Pervasive Computing at Scale: Transforming the State of the Art

Diane J. Cook
Washington State University
Pullman, WA 99164-2752
cook@eecs.wsu.edu

Sajal K. Das
University of Texas at Arlington
Arlington, TX 76205
das@uta.edu

Abstract

The remarkable recent progress in computing power, sensors and embedded devices, smart phones, wireless communications and networking technologies, combined with emerging data mining techniques, cloud computing and social networking paradigms have enabled us to create pervasive computing systems and services with diverse applications and global accessibility. In this paper we assess the current state of the art in of pervasive computing at scale (PeCS) and look ahead to future directions the field can pursue together with challenges it will need to overcome.

Keywords: smart environments, energy harvesting, cloud computing, smart phones, behavior modeling, internet of things

1. Introduction

The remarkable recent growth in computing power, sensors and embedded devices, smart phones, wireless communications and networking combined with the power of data mining techniques and emerging support for cloud computing and social networks have enabled researchers and practitioners to create a wide variety of pervasive computing systems that reason intelligently, act autonomously, and respond to the needs of the users in a context- and situation-aware manner. The field of pervasive computing is at an interesting and critical point in its development. On the one hand, the field has matured to the point where tangible, beneficial prototype testbeds such as smart homes, body area networks, health monitoring systems, and mobile social networking media are becoming fairly commonplace. These visible successes are built on mature underlying technology that performs smart device communications, resource discovery, information fusion, dissemination and routing, location tracking, activity recognition, and learning of user preferences. On the other hand, however, these systems have been mostly designed and tested on small to medium-scale applications with limited dissemination of the tools, results, and datasets.

This year is the 21st anniversary of Mark Weiser's landmark paper on ubiquitous computing [47]. While there has been significant progress toward his vision, most research has focused on the development of small-scale pervasive systems, tested by a handful of users, interacting with a limited number (e.g., tens or at most hundreds) of devices. In order to advance the field and make technology truly pervasive, the research community needs to address the issue of scale. The good news is that the trajectory from small to massive scale pervasive computing systems is underway. However, future large-scale pervasive systems still need to operate over different spatial and temporal scales and encompass a large number of computational platforms, users, devices and applications dealing with massive amounts of data. Diverse devices such as smart phones, tablets, laptops, desktops, wearables, RFIDs, and embedded wireless sensors on the order of thousands, millions and even billions as envisioned in the Internet of Things (IoT) will enable a wide spectrum of applications from predicting traffic jams and modeling human activities, to facilitating social interactions, tracking community health trends, and responding to disasters [11]. They will handle very large amounts of data distributed over heterogeneous networking and computing platforms (e.g., clouds, data centers), and support 100s of millions of users, including those of mobile phones and social media.

Sponsored by the US National Science Foundation, in January of 2011, a group of 72 researchers gathered to discuss challenges and issues for scaling future pervasive applications, architectures, algorithms, models, data and systems [34]. The problem of scaling pervasive systems is multi-disciplinary in nature, including challenges in human-computer or machine-to-machine interaction (HCI or MMI), machine learning, data mining, mobile systems, wireless and sensor networks, computing paradigms, smart environments, security and privacy, signal processing, information fusion, foundations (e.g., algorithms, stochastic control theory, information theory, game theory, optimization techniques), psychology, sociology and social networking. The workshop attendees represented expertise in these multi-disciplinary areas and brought perspectives from academia, industry, and funding agencies. This paper is based on the synergistic discussions, talks, panels and survey responses that took place at this workshop.

2. What is Pervasive Computing at Scale?

The goal of pervasive computing is to create ambient intelligence where networked devices embedded in the environment provide unobtrusive, continual, and reliable connectivity and also perform value added services. The result improves human experience and quality of life without explicit awareness of the underlying communications and computing technologies [11]. The field is closely related to smart environments in which computing and communications technologies employ artificial intelligence and machine learning techniques to reason about, control and adapt to our physical surroundings [10]. Cyber-physical systems, another related discipline which encompasses computer and information-centric physical and engineered systems as integration of communication, computation and control [31], may explore technologies outside of the human context. In contrast, pervasive computing necessarily focuses on sensing, interacting with and aiding humans at an individual and community level. While distributed and mobile computing supports information technologies such as remote information access and adaptive applications, pervasive computing extends this notion to provide computing and communication capabilities that are so gracefully integrated with users that it “disappears” [47]. Pervasive computing technologies are implicitly part of our everyday and social life and the environments with which we interact. While we are aware of the functionality they provide, we need not be aware of the underlying mechanisms by which that functionality is provided.

Because research has been directed toward core pervasive computing technologies and has done so with successful results over the last two decades, it is now appropriate to consider *pervasive computing at scale*, or PeCS. The notion of scalability here refers to the ability of a system to maintain some level of efficiency or functionality as the system dimensions increase. Generally, an increase in a system dimension adds capability to the system while incurring associated overheads. Capabilities and overheads can be measured by price tags, human time and attention, computation and communication power, storage capacity, accessibility, responsiveness, energy or other valuable resource usage. A pervasive computing system that is scalable provides a rate of increase in capability which is greater than the rate of increase in overhead; otherwise the overhead will eventually consume all resources, thus reducing the effective value added by the system to zero.

Not all pervasive computing applications are large scale or require large-scale resources and processing. However, as we look to the future of pervasive computing, ideas for large-scale use emerge and implementation of the ideas becomes more achievable. For example, while current pervasive computing systems have the ability to track individuals and analyze their behavioral patterns, future PeCS systems can scale to metropolitan area networks such as smart cities and smart communities that learn behavioral information and trends across a larger region. Similarly, current research is enabling smart vehicles, but future systems may scale to encompass an entire country's traffic system. Likewise, the Internet of Things (IoT) implies that every tagged object could be part of a very large-scale pervasively connected system across the globe. Finally, both IoT and (mobile) social networking have the potential to revolutionize as well challenge the scaling of pervasive systems.

Figure 1. The proliferation of pads, tabs, and boards.

Supporting PeCS will require vastly new approaches and paradigms to the design of hardware, networking, middleware, services and applications. When we envision pervasive computing at scale, the numbers of system dimensions that can increase are numerous. PeCS can scale to encompass a massive number of devices or an increasing heterogeneity of devices and the ways they communicate. The volume of data they generate, collect, store, transmit, and process will increase significantly. These systems can scale in the number and diversity of their applications and in the number of individuals and communities that utilize the systems. In order to scale effectively, PeCS must be able to increase these dimensions while maintaining or improving accuracy, efficiency, and reliability with cost effectiveness.

Figure 2. Pervasive computing devices are as diverse as currently lifestyles demand.

3. Where Does PeCS Currently Stand?

Computation and communication technology has evolved toward more pervasive and ubiquitous infrastructures over the past two decades. The PeCS community recognizes the wide-spread miniaturization and low-cost building of portable devices as well as myriads of applications for these devices. The current pervasive computing landscape includes massive numbers of portable devices (e.g., smart phones, “pads, tabs, and boards”) that gather and store information (see Figure 1). Devices are also increasingly diverse in their appearance, capability, portability, and use (see Figure 2).

Current mobile phones are as powerful as personal computers of old. The ability of these devices to collect and store information is well established. In addition, communication has become fast, fairly robust, and certainly pervasive. This is one reason why pervasive computing has already had an impact on the population in practice.

Another reason why the vision of pervasive computing is so powerful is that is reaches much larger masses than technology has in the past. As stated in the movie *The Social Network*, developing countries like Bosnia lack roads but “they have Facebook” [18]. Mobile phones are truly pervasive; they are accessible and reach around the world. High-speed Internet is out of reach for many low-income countries but mobile devices are truly ubiquitous. These devices provide accessibility to over 90% of the global population [25]. In 2009, 0.5 billion people accessed the Internet from mobile devices, and this number is expected to double by 2015 as mobile devices overtake the PC as the most popular way to get on the Web [25].

Cost of the devices steadily decreases while access to the devices and diversity of the applications steadily increases**.** As an example, when the App store started in July 2008, only 500 apps were launched. By June 2010, the number of available apps was over 225,000 and 3 trillion apps had been downloaded [16]. Another reason for this influence is that the technology appeals to the fundamental human trait of wanting to minimize effort in accomplishing a task, in much the same way as Facebook appeals to a fundamental need for social contact.

Each of the well-established areas of pervasive computing is partnered with a conceptual gap or area that needs to be better explored. While devices and applications are being increasingly manufactured and used, they are demanding more user time and attention rather than alleviating a user’s burden by means of context awareness [20]. Users need to spend more time understanding the data and educating themselves about the latest hardware and software features. Pervasive computing devices are changing our fundamental way of communicating and of gathering information. As evidence, consider the statistic that 85% of children in the

Figure 3. With pervasive computing comes increasing demand for consumer attention and, through pop-up advertisements, demand for consumer resources.

United States own a mobile phone while 73% of these children own books [16]. Technology designers and users need to keep this dynamics in mind as they scale devices, applications, and uses in everyday life. The proliferation of sensor and data modalities also increases the risk of various types of privacy invasion and security threats (adversarial attacks) the users may experience (see Figure 3).

Another breakthrough that has occurred in recent years is the ability of pervasive computing devices to perform their own energy harvesting. Small objects are fairly capable of harvesting energy. In order to push the field further, however, the ability to harvest energy needs to scale for thousands of such devices and for renewable energy sources including solar radiation, wind power, water power, vibrations, radio frequency transmissions, thermal gradients, and kinetic energy. Researchers need to understand the limits of energy production models and to design energy-aware hardware and software systems. They also need to be aware of the dangers that are posed by the proliferation of devices, including hazardous trash that results from people replacing phones and discarding old devices.

Figure 4. Social networks and crowd sourcing shape the state of the art and future for pervasive computing at scale.

Although some of the original pervasive computing goals have become reality, there are certainly emerging technologies and impacts that were not foreseen. Examples of these include the world-wide web, crowd sourcing, and social networking. In some respects the computing vision that Charles Babbage put forward does not align with the state of the art of pervasive computing. In this field we do not just consider a single user and device, but need to support communities of users and systems.

Pervasive computing has become such a large field that particular attention must be given to some of the components and influences of PeCS. In the next section, we reflect on the state of the art and future directions for these areas, and then offer some grand challenges and opportunities for the field as a whole.

4. Subfields of Pervasive Computing at Scale

The areas of research that are influenced by pervasive computing and in turn influence pervasive computing research itself are diverse. In this section we look at the future of pervasive computing at scale with respect to some of these focus areas.

4.1. Scaling models of individual and group behavior

The rapid advances in pervasive computing resulted in a proliferation of a wide variety of sensors deployed at a large scale. These in turn generate huge amounts of data that must be analyzed to extract relevant information. Data mining plays a pivotal role in the process of seeking bits and pieces of relevant information from such data explosion. The long term vision is that data mining and machine learning strategies will grow to handle spatio-temporal data at large scales, will extract necessary and relevant information, and will automatically build data models to understand human behavior. While the current progress is promising, there are a number of challenges that have to be addressed to achieve this vision.

The last couple of decades have seen rapid strides in the area of machine learning and data mining for modeling human behavior. Developments of novel algorithms and methodologies are reflected in many application areas including (but not limited to) recognition of activities, emotions, body mannerisms and gestures, and detection of abnormal behavior as well as physiological and psychological states. Most of these technologies are currently limited to data gathered from a laboratory or controlled real-world setting. While these are promising developments that have initiated inter-disciplinary research between computer scientists, behavioral and cognitive psychologists and social scientists, much work needs to be done to take the state of the art to the next level for dealing with large scale data sets and users in a real-world setting.

Large scale data sets. One of the key challenges for the future is making available well annotated, large scale data. There is truly a lack of large scale data sets available for experimentation and analysis. At present researchers collect data in silos, most often focused towards a very narrow problem. Large scale data collected through multiple modalities (such as vision, speech, wearable, phone and environmental sensors) is essential for the design, development and prototyping of architectures and algorithms to work in real-world settings. Furthermore, multi-modal data is necessary for behavior modeling as it captures the inherent multi-dimensional characteristics of human behavior. While it is sometimes impractical to collect data using different modalities, development of standardized data formats will facilitate in sharing and fusing heterogeneous data.

An important component for behavior modeling is the availability of longitudinal PeCS data. Be it physical, mental or social behavior, all tend to change over time and data collected over time facilitates analysis of dynamic behavior trends. Once PeCS research studies scale to collect data over time, the models can then be designed which capture the dynamics of evolving patterns of behavior or interaction with devices and with other humans. The study of emergent social groups and other similar phenomena is of great interest in psychology and sociology, and could be enabled by PeCS at an unprecedented scale.

Collecting large scale data sets also results in a fundamental problem of annotating them with expert-provided labels. The annotation process can be expensive and time consuming. Developing novel means of annotating data and building ontologies that can alleviate these problems will be a new direction to pursue. Developing crowd sourcing and interactive machine learning algorithms that can actively query for relevant data samples will help address these challenges.

Capturing context. The context for a PeCS application may refer to computing/communication context (network connectivity, communication costs, resource accessibility), user context (user profile, location, activity, social situation, preference), physical context (lighting, temperature, noise, traffic conditions), or time context (hour of day, day of week, season, year) [9]. By combining heterogeneous sources of information including geographical positioning system (GPS), satellite maps, recognized activities, and social media (e.g., Facebook) information, a pervasive computing system can build and use a targeted contextual picture of the current situation to be used strategically by vendors. For example, electronic kiosks in Japan already tailor advertising to potential customers as they walk by vender locations and Facebook can sell information about "change in relationship status" to wedding photographers [8].

Analysts predict that context-aware computing will very rapidly scale in popularity, exceeding $140 billion in net economic impact by 2015 [30]. However, there are several large hurdles to this targeted scale. The amount of data required to learn relevant contexts is one such hurdle: media companies are already uploading 100GB of content each day. Another difficulty is addressing the security and privacy concerns that are raised by such potentially-intrusive uses of information.

While context-aware services are crucial for PeCS, current algorithms frequently have a narrow understanding of the context recognition problem. Sensor data fusion techniques can combine disparate sources of information into a concise, usable contextual description [44]. A new direction to pursue would be to develop pervasive computing technologies that provide context information and mechanisms for integrating this information into a traditional user modeling paradigm. The natural evolution of context-aware services is then to seamlessly adapt to changing context or behaviors, at an individual, social group, or community level. This context awareness can bring the performance of PeCS to a new level by introducing non-traditional context-aware database querying techniques [27], context-aware and activity-aware network management, and a host of context-aware services.

Learning from large, noisy data sets. When pervasive computing systems scale to incorporate thousands or millions of sensors, data mining algorithms will have to deal with massive data collected from these devices. This necessitates the development of high-performance algorithms that can process information in real-time. Compressed sensing approaches [1] adopted in the signal processing and computer vision community can provide insights for developing such algorithms.

Machine learning algorithms have to often make decisions based on insufficient, incomplete and noisy data samples (i.e., in the presence of uncertainty), which is a likely scenario for PeCS. Design and development of robust algorithms, capable of making decisions in such uncertain conditions, and confidence measures that quantify the uncertainty have to be explored.

Figure 5. Modeling social interactions can identify conditions which require specialized assistance and education approaches.

Incorporating social network data. Pervasive mobile devices provide a new capability for measuring and modeling social networks. While social webs such as Facebook or Twitter provide a means for estimating the structure and strength of social connections, to a large extent these on-line sites reflect the social bonds and connections that are constructed from face-to-face social interactions in the physical world, whether at work, school, or home. Pervasive mobile devices provide the potential to gauge the strength of these interactions directly, through pervasive sensing of social interactions under naturalistic conditions. This can be viewed as a new deployment of pervasive sensing technology to directly measure the substrate of interactions from which social networks arise. We view this as leading to the development of a new paradigm of computational behavioral science.

Modeling social interactions can open up new applications for PeCS. For example, modeling interactions among young children can help screen for risk of Autism Spectrum Disorder and identify bullying. Employing mobile media can promote healthy behavior by providing context-aware suggestions and garnering social support of positive behavioral choices. This has the potential for reducing obesity among youths and helping kids with diabetes to adhere to a prescribed behavioral regimen.

4.2. Scaling PeCS devices

Pervasive computing at scale encompasses the ideas of scaling the number of devices employed by a pervasive computing system, scaling the diversity of devices, and scaling the inherent capabilities of the devices themselves. Focusing on the last characteristic, we consider that physical objects can be enhanced with computational and communication elements. This includes not only mobile devices but also familiar objects and buildings. Smart objects and smart tags can not only register their presence but also can record histories of recent interactions and route salient information through a network. Advances in smart objects, when deployed at scale, can profoundly influence our daily lives. For example, these objects can assist the elderly with tasks of daily living, provide critical monitoring of our nation's infrastructure, help with food safety, support safe public spaces, make buildings proactive, responsive, and energy smart, and design sustainable agriculture.

Our ability to develop and scale smart objects and embedded devices is rapidly improving: we can now manufacture tiny, inexpensive sensors. These small sized sensors — attached to physical objects — do not alter the object affordances. It is now possible to harvest energy for smart objects, which is an important factor to enable scaling. Industry has adopted the ZigBee specification for wireless monitoring devices and the university-industry collaboration has adapted IPv6 to run on emerging smart objects and sensor networks [48]. The IEEE 1451 standard facilitates network-independent, vendor-independent plug-and-play sensor design by specifying physical and functional interfaces between sensors/actuators and instruments/microprocessors/networks [33]. In addition, the sensor modeling language (SensorML) provides a mechanism to explicitly encode characteristics of sensors and sensor systems [35]. Despite these advances, there are still several conceptual gaps for the community to explore.

Object ecosystems. Pervasive computing research has resulted in point solutions which enhance a specific object, but cannot yet create smart object (or building) ecosystems. Some researchers have investigated middleware solutions for seamless handoff management between disparate devices [3]. However, seamless, scalable integration across devices has proven to be challenging for several reasons. First, while we have the knowledge to build smart objects, how to design scalable communication architecture for smart objects is less clear. Second, programming languages, tools, and abstractions must be specified to work across multiple smart objects in a device independent way. Because such objects are typically energy constrained, the developed programming models must view energy complexity as fundamental to the operation of the smart objects. Finally, a consistent semantics must be designed across devices to enable service composition and integration.

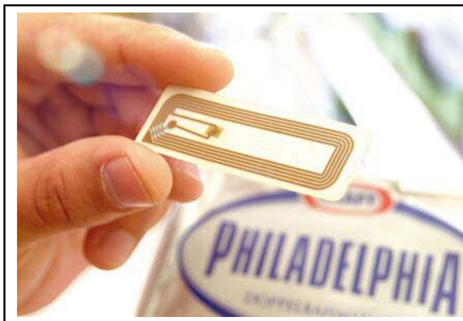

Figure 6. Smart tags on food can help ensure food safety.

Smart object applications. Applications for smart object research are diverse and immensely creative. A driving theme is that smart objects should cause change, either through actuation of other devices and services or by changing human behavior through notifications. For example, smart food can be envisioned as food items with embedded tags that record interaction history. Customers can examine the tags to determine food origin, expiration date, and safety. Similarly, embedded tags in vehicles and traffic lights can enable public safety by alerting people who are crossing the street to potential threats. Adding tags to parking infrastructure as well can enable more efficient parking. When developing smart buildings, including warehouses and hospitals, both sensing and actuation are necessary. Robots can then actuate changes to the state of a building based on events detected by embedded tags. New physical objects can also be designed to take advantage of pervasive computing. For example, a "smart swarm" of bees can conceivably help with pollination and agriculture.

If all manufactured objects are enhanced with tags, then decentralized physical object search is possible. A user can "ask" the table about a misplaced book, sparking the creation of a search engine for the physical world. Such an engine can include relational information in the tag, including cyber-physical links, which can also be queried. These types of search and query mechanism can be particularly important for individuals with cognitive and physical limitations. In addition to designing these smart objects and buildings, research will be needed to retrofit, monitor, and maintain existing physical infrastructure.

Smart phones. Phones, in many senses, represent the first pervasive mobile computing technology. Between 1990 and 2010 the number of mobile phone subscriptions grew by two orders of magnitude and in many parts of the world the current ratio is as high as 10 mobile phone users to one PC user [22]. Today's smart phone is as powerful as larger mobile devices were several years ago. They integrate powerful processors, multiple communication technologies, multimedia capability (e.g., audio, image) and ample storage and a multitude of sensor suites. The array of available wireless access and communication technologies means that phones may provide last-hop communication to body area and other deployed sensors that lack the power required for long-distance communication. Current smart phones can cache a great deal of information and the growing power of the cloud allows them to offload expensive computation. The emergence of application distribution channels has accelerated smart phone innovation by providing access to millions of deployed smart phone devices.

Figure 7. Futuristic wearable smart phone [17].

In moving ahead to what smart phones will look like in the next decade, one can imagine a device that continuously tracks our lives, including location traces, readings from internal and external sensors, and logs of our mobile-based activities. In addition to contributing to data deluge, they will also help analyze and interpret the data streams to maximize their value. By learning our activity and behavior patterns, these phones will make suggestions about our daily lives, anticipate our actions, and weave themselves into the fabric of our existence.

This vision necessitates that future smart phones be more powerful than current devices, communicate more quickly, store more data, and integrate new interaction technologies. Unfortunately, these goals are at odds with wireless data bandwidth and battery capacities, both of which are scaling rather slowly. Future smart phones are expected to deploy opportunistic algorithms that multiplex both time and space in order to improve performance. The overall heterogeneity of deployed devices and standards is another challenge that has the potential to limit device-to-device inter-operation and the potential for smart phones to interact with all of the objects, devices, and buildings they encounter.

4.3. Scaling pervasive computing applications

Pervasive computing finds application in almost every aspect of human life and activity. Here we focus on one specific domain of application, namely smart health. Health care is one of the most significant pervasive computing applications and is in great demand with the aging of the population. In the United States there will be an estimated 88.5 million individuals age 65+ in 2050 compared with 40.2 million in 2010 [45], and in the next 15 years there will be predicted shortfall of 800,000 nurses and 200,000 doctors [41]. This situation is reflected around the world and demonstrates the need for pervasive health care with the help of assistive technology to scale in a way that meets or exceeds this growing demand.

As acknowledged by the National Science Foundation [31] and the National Institutes of Health [15], information and communication technologies are poised to transform our access to health information and participation in our own healthcare and wellbeing. Pervasive computing technologies can conceivably reduce health care and rehospitalization costs as well as improve independence and quality of life in home settings. Similar performance measures can be crafted for scalable pervasive computing applications in other critical domains.

The last half decade has seen significant progress in the area of smart health care. These advances include increased adoption of electronic personal health records, availability of web-based tools to monitor personal health, commercialization of wearable/body sensors and integration with smart phones, monitoring and management of physical activity via sensors and wireless mobile devices, and increased support for aging in place. Such encouraging developments constitute only a modest start, with much more remaining to be done. The positive news is that there is wide-ranging interest and enthusiasm in all relevant scientific communities, ranging from healthcare providers and patients to policy makers. PeCS can improve the health and wellbeing of society in several meaningful ways, which we discuss in detail here.

Ubiquitous access to health information on demand. Health records and information today is scattered across various clinics and hospitals. Similarly, in-office access to providers involves significant overhead and delays that can prevent timely delivery of care. We envision a future in which health information will be readily available to designated individuals anytime, and be readily shared with providers located anywhere. Electronic health record (EHR) and personal health record (PHR) technologies support this vision, but PeCS challenges such as data integrity, security, privacy, reliability, interface standardization, data processing and visualization need to be addressed to realize this vision.

Attention to mental health. There is now a growing awareness about the prevalence of mental health issues in our society. These issues include depression, autism, post traumatic stress disorder (PTSD), chronic stress, and cognitive decline. Smart health technologies can effectively enable screening and treatment for mental health problems in a timely manner. For example, self-care methods that can be delivered privately to individuals without hospital visits may reduce the social stigma that is usually associated with mental health. Major challenges in realizing this vision include development of sensors, algorithms, models, and user interfaces for screening of mental health issues, thereby preserving the privacy of participants during treatment, and evaluating the efficacy of treatments.

Figure 8. Pervasive computing at scale can facilitate preventive interventions to promote physical and cognitive health.

Delivering timely interventions. Smart and pervasive health at scale will offer a unique opportunity to deliver intervention when and where needed, especially if it can be delivered via mobile devices. Furthermore, PeCS technologies can be used to monitor and encourage physical activity, which may help reduce the trend toward obesity in many countries. In the future, self-monitoring and real-time intervention can be developed to help people to reduce stress, address addictive behavior, depression, social anxiety, cognitive declines, autism, and PTSD, among others. Addressing each of these health issues will require the development of needed sensors, algorithms, models, middleware systems, tools, and user interfaces.

Predictive assessment and prevention. A new direction that can be pursued is to not only provide assessment of an individual's current well being, but also to perform longitudinal studies that support predictive analysis of well being and the course of disease. Knowledge of how and when symptoms become disease is crucial to appropriate intervention and prevention. Predictive analysis can also facilitate research to prevent disease morbidity and mortality.

Validation of hypotheses and technologies. We envision a future where pervasive and smart healthcare support technologies embed themselves in the infrastructure, in our environment, in the fabrics we wear, and in mobile devices we carry, thus becoming so ubiquitous that they essentially disappear from our explicit cognition. At the same time that we become less aware of them, they also become more attentive to our health needs anytime and anywhere. Extensive research is needed to make this dream a reality. Research on smart health is expensive because it needs to be validated for real-life adoption while ensuring it does not affect health adversely. To this end, the research community needs to design open extensible platforms for smart health in the same way that the sensor network community benefited from the open hardware mote platforms and the TinyOS platform. Smart health testbeds need to be designed and made available to the community along with smart health datasets. The MIT arrhythmia dataset [22], the PhysioNet physiologic signal dataset [21], and the CASAS dementia assessment dataset [43] are first steps in a direction we hope many other research groups will follow.

4.4. Scaling through cloud computing

Pervasive computing at scale via mobile phones and sensors may face fundamental limitations of storage capacity and compute power limitations. Keeping up with increasing mobile capability also introduces implications for electronic waste because mobile phones are usually replaced every 18-24 months. An alternative view to PeCS is to tap into cloud computing. Now portable or wearable devices can plug into the data collection, storage, and analysis power of the thousands of computers that are located around the world.

Cloud accessibility. A major issue in using clouds for pervasive computing is the locality of the cloud. The cloud may be placed one hop away via WiFi (a "cloud-let") or may be farther away and reachable via the Internet. Research can be pursued to define metrics for this accessibility and plan strategically to make use of both types of resources. Metrics will need to consider the demand for a data source or service, the latency of the cloud and its response to mobile devices, as well as reliability of wireless channels and availability of wireless bandwidth. Strategically, cloud-lets can be used as temporary cloud holders (e.g., the home computer, car computer, transition PC) because they can help maintain state without having to pay for the service. This raises the question of how to maintain mobile cloud-lets and manage their role with respect to mobile devices.

Mobile clouds. Latency presents an enduring, and worsening, challenge to mobile systems designers. Humans are acutely sensitive to delay and jitter and the user experience degrades further for highly interactive applications incurring even a few hundred milliseconds of latency. One possible approach to addressing this problem is to design a mobile cloud in which a proactive data delivery framework leverages a user's mobility route fingerprints with the user's contextualized behavior of data access for predictive data placement. This paradigm trades bandwidth and storage for latency and enables a number of networking application scenarios including mobile resource augmentation, on-

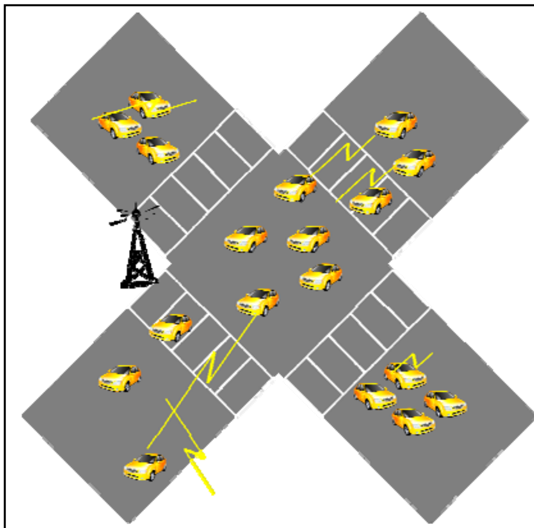

Figure 9. Inter-vehicular network.

demand social networks, smart grids, mobile health delivery, personal content distribution and intelligent transportation systems.

Network in the clouds. Future PeCS can utilize the immense cloud compute centers that communicate through ultra-high capacity wireless networks to remote mobile units in a metropolitan area. In the cloud computing scenario, physical processing is performed at the cloud data center but will virtually appear as a local service task. In addition, the wireless services of the future can themselves reside in the cloud. Computation will be displaced over speed of light wireless links and Radio over fiber links (RoF) to the cloud. Data center antennas may reside on towers far away from the cloud. Instead, RoF links displace the processing of the radio frequency signal from the cell tower to the cloud over the broadband RF fiber.

In addition to wireless networking in the clouds, other embodiments of network forms are taking shape. Vehicular and aerial networks are defining components of PeCS. Researchers can investigate inter-vehicle as well as intra-vehicle communication possibilities and use transmitted information in the context of intelligent transportation as well as general computation (see Figure 9). Similarly, research in miniaturization of sensors, micro-aerial vehicle flight, novel computation platforms and high-density power sources are enabling the design of micro-aerial vehicle swarms at an unprecedented size and scale. This research will enable a new class of applications including commercial pollination, search and rescue, surveillance, and environmental monitoring.

Clouds and crowds. In the same way that cloud computing harnesses the power of distributed data and compute services to perform local tasks, so future PeCS can make use of the massive insights of crowds to help with data-intensive tasks. Crowd sourcing, or the act of outsourcing a task to the public, has the potential to revolutionize data collection and processing by enabling in-depth, large-scale, cost-effective information gathering as well as more accurate techniques for information extraction from data. Scaling up data collection from the masses opens a spectrum of emerging research problems, such as providing incentives for efficient scale sensor data collection to millions of users, recognizing and compensating for malicious or inaccurate providers, and ensuring privacy.

We can consider a crowd as a mobile cloud where mobile devices are treated as first class entities. This strategy can be effective for addressing the latency problem since information can be transmitted between mobile devices and the cloud. Measuring crowds raises additional questions of how the crowds can be organized and defining metrics of interest. Load balancing between local phones, neighboring crowds, cloud-lets, and remote clouds lead to an interesting avenue of research and should be accomplished without letting the user know that these dynamics are occurring.

4.5. Scaling through energy analysis and harvesting

One of the biggest challenges for PeCS is the problem of powering the thousands (or more) pervasive computing devices that will be embedded in everyday surroundings and objects. Today's pervasive computing devices are primarily powered by battery. Next-generation systems will be expected to operate for several months to years without the need for battery replacement, because frequent battery replacement is not only infeasible at this scale but also implies that the devices are not "invisible". Limited battery capacity will present a significant challenge to the viability of PeCS. A promising alternative for powering PeCS systems is to scavenge energy from ambient sources such as solar radiation, wind, vibrations, radio frequency transmissions, or thermal gradients (micro-scale energy harvesting). Judicious design of pervasive systems to operate off scavenged energy has the potential to result in near-perpetual (also referred to as net-zero energy, self-sustained, or energy-neutral) system operation.

In contrast with the slow improvement of battery capacity, the rapid advancements in electronics, embedded systems, and integrated circuit design enables the important direction of practical power harvesting. While energy harvesting has been explored in the context of large systems such as solar farms and windmills, micro-scale harvesting as a systematic discipline is not as mature. The existing approaches represent point solutions rather than generalized designs because they are often tied to specific physical phenomena and applications.

However, realizing highly efficient micro-scale energy harvesting systems is challenging due to three main reasons. First, the form-factor constraint in these systems mandates the use of highly miniaturized energy transducers (often only a few cm^3, and in some cases, even mm^3). As a result, the output voltage of the transducer is very low, often far less than 1 Volt. Extracting energy from such ultra-low voltage sources is a non-trivial task. Second, the maximum power output of these micro-scale transducers is also extremely small, often in the μW range. Therefore, it is particularly important to ensure that the energy harvesting subsystem is as efficient as possible to minimize losses. Third, environmental energy supply is highly time varying in nature (e.g., changing light intensity significantly impacts the output power from solar cells) and exhibits a large dynamic range. Energy availability can be intermittent in nature. PeCS systems that are powered by these micro-scale energy harvesters should be able to adapt to such vagaries using intelligent resource management techniques. Embedded systems must be designed to address many of these new challenges. There are a few specific identified research directions and questions that will be a key to the analysis, design, and management of environmentally-powered micro-scale systems.

Foundations and basic concepts. The most common energy-related metric used to evaluate battery-powered systems is "lifetime." This metric makes sense for systems that are powered from a source that has a fixed, finite amount of energy. However, in the context of energy harvesting, where energy availability is essentially infinite along the temporal dimension (e.g., solar cells will produce electrical energy every day as long as there is sunlight), it is unclear what lifetime even means. Researchers must carefully define metrics for evaluating energy harvesting systems. One possible metric might be the ability to be energy-neutral or self-sustained, essentially evaluating whether the system scavenges enough energy per day to satisfy all of its computation and communication requirements. In addition, the metric may also incorporate the various design elements of the system including size and output power. Currently, there is no way to compare various solutions.

Efficient power extraction. A variety of interesting energy harvesting transducers has recently become available, such as thermoelectric generators, piezo-electric, and photovoltaics. Researchers need to understand the fundamental limits of these various harvesting modalities and transducers in terms of the amount of energy that they can provide per unit size. Energy transducer models must abstract away the transducer devices while designing higher layers of an energy harvesting system. A key aspect to developing these models is to decide what information needs to be captured by these models (e.g., dynamic range of available power, temporal and spatial dynamics) and inform the design of the overall system. This kind of simulation mechanism also enables the development and evaluation of hybrid solutions, which will likely be necessarily in realizing many of the pervasive computing applications. Finally, the extracted power from the transducer needs to be stored efficiently using energy storage elements such as rechargeable batteries or capacitors. New energy storage architectures should also be explored that synergistically combine heterogeneous energy storage elements (e.g., thin film batteries and ultra-capacitors) to minimize losses during energy storage.

Efficient resource management. PeCS necessitates the design of hardware and software systems that are "harvesting aware." In addition to pushing the limits on ultra-low power design, a key challenge is to design systems that explicitly consider the spatial and temporal variations in energy availability and modulate system performance/power consumption accordingly, with the goal of self-sustained operation. A key requirement to enable such harvesting aware power management is that the energy harvesting hardware must expose various control points (e.g., the amount of energy currently available from the transducer) to software. At the network level, researchers can address how to build large networks out of intermittently available devices, how to allow these devices to bootstrap and join a network, and how to synchronize interconnected devices in this kind of environment. Power asymmetry between different networked devices may be leveraged and coordinated to accomplish a task. In addition, researchers can define new methods to accurately predict future energy availability and use that information for resource planning.

Sustainability and energy management. PeCS research can also be directed toward sustainability. Smart grid research is gaining ground. In addition to this important field, pervasive computing can be used to monitor and measure power consumption, and carbon footprints, and model user behavior as they move and interact with the environment. This is particularly intriguing in environments where groups of individuals share energy resources and information can be combined with the individual's home data. From this local information algorithms can infer consumption of communities and cities as well.

Coupling information dissemination with social networks can provide comparison of consumption and incentives to minimize energy wastage.

4.6. Effect of PeCS on human factors

Researchers and commercial designers anticipate the scaling of pervasive computing to an even greater number of users, devices and applications. However, the future of PeCS will impact humans in many unforeseeable ways. There is a need to carefully consider possible implications of PeCS for privacy, security, ethics, and human interfaces as they move forward in this area.

Security and privacy. Many heavily-deployed Internet gadgets are nearly devoid of security against adversaries, and many others (including most smart phones) employ only crude methods for securing the platform from internal or external attacks. In addition, pervasive computing systems collect and aggregate large volumes of data, including location, but currently face uncertain privacy implications and offer too-limited control to users whose data is collected. The definition of privacy will continue to evolve as pervasive computing systems scale. Usually people use multiple sources of knowledge including physical, social, and experiential information, to determine the utility and the safety of technologies [22]. Designers need to keep all of these aspects in mind as they create safe, privacy-preserving scalable systems. Individuals are not the only entities that need to consider security and privacy. Organizations, including schools, corporations, banks, and governments, have a need to secure their systems and to protect proprietary interests. Researchers need to gain a better understanding of how pervasive systems reflect the needs, priorities, and structure of the organization as well as the preferences of the individuals within the organization.

A critical consideration is the fact that large-scale pervasive systems, especially those that “disappear” into daily life and lifestyle, may be the target of (or tool for) cyber warfare. An adversary may disrupt such a system as a method of disturbing, misleading, or even terrorizing a large population. Consider, for example, an attack on home heating systems in midwinter, on businesses when all transactions are conducted via mobile devices, or on public health when clothing is connected to the Internet.

While privacy-preserving data mining has been investigated, much of that work has been performed in the context of databases where each person is represented by a single record. In PeCS, data may represent time series of observations about a user or collection of users and their collective behavior including social interactions in multiple pervasive systems. Because information collected in this context may consist of complex structured sensor data and not just a collection of linear-value attributes, existing methods may no longer apply. Moreover, new paradigms need to be developed to preserve relational privacy of user's online social data at scale [28].

Scaling interfaces to field applications. Because pervasive computing applications are intimately integrated with our physical spaces, device usability is an essential concern. This becomes increasingly evident as pervasive computing systems increase in scale to include more devices, more capabilities, more humans, and more data. Commonly used interface techniques have been designed and tested in lab settings and do not always scale well to situations outside the lab. Many of these methods do not also scale over time, nor do they scale to multiple devices, multiple locations, or high volumes of data. These features characterize PeCS applications, the interface aspects of which must be evaluated in situ.

Scalability of HCI. A key issue in human-computer interaction (HCI) or even machine-to-machine interaction (MMI) for pervasive computing deals with the relationship between scalability and usability. Some aspects of the increasing scale of pervasive computing systems may in fact dampen usability challenges. For example, the increasing ubiquity of interactive applications may make them easier to learn and use. At the same time, other aspects of scalability may make pervasive computing less usable. This dichotomy demonstrates the need to identify and formalize these differences and incorporate them into the design of PeCS systems.

4.7. Scalable PeCS support

To support scalability at large in future pervasive computing scenarios, where several thousands of heterogeneous sensors and mobile nodes are expected to be involved over heterogeneous computing and communications platforms, several important systems issues and challenges need to be tackled. Such challenges include scalable middleware design to mitigate device and application heterogeneity; scalable fusion, recognition and mediation of context arising from multimodal sources, even in the presence of incomplete and uncertain information; and scalable resource discovery, composition and management in a dynamic and opportunistic manner. These scalability issues are ever more important for mobile social networking applications [5].

Context awareness. At the heart of any pervasive computing system design is the recognition and characterization of context which could be tangible or intangible. Context like user location, activity, preference, or social interaction are tangible because they can be captured or derived with the help of data collected from sensors and other devices. Whereas user intention, mood or behavior are intangible; they are as important as tangible context parameters but are not so easy to capture accurately and unambiguously. It is the notion of context and situation awareness (a sequence of contexts with semantic interpretation defines a situation) that makes a pervasive computing environment unique and perhaps different from cyber-physical systems, for example. There exists a significant body of literature dealing with context aware modeling and reasoning in pervasive computing [4,6,8,9,40]. A major challenge for PeCS is multi-context recognition and multi-context ambiguity mediation in the presence of uncertain (noisy or incomplete) information arising from a huge number of resource-constrained devices. Although some recent effort in this direction develops probabilistic models for improving multi-context recognition accuracy with the help of information theoretic reasoning [38], still a lot needs to be done for fundamental understanding of scalable interaction of a multitude of contexts, not only for single users but also for multiple users sharing the same pervasive computing environment, often with conflicting goals [14]. Naturally, context-aware algorithms and solutions involve how efficiently device and network resources are discovered and managed dynamically at scale.

Scalable middleware. Over the last decade, efforts have been made to design and develop efficient middleware systems for pervasive computing [12,26,46]. However, most of the existing middleware solutions are developed for small to medium scale systems that can handle at most 100’s of devices and/or users without performance degradation. Therefore, PeCs issues call for novel algorithmic design and modeling frameworks and middleware tools that are capable of capturing and processing large amounts of context metadata, supporting large scale collaborative sensing and computing, and adaptive decision making through spatio-temporal information fusion. These scalable middleware solutions must be lightweight, interoperable, secure, and also provide tradeoff between information accuracy, context quality, latency, and resource consumption [39]. Then only there will be wide-scale deployment of pervasive computing at scale, with the possibility of industry standard evolving.

4.8. Theoretical foundations of PeCS

PeCS is not yet a mature theoretical field and few metrics have been formalized (e.g., scaling of network capacity for wireless networks). Moreover, PeCS needs solid analytical tools that can effectively deal with scalability issues. Researchers can adapt theories from human factors, machine learning, algorithm design, information theory, control theory, game theory, sampling theory, sensor fusion, and mobile computing to PeCS, such as examining how Fitt's law [19] may be used to analyze mobile devices and interfaces.

Unifying Theory. Researchers need to develop a unifying theory for PeCS that can scale to very large numbers of heterogeneous devices, applications, and interaction methods. Because PeCS necessarily exacerbates the current scalability challenges in system architectures, there is a need to provide insights regarding the scalability as a function of the number of nodes, users, and volume of data considered. Performance measures should not only be limited to speed but should also consider reliability, accuracy, resiliency, adaptability, and usability.

Theory of useful information. PeCS systems will collect large amounts of data. They also make a great variety of data available to users and force new pieces of information, tasks, and advertisements on users. PeCS theory can include an analysis of the usefulness and usability of information and on the value of information for specific tasks and goals. Advanced machine learning and data mining techniques can be used to gather and analyze data in an unsupervised fashion in order to reduce the burden of data analysis and actionable decision making for users.

5. Additional Interdisciplinary PeCS Challenges and Opportunities

Because pervasive computing, at its heart, pervades every aspect of our lives, the field is by nature multidisciplinary. Here we highlight additional interdisciplinary opportunities for PeCS collaboration that have been minimally tapped to date. For example, partnerships between engineers and economists can be valuable as industry designs business models for PeCS. While researchers usually define for and measure performance factors such as delay, message overhead, and recognition accuracy, they also need to factor in system design, management, and usage cost.

In addition, PeCS researchers can learn from psychologists and social scientists how to conduct human subject experiments. Computer scientists need to work with psychologists and sociologists to understand and automate modeling of human dynamics and behavior, and also to understand the impact of pervasive computing on users. For example, Steve Jobs drew from his training in calligraphy to design Apple’s typography [32]. In a similar way, PeCS researchers learn from artists and from experts in sociology, law, and public policy to design pervasive computing interfaces and to understand privacy, technology acceptance, and to define policy and regulations for ethical research in pervasive computing. These collaborations can help with when designing applications that are sensitive to socio-economic and cultural differences.

Educators can also make use of PeCS research to improve the quality of education and training for future PeCS researchers. Building on the observation that students enjoy playing with new gadgets and respond well to competition, curriculum developers may make use of mobile devices in the classroom and design competitions such as designing smart phone applications to minimize power consumption.

Interdisciplinary training is also necessary for students in pervasive computing, particularly as we scale the field. Many academic institutions assume that students will obtain this training by taking classes in each of the contributing disciplines. However, this approach to education may only increase disciplinary isolation and hence may prevent true multi-disciplinary collaboration. The future of PeCS education relies upon schools offering interdisciplinary courses that actually integrate information across disciplines and focus on defining a common vocabulary.

Finally, pervasive computing has become a field that not only attracts great interest from researchers but also dramatically impacts everyday lives. With dramatic successes having been achieved in this field, researchers can now look toward the next step. We anticipate dramatic changes in the field as it begins to scale and look forward to seeing the field continue to grow.

6. Acknowledgements

We would like to acknowledge the roles of Andrew Campbell, Roy Want, and Shwetak Patel in organizing (along with the authors of this article) the NSF PeCS workshop and thank them for their help in defining related issues and areas of research. We are also grateful to the input provided by all of the PeCS workshop attendees. This material is based upon work supported by the National Science Foundation under Grant No. 1059280.

References

1. L. Atallah and G.-Z. Yang. The use of pervasive sensing for behavior profiling – a survey. Pervasive and Mobile Computing, 5(5):447-464, 2009.
2. R. Baraniuk. Compressive sensing. IEEE Signal Processing Magazine, 24(4):118-121, 2007.
3. P. Bellavista, A. Corradi, and L. Foschini. Context-aware handoff middleware for transparent service continuity in wireless networks. Pervasive and Mobile Computing, 3(4):439-466, 2007.
4. P. Bellavista, A. Corradi, and C. Giannelli. A unifying perspective on context-aware evaluation and management of heterogeneous wireless connectivity. IEEE Communications Surveys and Tutorials, 13(3):337-357, 2011.
5. P. Bellavista, R. Montanari, and S.K. Das. Mobile social networking: a survey. Under review, 2011.
6. C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, and D. Riboni. A survey of context modeling and reasoning techniques. Pervasive and Mobile Computing, 6(2):161-180, 2010.

7. A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musoles, K. Fodor, and G.-S. Ahn. The rise of people-centric sensing. IEEE Internet Computing, 12(4):12-21, 2008.
8. L. Case. IDF 2010, Research day: Context-aware computing. http://www.tomshardware.com/reviews/idf-2010-context-aware-computing-keynote,2750.html, September 24, 2010.
9. G. Chen and D. Kotz. A survey of context-aware mobile computing research. Technical Report TR2000-381, Dartmouth College, 2000.
10. D. Cook and S. K. Das. Smart Environments: Technologies, Protocols and Applications. Wiley, 2004.
11. D. Cook and S. K. Das. How smart are our environments? An updated look at the state of the art. Pervasive and Mobile Computing, 3(2):53-73, 2007.
12. A. Corradi and P. Bellavista. Mobile Middleware. CRC Press, 2006.
13. S. K. Das, K. Kant, and N. Zhang. Handbook on Securing Cyber-Physical Critical Infrastructure: Foundations and Challenges, Morgan Kauffman, 2012.
14. S. K. Das, N. Roy, and A. Roy. Context-aware resource management in multi-inhabitant smart homes: A framework based on Nash H-learning. Pervasive and Mobile Computing, 2(4):372-404, 2006.
15. Department of Health and Human Services. Translational Research to Help Older Adults Maintain their Health and Independence in the Community. http://grants1.nih.gov/grants/guide/pa-files/PA-11-123.html, October, 2011.
16. Digitalbuzz. Statistics: The growth of mobile into 2011. http:// http://www.digitalbuzzblog.com/mobile-statistics-2011-growth-of-mobile/, October, 2011.
17. Fgadgets. PCOD – Futuristic Wearable Smart phone by Karl Thompson. http://fgadgets.com/future-gadgets/pcod-futuristic-wearable-smart phone-by-karl-thompson, October, 2011.
18. D. Fincher. (Director). The social network [Film]. Columbia Pictures, 2010.
19. P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. Journal of Experimental Psychology, 47(6):381-391, 1954.

20. W. Gibbs. Considerate computing. Scientific American, 292(1):54-61, 2005.
21. A. Goldberger, L. Amaral, L. Glass, J. Hausdorff, P. Ivanov, R. Mark, J. Mietus, G. Moody, C.-K. Peng, and H. Stanley. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. Circulation, 101(23):E215-220, 2000.
22. P. Hamilton and W. Tompkins. Quantitative investigation of QRS detection rules using the MIT/BIH arrhythmia database. IEEE Transactions on Biomedical Engineering, 33(12):1157-1165, 1986.
23. G. Hayes, E. Poole, G. Iachello, S. Patel, A. Grimes, G. Abowd, and K. Truong. Physical, social, and experimental knowledge of privacy and security in a pervasive computing environment. IEEE Pervasive Computing, 6(4):56-63, 2007.
24. F. Hillebrand. GSM and UMTS: The Creation of Global Mobile Communications, Wiley, 2001.
25. International Telecommunication Union. Newsroom. http://www.itu.int/net/pressoffice/press_releases/2010/39.aspx, October, 2011.
26. M. Kumar, B. Shirazi, S. K. Das, B. Sung, D. Levine, and M. Singhal. PICO: A Middleware Framework for Pervasive Computing. IEEE Pervasive Computing, 2(3):72-79, 2003.
27. Y. Li, L. Feng, and L. Zhou. Context-aware mobile and ubiquitous computing for enhanced usability: Adaptive technologies and applications. In Context-Aware Mobile and Ubiquitous Computing for Enhanced Usability (D. Stojanovic, ed.), IGI Global, pages 147-168, 2009.
28. N. Li, N. Zhang, and S. K. Das. Preserving relation privacy in online social network data. IEEE Internet Computing, 15(3):35-42, 2011.
29. H. Lu, N. D. Lane, S. B. Eisenman, and A. T. Campbell. Bubble-sensing: binding sensing tasks to the physical world, Pervasive and Mobile Computing, 6(1):58-71, 2010.
30. E. Messmer. Garner sees ‘context-based computing’ on the horizon. Network World, http://www.networkworld.com/news/2010/102110-itxpo-content-based-computing.html, October 21, 2010.
31. National Science Foundation. CISE Cross-Cutting Programs: FY 2011. http://www.nsf.gov/pubs/2010/nsf10575/nsf10575.htm, October 2011.

32. E. Nisperos. Steve Jobs: Calligraphy lover. South San Francisco Path, http://southsanfrancisco.patch.com/articles/steve-jobs-calligraphy-lover, October, 2011.
33. NIST. IEEE P1451: Draft standard for a smart transducer interface for sensors and actuators. http://ieee1451.nist.gov/intro.htm, October, 2011.
34. NSF Workshop on Pervasive Computing at Scale. Seattle, Washington, January 27-28, 2011. http://sensorlab.cs.dartmouth.edu/ NSFPervasiveComputingAtScale/, October 2011.
35. OGC. Sensor model language (SensorML). http://www.opengeospatial.org/standards/sensorml, October 2011.
36. R. Rajkumar and I. Lee. NSF Workshop on Cyber-Physical Systems. http://varma.ece.cmu.edu/cps/, October 2011.
37. G. Roussos, M. Musolesi and G. D. Magoulas. Human behavior in ubiquitous environments: modeling of human mobility patterns. Pervasive and Mobile Computing, 6(4):399-496, 2010.
38. N. Roy, S. K. Das, and C. Julien. Resource-optimized quality-assured ambiguous context mediation framework in pervasive environments. IEEE Transactions on Mobile Computing, to appear.
39. N. Roy, A. Misra, C. Julien, S. K. Das, and J. Biswas. An Energy Efficient Quality Adaptive Multi-Modal Sensor Framework for Context Recognition. Proceedings of the IEEE International Conference on Pervasive Computing and Communications, pp. 63-75, 2011.
40. N. Roy, G. Tao and S. K. Das. Supporting Pervasive Computing Applications with Active Context Fusion and Semantic Context Delivery. Pervasive and Mobile Computing, 6(1): 21-42, 2010.
41. A. Salzhauer. Is there a patient in the house? Harvard Business Review, November 1, 2005.
42. M. Satyanarayanan. Pervasive computing: Visions and challenges. IEEE Personal Communications, 8:10-17, 2001.
43. M. Schmitter-Edgecombe, C. Parsey, and D. Cook. Cognitive correlates of functional performance in older adults: Comparison of self-report, direction observation and performance-based measures. Journal of the International Neuropsychological Society, 2011.

44. M. Siegel and H. Wu. Confidence fusion. Proceedings of the IEEE Workshop on Robotic Sensing, 2004.
45. G. Vincent and V. Velkoff. The next four decades – the older population in the United States: 2010 to 2050, US Census Bureau, May 2010.
46. M. Wang, J. Cao, J. Li, and S. K. Das. Middleware for wireless sensor networks: a survey. Journal of Computer Science Technology, 23(3): 305-326, 2008.
47. M. Weiser. The computer for the twenty-first century. Scientific American, 265(3): 94-104, 1991.
48. ZigBee Alliance. ZigBee standards overview. http://www.zigbee.org/Standards/Overview.aspx, October, 2011.