (ii) applications where the system passively monitors humans and takes appropriate actions, (iii) applications where physiological parameters of the human are modeled, and (iv) hybrids of (i), (ii), and (iii). Applications where humans directly control the system primarily use supervisory control. In supervisory control, involvement of humans takes place in two ways. In one case, the process runs autonomously. Humans intervene with the control algorithm when it is necessary typically by adjusting set points. These control problems are well understood. In the second case, the behaviors of a human are observed, e.g., eating behaviors, and interventions are controlled to improve their life. In the third case, the process accepts a command, carries out the command autonomously, reports the results and waits for further commands to be received from the human. As a concrete example, in [30], human-in-the-loop control is used in a wheelchair-mounted robotic arm to retrieve an object from a shelf. In this feedback control system, human provides input via a touch screen or joystick which is analyzed by a vision processing system to position the robotic arm to retrieve the object. In this application, a human directly controls the controller of the feedback control system and guides it to take appropriate action. Applications such as [19] are similar.

Challenge 2: The need for extensions to system identification or other techniques to derive models of human behaviors.

System identification is a powerful technique to create system models. It is a new challenge to apply it to human behaviors. The order and types of equations to use, how to produce adequate testing inputs, what output variables are required, and how such a model accounts for human traits are unknown. If we were to use system identification technique to model a human being who is suffering from depressive illness, it is not clear what are the inputs, what are the states and how the state transitions occur based on different physiological, psychological and environmental factors. If there was a formal model of human behavior or even an estimated model, then by combining all the factors that affect depression, we could close the loop by changing the factors in a way that helps the patients and that is based on an established methodology rather than ad hoc rules. Clustering, data mining, inference, first principle models based on human physiology [14] and behaviors may all be necessary techniques to be enhanced and applied for different applications. Robust systems will likely require predictive models to avoid problems before they occur. Advances to stochastic model predictive control are also required. It is also unlikely that any models developed initially to design the controllers will remain accurate as the system and human behaviors evolve over time. Hence, adaptive control with humans-in-the-loop will be necessary.

Challenge 3: Determining how to incorporate human behavior models into the formal methodology of feedback control.

In the formal methodology of feedback control there are several areas where a human model can be placed:

- Outside the loop,
- Inside the controller,
- Inside the system model,
- Inside a transducer, and
- At various levels in hierarchical control.

The newest challenge seems to be how to incorporate the human behavior **as part of the system itself**. Can we define/guarantee/learn the stability, accuracy, settling time and overshoot properties of such systems, initially and as the system and human behavior evolves? As an example, [6] proposes a procedure to refine user behavior models based on reports of accidents and incidents that occur during the operation of electrical power system. This work mainly focuses on using Components Model of Emotion (CME) for observing, recording and analyzing the emotional components of the operator behavior, which can be eventually useful for simulating dynamic behavior of an operator performing tasks in a context that leads to an error. If we can model such an operator behavior using formal methodology of feedback control and if we can incorporate these operator models into the system, we will be able to analyze various safety properties of the overall system. See also [26] for additional descriptions of human-in-the-loop systems.

IV. Conclusion

In summary, one vision of the future is that IoT becomes a utility with increased sophistication in sensing, actuation, communications, control, and in creating knowledge from vast amounts of data. This will result in qualitatively different lifestyles from today. What the lifestyles would be is anyone's guess. It would be fair to say that we cannot predict how lives will change. We did not predict the Internet, the Web, social networking, Facebook, Twitter, millions of apps for smartphones, etc., and these have all qualitatively changed societies' lifestyle. New research problems arise due to the large scale of devices, the connection of the physical and cyber worlds, the openness of the systems of systems, and continuing problems of privacy and security. It is hoped that there is more cooperation between the research communities in order to solve the myriad of problems sooner as well as to avoid re-inventing the wheel when a particular community solves a problem.

Acknowledgment

Some of the material in this invited paper is taken from my previous papers. Some of section II is from [29] which is available only in Chinese. Parts of Section III B. are from [21]. Parts of section III C, D, E and F are from [27]. Finally, some of III H. is from [20].

References

[1] M1 Security and Automation Controls. http://www.elkproducts.com/m1 controls.html.

[2] Apple app store. http://www.apple.com/osx/apps/app-store.html.

[3] Control4 Home Automation and Control. http://www.control4.com.

[4] http://www.phonearena.com/news/Androids-Google-Playbeats- App-Store-with-over-1-million-apps-now-o_ciallylargestid45680.

[5] T. Abdelzaher, S. Prabh, and R. Kiran, On Real-Time Capacity Limits of ad hoc Wireless Sensor Networks, *RTSS*, December 2004.

[6] Y. Aguiar, M. Vieira, E. Galy, J. Mercantini, and C. Santoni, Refining a User Behavior Model based on the Observation of Emotional States. *COGNITIVE* , 2011.

[7] V. Bradshaw. The Building Environment: Active and Passive Control Systems. John Wiley & Sons, Inc., River Street, NJ, USA, 2006.

[8] B. Brumitt, B. Meyers, J. Krumm, A. Kern, and S. A. Shafer. Easyliving: Technologies for Intelligent Environments. *HUC*, 2000.

[9] G. Burnham, J. Seo G. Bekey, A. Identification of Human Driver Models in Car Following. *IEEE Transactions on Automatic Control 19*, 6, 1974, pp. 911–915.

[10] J. Deng, R. Han, and S. Mishra, Secure Code Distribution in Dynamically Programmable Wireless Sensor Networks, *Proc. of ACM/IEEE IPSN*, 2006. pp. 292-300.

[11] R. Dickerson, E. Gorlin, and J. Stankovic, Empath: a Continuous Remote Emotional Health Monitoring System for Depressive Illness. *Wireless Health* , 2011.

[12] C. Dixon, R. Mahajan, S. Agarwal, A. Brush, B. Lee, S. Saroiu, and P. Bahl, An Operating System for the Home, *NSDI*, 2012.

[13] T. He, J. Stankovic, C. Lu and T. Abdelzaher, A Spatiotemporal Communication Protocol for Wireless
Sensor Networks, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 16, No. 10, Oct. 2005, pp. 995-1006.

[14] M. Huang, J. Li, X. Song, and H. Guo, Modeling Impulsive Injections of Insulin: Towards Artificial Pancreas. *SIAM Journal of Applied Mathematics 72*, 5, 2012, pp. 1524–1548.

[15] M. Kay, E. Choe, J. Shepherd, B. Greenstein, N. Watson, S. Consolvo, and J. Kientz, Lullaby: a Capture & Access System for Understanding the Sleep Environment. *UbiComp*, 2012.

[16] A Liu, and D. Salvucci, Modeling and Prediction of Human Driver Behavior, *Intl. Conference on HCI* , 2001.

[17] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben J. Stankovic, E. Field, and K. Whitehouse, The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes, *ACM SenSys*, 2010.

[18] M. Maroti, B. Kusy, G. Simon, and A. Ledeczi, The Flooding Time Synchronization Protocol, *ACM SenSys*, November 2004.

[19] S. Mohammed, P. Fraisse, D. Guiraud, P. Poignet, and H. Makssoud, Towards a Co-contraction Muscle Control strategy for Paraplegics. *CDC-ECC*, 2005.

[20] S. Munir, J. Stankovic, C. Liang, and S. Lin, New Cyber Physical System Challenges for Human-in-the-Loop Control, *8th International Workshop on Feedback Computing*, June 2013.

[21] S. Munir and J. Stankovic, DepSys: Dependency Aware Integration of Systems for Smart Homes, submitted for publication.

[22] S. Ravi, A. Raghunathan, S. Chakradhar. Tamper Resistance Mechanisms for Secure, Embedded Systems, *Proc. of 17th International Conference on VLSI Design*, 2004. p. 605.

[23] B. Rong Chen, G. Peterson, G. Mainland, and M. Welsh, LiveNet: Using Passive Monitoring to Reconstruct Sensor Network Dynamics, *DCOSS 2008*, June 2008.

[24] S. Rost and H. Balakrishnan, Memento: A Health Monitoring System for Wireless Sensor Networks. *SECON 2006*, September 2006.

[25] L. Ruiz, J. Nogueira, and A. Loureiro, MANNA: A Management Architecture for Wireless Sensor Networks. *IEEE Communications Magazine*, February 2003.

[26] G. Schirner, D. Erdogmus, K. Chowdhury, and T. Padir, The Future of Human-in-the-Loop Cyber-Physical Systems. *Computer 46*, 1, 2013, pp. 36–45.

[27] J. Stankovic, When Sensor and Actuator Networks Cover the World, invited Keynote Article, Special Issue on Ubiquitous Sensor Networks, *ETRI Journal*, Korea, Vol. 30. No. 5, October 2008, pp. 627-633.

[28] J. Stankovic, I. Lee, A. Mok, R. Rajkumar, Opportunities and Obligations for Physical Computing Systems, *IEEE Computer*, Vol. 38, No. 11, Nov. 2005, pp. 23-31.

[29] J. Stankovic, A Vision of a Smart City in the Future, *Smart Cities*, Vol. 1, Issue 10, Oct. 2013.

[30] K. Tsui, D. Kim, A. Behal, D. Kontak, and H. Yanco, I Want That : Human-in-the-Loop Control of a Wheelchair-Mounted Robotic Arm. *Journal of Applied Bionics and Biomechanics 8*, 2011.

[31] P. A. Vicaire, E. Hoque, Z. Xie, and J. A. Stankovic, Bundles: a Group Based Programming Abstraction for Cyber Physical Systems, *ICCPS*, 2010.

[32] P. A. Vicaire, Z. Xie, E. Hoque, and J. A. Stankovic, Physicalnet: A Generic Framework for Managing and Programming across Pervasive Computing Networks, *RTAS*, 2010.

[33] A. Wood, J. Stankovic, G. Virone, L. Selavo, T. He, Q. Cao, T. Doan, Y. Wu, L. FANG, and R. Stoleru, Context-Aware Wireless Sensor Networks for Assisted Living and Residential Monitoring. *IEEE Network 22*, 4, Jul./Aug. 2008, pp. 26–33.

[34] A. Wood, L. Fang, J. Stankovic, and T. He, SIGF: A Family of Configurable, Secure Routing Protocols for Wireless Sensor Networks, *ACM Security of Ad Hoc and Sensor Networks*, Best Paper Award, October 31, 2006.

[35] Y. Wu, J. Li, J. Stankovic, K. Whitehouse, S. Son and K. Kapitanova, Run Time Assurance of Application-Level Requirements in Wireless Sensor Networks, *IPSN*, April 2010.

[36] W. Xu, W. Trappe, Y. Zhang, T. Wood, The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks, *Proc. of MobiHoc*, 2005. pp. 46-57.

[37] Y. Yu, L. J. Rittle, V. Bhandari, and J. B. LeBrun, Supporting Concurrent Applications in Wireless Sensor Networks, *ACM SenSys*, 2006.

[38] W. Zeiler, R. Houten, G. Boxem, D. Vissers, and R. Maaijen, Indoor Air Quality and Thermal Comfort Strategies: the Human-in-the-Loop Approach, *ICEBO*, 2011.

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. He also won the IEEE Technical Committee on Distributed Processing's Distinguished Achievement Award (inaugural winner). He has six Best Paper awards, including one for ACM SenSys 2006. He also has a Best Paper Runner Up Award for IPSN 2013. Stankovic has an h-index of 98 and over 34,000 citations. He has also won Distinguished Faculty Awards at the University of Massachusetts and at the University of Virginia. He has given more than 30 Keynote talks at conferences and many Distinguished Lectures at major Universities. Currently, he serves on the National Academy's Computer Science Telecommunications Board. 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. His research interests are in real-time systems, distributed computing, wireless sensor networks, wireless health, and cyber physical systems. Prof. Stankovic received his PhD from Brown University.