NSF Workshop on Pervasive Computing at Scale, PeCS

The NSF sponsored workshop on Pervasive Computing at Scale was held in Seattle, Washington on January 27 and 28, 2011. The workshop was attended by 75 researchers from academia, industry, and federal funding agencies. In this report we highlight visions, challenges, and recommendations for the field that were discussed in individual presentations, in breakout sessions, and in a plenary meeting.

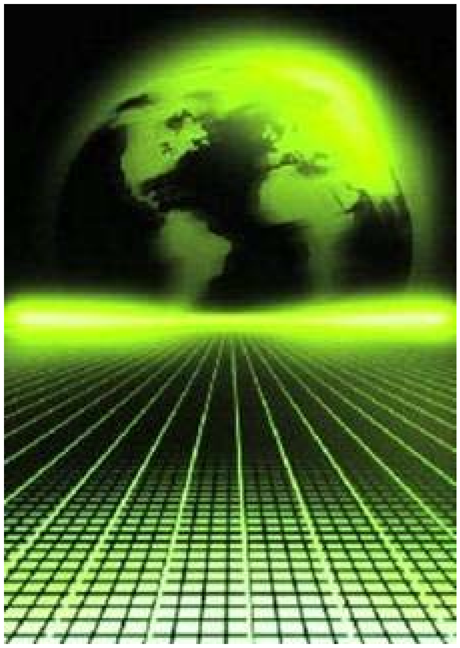

Diane Cook, Andrew Campbell, and Roy Want

February 22, 2011

NSF Workshop on Pervasive Computing at Scale

Workshop Report

Prepared by Diane Cook, Andrew Campbell, and Roy Want, PIs

Organizing Committee

Andrew Campbell, Dartmouth College
Diane Cook, Washington State University
Roy Want, Intel Labs
Shwetak Patel, University of Washington
Sajal Das, National Science Foundation
Krishna Kant, National Science Foundation
Silvia Spengler, National Science Foundation

Workshop Web Page

http://sensorlab.cs.dartmouth.edu/NSFPervasiveComputingAtScale/program.htm

Table of Contents

Introduction

The remarkable recent progress in computing power, networking, and sensors combined with the power of machine learning and data mining techniques have enabled researchers and practitioners to create pervasive computing systems that reason intelligently, act autonomously, and respond to the needs of users in a context aware manner. The field of pervasive computing and smart environments is at an interesting and critical point in its development. On the one hand, the field has matured to the point where tangible, beneficial testbeds such as smart homes, body area networks, health monitoring systems, and social networking media are becoming fairly commonplace. These visible successes are built on mature underlying technology that performs device communication, information routing, sensor fusion, activity recognition, and user preference learning. On the other hand, however, these systems have been designed and tested on small to medium-scale applications with limited dissemination of the tools, results, and datasets.

Next year is the 20th anniversary of Mark Weiser's ubicomp vision landmark Scientific American paper on ubiquitous computing: "The Computer for the 21st Century". While there has been significant progress towards 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 of devices. In order to advance the field and make technology truly pervasive, the research community needs to address the issue of scale. Future large-scale pervasive systems need to operate over different spatial and temporal scales, encompass a large number of diverse devices (e.g., mobile phones, tablets, wearables, embedded wireless sensors) that enable a spectrum of applications, deal with very large amounts of data distributed over a diverse set of networking platforms and devices, and support 100s of millions of users.

The trajectory from small to massive scale pervasive computing systems is underway. Recently, smart environments, body-area sensor networks, and smartphones with embedded sensors are enabling the delivery of a wide variety of applications from predicting traffic jams and modeling human activities, to social interactions, behavioral and mobility patterns, to community health tracking, public safety and large-scale environmental sensing. These recent developments are being driven by the availability of embedded sensors; the ease at which researchers and developers can distribute new applications to millions of users; and the emergence of the mobile computing cloud.

The goal of this NSF sponsored workshop was to discuss the challenges for scaling our future pervasive applications, algorithms, models, data and systems. The problem of scaling pervasive systems is multi-disciplinary in nature, including challenges in HCI, machine learning, data mining, mobile systems, wireless and sensor networks, smart environments, security and privacy, signal processing, control theory, information theory, game theory, optimization techniques, psychology and social networking.

The workshop program included four keynote talks by Gaetano Borriello (University of Washington), Andy Hopper (Cambridge University), David Culler (University of California at Berkeley) and Shwetak Patel (University of Washington); NSF perspective by Keith Marzullo (Division Director of CNS Division, NSF); a number of panels and breakout sessions related to scaling ubicomp; and a plenary discussion on 20 years after Mark Weiser's ubicomp vision moderated by Mahadev Satyanarayanan (Carnegie Mellon University) and Roy Want (Intel).

Workshop Format

The goal of this workshop was to provide researchers with a forum to discuss the future of Pervasive Computing at Scale. Specific objectives included:

- identify open problems and fundamental challenges that must be addressed to enable deployment of pervasive computing systems at massive scale;
- understand the needs of new applications capable of complex inferences about personal, social, and urban spaces across a set of domains, including but not limited to: smart health and well-being, social networks, smart environments, human behavioral modeling and persuasion, cyber-physical systems, and environmental and critical infrastructure monitoring and protection (e.g., smart grids);
- explore infrastructures, algorithms and tools necessary for the collection and analysis of data from large ensembles of pervasive heterogeneous and distributed devices (e.g., mobile phone sensing, wearable sensors, embedded sensors) and intelligent multi-scale decision making;
- understand the implications for privacy, security, trust and social aspects of large-scale pervasive computing systems;
- identify software challenges, including middleware and operating systems, for pervasive computing and associated applications;
- explore new interfaces and modes of interactions between people and pervasive computing devices, applications or environments;
- study the computing continuum and trade-offs where pervasive applications are self adaptive across a wide spectrum of devices and networking technologies, from the smallest embedded sensors to the computing cloud;
- provide theoretical foundations that define the "Science of Pervasive Computing"; for example, How to guarantee reliable pervasive computing at scale in the presence of uncertain and incomplete information? How to quantify and assess information quality for making accurate decisions?
- explore the nexus between scalability and application characteristics and context with the goal of identifying fundamental insights, models and methods;
- identify new networking challenges as very large pervasive systems become commonly integrated with the Internet;
- understand assurance and verification for critical applications (e.g., health-care or safety); and
- identify challenges in simulation, emulation and experimentation with pervasive systems at scale.

The organizing committee solicited two-page white papers from the community on various topics related to scaling in pervasive systems. We received 250 submissions and invited 57 participants from these submissions. In addition, we invited speakers and attendees from funding agencies and groups working in key areas.

The workshop consisted of keynote talks from four invited speakers: Gaetano Borriello, Andy Hopper, David Culler, and Shwetak Patel. Additional vision presentations were made from nine authors of the submitted white papers. The rest of the time was devoted to discussions of current and future directions in the field. Twelve breakout sessions were organized to discuss specific topic areas and a plenary discussion time was devoted to looking at the history of the field as a whole and where it could move in the next decade. Finally, participants were asked to complete an online survey to get their individual opinions and visions for the field.

Each breakout session facilitator was asked to direct discussions around the following set of questions. In addition, survey participants answered similar questions in order to get individual insights that could shape the field and provide recommendations to NSF. The questions include:

- What are the future directions of PeCS in the short term (3-5 years), medium term (5-10 years), and long term (>10 years)?
- What are the grand challenge problems in PeCS that could help the field move forward? What technical breakthroughs are necessary to solve these problems?
- Why should the federal government invest in PeCS, as compared to industry? Is there a need for a cooperative effort and what kind of technical advances could evolve out of PeCS that could lead to new industry jobs?

The rest of this report summarizes each of these discussions and summarizes overall conclusions and recommendations to NSF. A list of attendees, the workshop program, and individual survey responses are included as appendices to this report.

Overview of Keynote and Invited Talks

As part of the workshop program, the organizers invited four researchers who represent key areas of work to give keynote speeches to the group. These speakers were Gaetano Borriello (University of Washington), Andy Hopper (Cambridge University), David Culler (University of California at Berkeley), and Shwetak Patel (University of Washington).

The first keynote talk of the meeting, given by Gaetano Borriello, was titled "Looking forward to ubiquitous computing that looks ahead". In this talk Gaetano provided a survey of the pervasive computing landscape and then identified necessary technologies that are missing in this landscape. The current landscape includes a variety of portable devices that gather and store information (tabs, pads, and boards). Communication has become faster, more robust, and most pervasive as it is attached to a variety of everyday devices. The danger with the proliferation of these technologies is that we are not embedding computing into the fabric of life, we are instead making computing the fabric of life. Computing already occurs at an immense scale, to the point where individuals cannot keep informed and educated on where the data is, how to understand and analyze it, and how to use the latest apps. A challenge for the community is to let the devices do a lot of this work for us. This talk highlighted the Cell Biology laboratory at Intel Labs in Seattle as a prototype for how information can be propagated to both machines and scientists to assist with daily research tasks.

In the second keynote talk, titled "Computing for the future of the planet", Andy Hopper argued that sustainability computing is not just defense, but offensive: computing can be used to solve sustainability problems. Computing for the future of the planet must include optimal digital infrastructure that do not just design better transistors that use less energy but must focus on energy-proportional computing. Our systems can also optimize sensor type and placement to make use of renewable energy and surplus energy. Andy's group has developed a Zen package that facilitates reasoning about energy usage of computing components and searches through possible energy sources to power these components. Other key components are the ability to sense and react, such as using sensor information to provide sports analysts with insights on athlete movement in order to more effectively train, and digital alternatives to physical activities.

The third keynote talk, titled "Beyond the lamplight – lessons from making sensor networks real", was given by David Culler. This talk highlighted a number of general principles that are valuable in attempting to scale pervasive computing systems. These include learning from failure and rejection, making sure to "nail it before you scale it" and are punctuated with lessons learned from previous large-scale projects such as the Mote/TinyOS development, the NEST Open Embedded Platform, and live monitoring of the Golden Gate Bridge. The talk raised questions for the audience of what should be scaled (#nodes, extent, fidelity, accuracy, reliability, duration, realism), and why should it scale. There are many advances, tools, and testbeds available – research at scale should build on these projects rather than start from scratch.

In the final keynote speech, titled "Strategies for large scale deployment of energy monitoring and sensing in the home", Shwetak Patel described the design of devices for monitoring resource utilization in everyday environments. This technology, which allows for fine-grained monitoring of energy consumption, water consumption, and gas use as well as status of devices in the environment, is valuable for promoting energy-conscious behavior in everyday environments. Because the energy monitoring technology can be installed by the user, is small and efficient, and analyzes continuous noise generated on power lines, it is able to scale to use in every home and building. All of the monitoring devices support long-term deployments and use in hard-to-reach areas, which allow the ideas to truly scale to a variety of geographic areas and end-user applications.

In addition to the keynote talks, short (8 minute) vision statements were given by nine selected participants. Jim Rehg discussed how pervasive sensing of social behavior can be used to provide early-childhood screening for autism as well as a host of other behavioral and developmental disorders. Santosh Kumar discussed how mobile phone apps can enable stress monitoring and management for users and individuals with whom they interact. Each of these presenters talked about how to scale their own applications, and Svetha Venkatesh discussed how to scale a class of pervasive computing systems using feature extraction and analysis.

The next set of speakers introduced specific technologies that will benefit aspects of the scaling challenges. Gil Zussman talked about energy harvesting active networked tags. These tags, which enable Internet of Things research, are self-reliant because they harvest ambient light and make use of ultra-low-power communications. The WISPs that were described by David Wetherall also represent smart tags that harvest power but are also programmable, and are coupled with sensors. The tags can actually represent hosts to run apps, as routers or APs to provide power and network services, and as sensors to provide data to users. Karthik Dantu discussed how bee swarms are able to accomplish large-scale tasks at a colony level. By emulating this behavior with robotic bee swarms complex tasks can be tackled, without complex programming, precise sensors, and excessive energy usage at an individual node level.

In his talk, Oliver Brdiczka argued that contextual intelligence can help disambiguate meaning from the massive digital information that is collected and transmitted by current pervasive computing systems. The information can be converted into a personal semantic network. Thomas Little introduced the notion of ubiquitous networking through manipulation of light as a communications medium. By controlling light, sensors and mobile devices can participate in the optical field and can be maintained as part of the container in which humans live, work, and play. In the final talk, Mahadev Satyanarayanan argued that future pervasive computing research needs to conserve the most precious resource, human attention. This can be accomplished by eliminating system-induced distractions such as failures, poor performance, confusing output, and unnecessary interactions. He discussed the idea that hardware virtual machine technology can lower the external complexity of a software system by transforming it into internal complexity. By decoupling personal computing state from hardware and using VMs to encapsulate and recreate the state, migrating processes and transitioning computing from personal to pervasive can be graceful and non-disruptive.

Executive Summary of the PeCS Workshop and its Recommendations to NSF

Throughout the various workshop discussions and presentations, a number of salient issues, observations, and visions recurred. These general insights are summarized in this section along with specific challenges and recommendations to NSF for future support of Pervasive Computing at Scale.

State of the Art

First, workshop attendees were fully convinced that pervasive computing is a recognized area and is distinct from related areas such as cyber-physical systems. Aspects of pervasive computing that are unique and define the field include pervasive communication (one could argue, in fact, that the name of the workshop should be PeCCS, Pervasive Computing and Communication at Scale). In addition, fields such as cyber-physical systems explore technologies outside of a human context. In contrast, pervasive computing focuses on sensing, interacting with and aiding humans at an individual and community scale throughout their lives.

The PeCS community recognizes the wide-spread miniaturization and low-cost building of portable devices and applications for these devices. The number of such devices has certainly scaled to massive numbers. 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.

Each of these 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 rather than alleviating a user's burden. Users need to spend more time understanding the data and educating themselves about the latest hardware and software. The proliferation of sensor and data modalities also increases the risk of various types of privacy invasion and attack by adversaries, yet researchers and industry are largely unaware even of what the potential issues are.

Another area that has been heavily explored recently is ambient energy harvesting; a critical capability that would allow pervasive systems to be deployed at scale without constant maintenance to recharge, or change, their power source. Small computing devices that consume measured amounts of power can be designed to harvest enough ambient energy for some pervasive tasks. However, their ability to harvest energy needs to scale to thousands of pervasive computing devices and for alternative energy sources including solar radiation, vibrations, radio frequency transmissions, thermal gradients, and kinetic energy. Researchers need to understand the limits of energy production models and how to design energy-aware hardware and software systems that optimize the availability of a power source. They also need to be aware of the dangers that are posed by the proliferation of devices, including hazardous trash at scale that results from people replacing phones and trashing old models.

Workshop attendees also recognized that social network tools for adults are sophisticated and have scaled to massively large networks of users. However, less attention has been given to social networks for kids and their potential benefit for healthy behavior development. Additional attention could be given to integrating virtual information such as online social networks with physical information collected by sensors. Researchers can also use social network information for additional analyses such as identifying opinion leaders and trend analysis.

Vision and Grand Challenges

The workshop theme is pervasive computing at scale. In the spirit of the workshop, attendees suggest that this idea of scale be attacked more aggressively. We would like to see pervasive computing systems scale up to home area networks, metropolitan area networks / systems and smart communities that learn behavioral information and trends across a larger region.

Researchers at the workshop realize that scaling pervasive computing may result in an increased amount of data and variety of applications that demand users' time and attention. They pose challenges for the community to define metrics that quantify this demand such as ease of use and expected value of information and tasks. They also recommend that researchers learn and identify contextual information that can be integrated into more automated decision making with the goal of reducing demands on user's time. They would also like to see a unifying theory for pervasive computing that can scale to large numbers as well.

There was a workshop-wide demand for the creation of publically-available testbeds, datasets, simulation models, and open source software. These common tools should be easy to use and allow researchers to benchmark and compare the performance of pervasive computing components. They also challenge researchers to create datasets that scale to multiple age groups, demographics, and are longitudinal in nature.

A number of grand challenges were posed by PeCS workshop attendees across a number of different research domains. In what follows, we summarize the key challenges in each area.

Smart Health. The first set of challenges focuses on pervasive computing for healthcare. Current research has focused on collecting sensor information that could potentially be used for health assessment. The challenge is to take this to the next step and perform automated assessment using this information. In addition to assessing physical health, researchers can design algorithms to assess mental health including detection of dementia, depression, and PTSD. In addition to assessing well being for adults, researchers can perform early screening of children in order to detect conditions for which early intervention is critical. In addition to assessing current health status, researchers are challenged to identify health trends, and perform predictive assessment and prevention. Researchers are already aware of the need to integrate information into personal health records – the availability of pervasive computing information only highlights this need and makes it even more important to address this challenge.

Machine Learning and Data Mining. There is a need for robust tools capable of analyzing large scale spatio-temporal real-world data for physical, social and mental behavior of an individual, community or population. There is a lack of large scale annotated data sets available for experimentation and analysis. An important component for modeling is the availability of longitudinal data to study behavioral trends. The annotation process can be expensive and time consuming. Developing novel means of annotating data can alleviate these problems, and will be a new direction to pursue. With pervasive computing at scale, machine learning and data mining algorithms have to deal with data being generated from thousands of multimodal sensors. Development of distributed machine learning algorithms can collaboratively extract relevant information in real-time. Compressed sensing approaches adopted in the signal processing/vision community can provide insights for developing such algorithms. Furthermore, machine learning algorithms have to often make decisions based on insufficient and noisy data samples, which is a likely scenario for pervasive computing at scale. Design and development of robust algorithms, capable of making decisions in such uncertain conditions, and confidence measures that quantify the uncertainty need to be explored.

Smart Objects. Current research has allowed objects to become smarter, have greater memory, propagate information, and power themselves. Workshop attendees posed applications that challenge the design and use of these smart objects, such as smart eating that requires tagged food items, smart streets that utilize tag vehicles, roads, and traffic lights, smart buildings that build on tagged items for efficient search and utilty, and smart education that makes use of tagged items throughout schools to provide a more immersive educational experience.

Smart Phones 2020. To best understand the future of smartphone research we envisioned the smartphone in 2020. In a distracted world, Phone 2020 will help us deal with the data deluge by offloading much of the current human burden caused by information overload. Phone 2020 is itself continuously capturing large quantities of data about our lives—including location traces, readings from internal and external sensors, and logs of our mobile-based activities—and contributing to the steady increase in data collection. But it will also help analyze and interpret these new data streams to maximize their value. By learning our patterns, Phone 2020 will make suggestions about our daily lives, anticipate our actions, and become woven into the fabric of our existence. In order to assist us, the future smartphone will interact with everything—other phones, the cloud, nearby sensors and actuators, vehicles and buildings—and display information in ways tailored to each user.

Security, Privacy, and Ethics. The security and privacy of a pervasive system must scale. We need to understand the privacy implications of such long-term historical records, and develop usable privacy abstractions and interfaces so people are aware of the (evolving) risk and the opportunities for personal choice to manage those risks. We need meaningful behavioral data-mining with privacy support. Embedded components must be secure and thus must be adaptive to new threats. For example, the emerging area of mHealth raises new risks (e.g., attacks on medical instruments such as pacemakers). Usable security and privacy for pervasive systems remains a challenge. Usability becomes more difficult along many dimensions of scale: as the number of devices expands in a person's life, the number of interactions is quadratic and the configuration challenge explodes. In addition, sound models are needed for trust in safety-critical pervasive applications and applications that include actuators.

Social Networks and Modeling. The dynamic evolution of social networks over time necessitates an approach to PeCS in which data is gathered continuously and analyzed using models, which capture the dynamics of evolving patterns of interaction. One example in which dynamic phenomena come to the forefront is in the formation of social groups, for example when students meet for the first time at the start of the school year. The study of these emergent socialization phenomena is of great interest in psychology and sociology and could be enabled by PeCS at an unprecedented scale. Another area with significant research challenges is the use of PeCS to influence social behavior, both collectively and at the individual level. This leads to the notion of developing socially adaptive systems, which can be viewed as a logical extension of the more common idea of context-aware computing. These adaptive systems need to be informed by the evolving social context in which their users are living out their lives.

HCI. We believe that HCI is still an important aspect of pervasive computing, but it should be, ideally, an invisible aspect. That is, users should notice, as little as possible that they are interacting with a novel system. We do want the technologies to be invisible, but sometimes we want the data to become visible again; that is, we want the data in pervasive computing systems to be able to impact user behavior. With respect to understanding and defining what good HCI is in pervasive computing at scale, two competing options exist: we can either train every user to use the same interface or we can enable interfaces to tailor themselves to each user's preferences. In addition, the emergence of crowdsourcing may provide techniques to help applications and their users deal with the deluge of data. At the same time, crowdsourcing introduces new HCI issues, including introducing questions of how to deal with cultural differences among users.

Energy Analysis, Harvesting, and Storage. Micro-scale energy harvesting will greatly reduce the reliance on traditional batteries in next-generation pervasive computing systems, removing one of the biggest showstoppers to their large-scale adoption and greatly reducing the large number of batteries that are discarded every year. It is important to define a new metric for evaluating energy harvesting systems beyond "lifetime". 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. It is also important 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. This research direction involves the design of efficient hardware and software systems that are "harvesting aware." There is also a need to study design methodologies and tools that enable systematic design space exploration of micro-scale energy harvesting systems.

Intelligent Transportation. Vehicular and aerial networks are important components of future pervasive systems. Today, traffic lights, on-board navigators, and city traffic centers do not talk to each other. The challenge is to connect existing solutions via state-of-the-art communications and networking to provide efficient, coordinated real-time traffic and air quality control. The closing of the loop between traffic and air quality data sensing and vehicle routing will enable an urban traffic management that can adjust to the rapidly changing traffic and air quality conditions typical of large cities. Urban transportation is essentially multi-modal and future intelligent transportation solutions must address traffic and pollution issues assuming multiple cooperative and competitive transportation means. Future highways and vehicles will communicate with one another, making the highway system aware of the drivers' travel plans and allowing it to cooperate with, and actively instruct, the driver on achieving them. Aerial networks also present distinct challenges in computing at scale. Unlike traditional sensor networks, such swarms not only collect data but need to have the ability to make decisions using the data collected in real time. Actuation (flight) is much more expensive in terms of energy requiring a rethinking of trade-offs to be made in terms of using communication to make better actuation decisions.

Sustainability and Energy Management. The research in the area of sustainability and energy management aims to reduce the energy usage and to improve the energy efficiency of a system by monitoring and control. Sustainable systems typically rely upon heavy penetration of renewable energy sources with intermittent power generation characteristics. For energy management, one of the challenges is the use of pervasive computing and handheld devices to measure and monitor the power consumption (e.g., carbon footprint) of an individual as they move and interact with the environment. By applying statistical techniques, one can then infer the power consumption of the population in a city. If data about the power consumption of an individual is known, then adaptation of human activity as a result of guidance is likely to lead to altered human behaviors, which may not be easy to capture in models. A grand challenge in this area is to monitor and control the closed loop demand response of an entire city to within 5% of a target reference power with significant dynamic power variability due to high penetration rates of renewable sources.

Theoretical Foundations. There is no strong theoretical foundation for PeCS. Namely, there are almost no specific metrics (e.g., similar to scaling of network capacity in networking) or analytical tools that are tailored for PeCS, and that can effectively deal with the scalability issues. Moreover, existing theoretical elements in related areas have not yet been adapted for Pervasive Computing (e.g., Fitt's law, from HCI theory related to desktop PCs, type of analysis does not easily carry over to mobile devices and new interfaces). Such a theory should be able to deal with scalability to very large numbers and to a variety of devices, applications, and interaction methods (as a few examples). In particular, since PeCS necessarily exacerbates the current scalability challenges in system architectures, there is a need for insights regarding the scalability as a function of the number of nodes, users, and quantity of data. There is also a need for better understanding of interactions among components, resources, and humans. Namely, there is a need for theory that would support the understanding of emergent behavior. Since the human interactions with the devices are interleaved with interactions among the devices and among humans, there is a need for theoretical tools that will take the users and their interactions into account.

Interdisciplinary Opportunities and Challenges

Workshop attendees identified a great diversity of potential interdisciplinary collaborations. As was pointed out, we need to collaborate with economists to design business models for PeCS. While researchers usually define for and measure performance factors such as delay, message overhead, and recognition accuracy, they need to also factor in system design, management, and use cost.

In addition, engineers can benefit in many ways from collaborating with psychologists and social scientists. First, engineers can learn how to conduct human subject experiments from these researchers. Workshop attendees would like to see experiences shared for writing Institutional Review Boards (IRB) applications and sample applications. They also see a need to educate IRBs on the nature of pervasive computing research. Computer scientists need to work with Psychologists to understand and automate modeling of human behavior and to understand the impact of pervasive computing on users.

Pervasive computing researchers can also work with individuals in sociology, psychology, law, and public policy to understand privacy, technology acceptance, and to define policy and regulations for ethical research in pervasive computing. These collaborations can help with designing pervasive computing interfaces that are sensitive to cultural differences, particularly for crowdsourcing.

Educational Opportunities and Challenges

Researchers at the workshop identified several educational opportunities that arise from research in pervasive computing at scale. Building on the observation that students enjoy playing with new gadgets and respond well to competition, they recommend that curriculum developers make use of pervasive computing devices in the classroom and design competitions such as designing applications to minimize power consumption.

As was pointed out, interdisciplinary training is necessary for students in pervasive computing. Many schools assume that students will obtain this training by taking classes in different disciplines. However, we recommend that schools offer actual interdisciplinary courses that integrate information across disciplines and focus on defining a common vocabulary.

Recommendations to NSF

By reading the grand challenges listed in this section, we can see the nature of funding opportunities that are recommended by this group of researchers. Workshop attendees note that federal funding is needed for this research because industry focuses on short-term advances and profit. As a result, long-term directions, theoretical foundations, and expanding to underdeveloped countries might be overlooked. We feel that these directions are important and therefore should be supported by NSF. Some of the key recommendations raised at the workshop are outlined below.

- *Inter-disciplinary research for behavior modeling.* Strong collaborations between computer scientists and psychologists are essential to develop a better understanding of the taxonomy and properties of human behavior.
- *Create, share and maintain large scale data sets.* Since large scale data sets are critical for the design and development of algorithms, the panel recommends supporting research that aims to develop pervasive computing solutions at scale that can assist in the process of unobtrusive real-life data collection.
- *Develop pervasive infrastructure testbeds.* These testbeds will support various cross-cutting PeCS challenges, for example, consider security and privacy research: It is often necessary to construct a testbed differently if one wants to conduct security-related research, e.g., because one needs to be able to attack the devices and services within the testbed and yet not cause negative consequences outside the testbed.
- *Work to develop strong cross-directorate collaboration between CISE & SBE*. Encourage more inter-disciplinary funding programs and inter-disciplinary projects.
- *Explicitly support cross-disciplinary, integrative proposals*. NSF should recognize that there is depth in integration. This is essential to advancing the state of the art in smart objects.
- *NSF should encourage proposers to develop cross-cutting curricula*. These curricula will foster the development of graduate students in the area.
- *Support the board goals of Phone 2020*. Support experimentation at a large scale via phone testbeds, continuous sensing of people and their environments, collaborative sensing, and smart security and privacy models.
- *Develop new tools and methodology for determining ground truth at scale in the wild.* These resources can help researchers tap into the strengths of machine learning and automated reasoning algorithms.
- *Develop new models and techniques to influence social behavior.* These influences can occur both collectively and at the individual level.
- *Gather and analyze continuous data which capture the dynamics of evolving patterns of interaction in social networks.* Collaboration with graph theorists and sociologists would be valuable for this work.
- *Support of interdisciplinary pervasive computing HCI projects*. It is essential that we broaden the number of researchers involved in and aware of HCI issues at scale.
- *Recognize the type of funding that is needed to perform this work.* These scaling studies are more traditionally considered "development efforts" that require engineers for system

development. These should be recognized as part of the research process, given the fundamentally different nature of HCI research.

- *Study and develop highly efficient micro-scale energy harvesting systems.* These approaches range from devices, circuits, and architectures to power management algorithms, design exploration frameworks, and new networking protocols.
- *Support cross cutting programs in vehicular/intelligent transportation issues.* Facilitate collaboration with industry, as well as international collaboration. Coordinate with other national research programs, such as the Transportation Research Board of the National Academies, Strategic Highway Research Program (SHRP 2), the Research and Innovative Technology Administration of the U.S. Department of Transportation, and the U. S. Department of Transportation Intelligent Transportation Systems, Joint Program Office.
- *Scaling of smart health technology*. Smart health technology needs to scale so it can be adopted widely by both the general public and by scientists in scientific studies of health issues.
- *Develop open extensible platforms and testbeds for smart health.* Smart health software platform and testbed support for both wearable sensors and mobile phones that is affordable and accessible to the larger scientific community.
- *Access to health datasets.* Datasets on real people (such as the MIT arrythmia dataset and the PhysioNet dataset) are critical to the development and evaluation of various algorithms and models in smart health.
- *Access of data from utility companies*. NSF can discuss with Utility companies and Utility-University Centers of Excellence groups terms for generating anonymous data on pricing information, load demand curve, etc.
- *Organize exchange programs or workshops between academics and engineers working in utilities.* This will provide an opportunity to exchange ideas and increase collaboration.
- *Foster collaborative relationships between multiple mature theoretical areas that provide the basis for PeCS theory.* Such collaborations include, for example, HCI, networking, and machine learning.
- *Promote research focused on theoretical foundations.* Such research should also include collaborations between foundational projects and systems/experimental projects.

In addition to the grand challenges and opportunities and recommendations listed above, many attendees emphasized the NSF support the creation of testbeds and datasets, particularly if they support the validation of what cannot be currently done or evaluated. They suggest that investigators be strongly encouraged to avoid designing systems from scratch, and instead to specifically build upon current or prior tools, datasets, testbeds, and results.

They also argue that NSF should fund disruptive technologies, those that look ahead to upcoming trends, changing paradigms, and how pervasive computing scales. As usual, they feel that NSF should fund projects that are interdisciplinary. Several joint programs were recommended, particularly ones that allow NSF to work together with NIH, DARPA, transportation boards, and industry.

Workshop attendees feel strongly that NSF should create a program specifically for pervasive computing and communication technologies. The range of vision papers and discussions highlight the unique contributions of this field and the need to support advancement of the state of the art in this area.

Workshop benefits

Workshop attendees uniformly commented that they dramatically benefitted from the workshop. Many commented that the workshop discussions were more beneficial than conferences, and want to see workshops of this type integrated into most major conferences (and for NSF to sponsor more workshops like this!). In particular, attendees appreciated the diversity of attendees, research areas, and opinions, and felt that they received new ideas and collaborations that will fuel their research programs and spark new directions for their own research efforts.

Appendix A – Attendees

Michael Anderson, University of Hartford
Peter Bajcsy, University of Illinois
Farnoush Banaei-Kashani, University of Southern California
Gaetano Borriello, University of Washington
Oliver Brdiczka, PARC
Andrew T. Campbell, Dartmouth College
Roy Campbell, University of Illinois
Geoffrey Werner Challen, Massachusetts Institute of Technology
Tanzeem Choudhury, Dartmouth
David Chu, Microsoft Research Asia
Diane Cook, Washington State University
David Culler, University of California at Berkeley
Karthik Dantu, Harvard University
Sajal Das, CNS, NSF
Dejing Dou, University of Oregon
Prabal K Dutta, University of Michigan
Niklas Elmqvist, Purdue University
James Fogarty, University of Washington
Mario Di Francesco, University of Texas at Arlington
Deepak Ganesan, University of Massachusetts
Mads Haahr, Trinity College, Ireland
Zygmunt Haas, ENG directorate
Ahmed Helmy, University of Florida
Andy Hopper, Cambridge University, UK
Liviu Iftode, Rutgers University
Christine Julien, The University of Texas at Austin
Krishna Kant, CNS, NSF
Brian T Kelley, University of Texas at San Antonio
Vassilis Kostakos, University of Madeira, Portugal
David Kotz, Dartmouth College
Narayanan Chatapuram Krishnan, Washington State University
Mohan J Kumar, The University of Texas at Arlington
Santosh Kumar, University of Memphis
Jim Kurose, University of Massachusetts
James Landay, University of Washington
Wang-Chien Lee, The Pennsylvania State University at University Park
Du Li, Huawei Innovation Center, California
Thomas Little, Boston University
Yu David Liu, State University of New York at Binghamton
Jie Liu, Microsoft Research
Paul Lukowicz, University of Passau, Germany
Keith Marzullo, CNS Division Director, NSF
Mirco Musolesi, University of St. Andrews, UK
Klara Nahrstedt, University of Illinois
Vinod Namboodiri, Wichita State University
Wendy Nilsen, National Institutes of Health
Brian Noble, University of Michigan
Shwetak N. Patel, University of Washington
Qinru Qiu, Binghamton University
Vijay Raghunathan, Purdue University
Umakishore Ramachandran, Georgia Institute of Technology
James Rehg, Georgia Institute of Technology
Mahadev Satyanarayanan, Carnegie Mellon University
Andreas Savvides, Yale University
Bill Schilit, Google
Joshua Smith, University of Washington/Intel,USA
Justin Y. Shi, Temple University
Sylvia Spengler, IIS, NSF
Ioannis Stavrakakis, National and Kapodistrian University of Athens, Greece
Erik Stolterman, Indiana University
Hari Sundaram, Arizona State University
Mario Sznaier, Northeastern University
Doug Terry, Microsoft Research
Mohan M Trivedi, University of California
Gustavo de Veciana, The University of Texas
Svetha Venkatesh, Curtin University of Technology, Australia
Roy Want, Intel Labs
David D Wentzloff, University of Michigan
David Wetherall, University of Washington
Vincent Wong, University of British Columbia, Canada
Qiang Yang, Hong Kong University of Science and Technology, Hong Kong
Pei Zhang, Carnegie Mellon University
Feng Zhao, Microsoft Research Asia, China
Gang Zhou, College of William and Mary
Gil Zussman, Columbia University

Appendix B – Workshop Program

Thursday, January 27, 2011

8:30- 8:45 Welcome by workshop chairs

8:45- 9:00 NSF Perspective
Keith Marzullo (Division Director of CNS Division, NSF)

9:00- 9:30 Keynote: Looking Forward to Ubiquitous Computing that Looks Ahead
Gaetano Borriello (University of Washington)

9.30- 10.10 Presentation Session 1: Human Sensing and Smart Devices at Scale
Moderator: James Landay (University of Washington / Microsoft Research Asia)

Pervasive Assessment of Social Behavior [5 mins]
James Rehg (Georgia Institute of Technology)
Scaling Personal Stress Assistance in Natural Environment [5 mins]
Santosh Kumar (University of Memphis)
Surviving the data deluge [5 mins]
Svetha Venkatesh (Curtin University of Technology)
Energy Harvesting Active Networked Tags (EnHANTs) [5 mins]
Gil Zussman (Columbia University),
Simple Scaling for RFID-based Pervasive Computing Systems [5 mins]
David Wetherall (University of Washington),
Q&A

10:30- 11:55 Breakout session 1

B1: Cloud computing/ crowdsourcing;
Facilitators/Scribes: Deepak Ganesan (University of Massachusetts)
B2: Machine learning/behavior modeling, data mining;
Facilitators/Scribes: Narayanan Krishnan (Washington State University), Qiang Yang (Hong Kong University of Science and Technology, Hong Kong)
B3: Privacy/security/ethics;
Facilitators/Scribes: David Kotz (Dartmouth), Roy Campbell (University of Illinois at Urbana-Champaign)
B4: Smart Objects / tags / buildings;
Facilitators/Scribes: David Wetherall (University of Washington), Hari Sundaram (Arizona State University)

11:55 - 12:40 Breakout reports and discussion

12:40- 1:40 Lunch and funding agency overviews
NSF: Sajal Das, Krishna Kant, Sylvia Spengler, Zygmunt Haas; NIH: Wendy Nilsen)

1:40- 2:10 Keynote: Computing for the Future of the Planet
Andy Hopper (Cambridge University)

2:10- 2:40 Presentation Session 2: Ubicomp Problems at Scale
Moderator: Roy Campbell, (University of Illinois at Urbana-Champaign)

Contextual Intelligence: Scalability Issues in Personal Semantic Networks
Oliver Brdiczka (PARC)
Coordinating Robotic Bee Swarms
Karthik Dantu (Harvard University)
Ubiquitous Networking for Human Containers
Thomas Little (Boston University)
Achieving Ubiquity through Hardware Virtualization
Mahadev Satyanarayanan (Carnegie Mellon University)
Q&A

2:40- 4:10 Breakout session 2

B5: Smart phones;
Facilitators/Scribes: Andrew Campbell (Dartmouth), Geoffrey Challen (University at Buffalo)
B6: Social networking / modeling;
Facilitators/Scribes: James Rehg (Georgia Tech), Paul Lukowicz (University of Passau)
B7: HCI;
Facilitators/Scribes: James Landay (University of Washington, USA & Microsoft Research Asia), Christine Julien (University of Texas at Austin)
B8: Energy analysis, harvesting, storage;
Facilitators/Scribes: Vijay Raghunathan (Purdue University), Shwetak Patel (University of Washington)

4:10- 4:30 Coffee break

4:30- 5:15 Breakout reports and discussion

5:15- 6:15 Plenary Discussion: 20 years after Mark Weiser's vision on ubiquitous computing - what next?

Facilitators/Scribers: Mahadev Satyanarayanan (Carnegie Mellon University), Roy Want (Intel)

Friday, January 28, 2011

8:00- 8:30 Keynote talk: Beyond the Lamplight - Lessons from Making Sensor Networks Real

David Culler (University of California at Berkeley)

8:30- 9:00 Keynote talk: Strategies for the Large Scale Deployment of Energy Monitoring and Sensing in the Home

Shwetak Patel (University of Washington)

9:00-10:30 Breakout Session 3

B9: Intelligent transportation / vehicle networks / aerial networks;
Facilitators/Scribes: Mohan Trivedi (University of California at San Diego), Liviu Iftode (Rutgers University)

B10: Smart Health;
Facilitators/Scribes: Santosh Kumar (University of Memphis), Diane Cook (Washington State University)
B11: Sustainability and Energy Management;
Facilitators/Scribes: Vincent Wong (University of British Columbia), Brian Kelley (University of Texas at San Antonio)
B12: Theoretical Foundations;
Facilitators/Scribes: Gil Zussman (Columbia University), Justin Shi (Temple University)

10:45-11:30 Breakout reports and discussion

11:30-12:00 Closing remarks, plans for written report

Appendix C - Breakout Session Reports

Breakout Session Report: Machine Learning, Behavior Modeling, and Data Mining

Narayanan Krishnan and Qiang Yang

Participants: David Chu, Gustavo de Veciana, Dejing Dou, Diane Cook, Farnoush Banaei-Kashani, Mirco Musolesi, Oliver Bridczka, Wang-Chien Lee, Tanzeem Choudhury, Wendy Nilsen, Michael Anderson, James Landay, James Rehg, Mohan Trivedi, Svetha Venkatesh, Andrew Campbell, Peter Bajcsy, Du Li.

Introduction

The rapid advances in pervasive computing will result in proliferation of a wide variety of sensors deployed at a large scale. This in turn results in huge amounts of data that has to be carefully analyzed to extract the relevant information. Data mining and machine learning have the potential to play a pivotal role in this process of seeking the bits and pieces of relevant information from the data explosion. The long term vision is that data mining and machine learning domains will grow to handle spatio-temporal data at large scales with optimal computing resources for extracting necessary and relevant information for understanding human behavior. While the current progress is promising, there are a number of research challenges that have to be addressed to achieve this vision.

State of the Art

The last couple of decades have seen rapid strides being made in the area of machine learning and data mining for modeling human behavior. These developments in the form of novel algorithms and methodologies are reflected in many application areas such as (but not limited to) activity recognition, emotion and facial expression recognition, abnormal behavior detection, recognition of body mannerisms and gestures, detecting physiological 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 needs to be done to take the state-of-the-art to the next level for dealing with large scale data sets in a real-world setting.

Vision and Challenges

Machine learning and data mining have the potential to impact behavior modeling at scale in many positive ways. There are a number of challenges that have to be overcome for realizing this goal. Following is a brief discussion of some of these research challenges.

- *Large scale data sets:* One of the key challenges for the future is making available large scale well annotated data. There is 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 collection through multiple modalities (such as vision, speech, wearable and environmental sensors) is essential for the design, development and prototyping of 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 impractical to collect data using different modalities at a single place, development of standardized data formats will facilitate sharing the data between researchers.
- *Access to longitudinal data sets:* An important component for behavior modeling is the availability of longitudinal data. Be it physical, mental or social behavior, all of them tend to change over time and data collectedfacilitates analyzing behavioral trends. At present there is no data available for conducting these types of studies. For modeling the behavior accurately, these data sets have to be collected as part of a longitudinal study.
- *Annotated data:* Collecting large scale data results in a fundamental problem of annotating the data. The annotation process can be expensive and time consuming. Developing novel means of annotating data can alleviate these problems and will be a new direction to pursue. Another approach for solving this problem would be developing interactive machine learning algorithms that can iteratively query for data samples that are relevant for learning.
- *Context recognition:* Current machine learning and data mining algorithms have a narrow vision of understanding of the problem. These algorithms can benefit tremendously from the contextual information available when the data is captured. A new direction to pursue would be to develop pervasive computing technologies that provide context information, and mechanisms for integrating this data tinto traditional learning paradigms.
- *Collaborative distributed machine learning and data mining for real-time information extraction:* With pervasive computing at scale, machine learning and data mining algorithms have to deal with data being generated from thousands of sensors. Development of distributed machine learning algorithms that can collaboratively extract relevant information in real-time. Compressed sensing approaches adopted in the signal processing/vision community can provide insights for developing these algorithms.
- *Decision making in uncertain conditions:* Machine learning algorithms often have to make decisions based on insufficient and noisy data samples, which is a likely scenario for pervasive computing at scale. Design and development of robust algorithms, capable of making decisions in such uncertain conditions have to be explored, along with confidence measures that quantify the uncertainty.

In summary, the machine learning and data mining community envisions development of robust tools capable of analyzing large scale spatio-temporal real-world data and make inferences on the physical, social and mental behavior of an individual or a community.

Recommendations to NSF

The following are the recommendations of the panel for addressing the challenges discussed above.

- *Support for inter-disciplinary research for behavior modeling:*
    - The panel recommends funding for inter-disciplinary research for behavior modeling. Strong collaborations between computer scientists and psychologists are essential to develop a better understanding of the taxonomy and properties of human behavior. Furthermore, these collaborations have the potential to crystallize the goals and objectives that drives innovation of new technologies.
- *Supporting research aimed at creating, sharing and managing large scale data sets:*
    - Since large scale data sets are critical for the design and development of algorithms, the panel recommends supporting research that aims to develop pervasive computing solutions at scale that can assist in the process of unobtrusive real-life data collection. In particular, the panel emphasizes the need for research that aims at collecting data over a long time. Design and development of novel data annotation mechanisms is also of interest to the community.
    - Another recommendation of the panel is for NSF to partner with industries providing real-world data-sets to the research communities. These data sets can be made available as part of competitions at conferences to benchmark state of the art algorithms, and facilitate incremental development of these algorithms; akin to the KDD-CUP (an annual data mining and knowledge discovery competition organized by the ACM special interest group on knowledge discovery and data mining)
- *Educational opportunities:*
    - The panel recommends support for inter-disciplinary courses that encourages development and learning of a shared vocabulary amongst the disciplines. This facilitates the translation of theories and ideas between the different domains. For example, theories and formulations from computational psychology can be used to develop novel machine learning/data mining algorithms; while computer science can provide valuable empirical and analytical evidences to support some of the psychology/behavioral theories.

Breakout Session Report:
Security, Privacy, Ethics

David Kotz and Roy Campbell

Participants: Roy Campbell, Geoffrey Challen, Sajal Das, Mario Di Francesco, Mads Haahr, Ahmed Helmy, Liviu Iftode, Vassilis Kostakos, David Kotz, Santosh Kumar, Thomas Little, Justin Shi, Ioannis Stavrakakis, Vincent Wong, Pei Zhang.

Introduction

In the following summary we provide the key nuggets extracted from our discussion, in bullet form, aligned with the seven questions asked in the Guidelines for breakout sessions.

State of the Art and Existing Conceptual Gaps

State of the art: What aspects do we understand well enough (i.e., mostly scope for refinements instead of breakthrough research).

Pervasive technology security and privacy is either non-existent or crude.

a. Many common (pervasively deployed) Internet gadgets have nearly no security against adversaries and many others (including smartphones) have only crude methods for securing the platform from a physically present or remote adversary.

b. Today we are good at collecting and aggregating lots of data, but with uncertain privacy implications and too-limited control given to the user whose data is being collected. In addition, we do not know how to scale user security and privacy across multiple devices, applications, and services that may source information to large aggregated data sets?

c. There are a few libraries of sensed data available to pervasive-computing researchers with which to study privacy or security issues, but these libraries are limited in their scale and scope.

d. Secure & privacy policy technologies exist and there are research papers that define privacy frameworks, policy languages, and privacy interfaces. However, these mechanisms have not caught on in main stream deployed systems.

e. The state of the art in *usable privacy interfaces* is extremely poor; as one put it, we live in a 'lawless land' where anything goes, and users are on their own when it comes to discovering privacy policies and specifying privacy choices.

f. The privacy challenges in other aspects of pervasive computing are poorly understood. For example: the emerging area of mHealth; the evolving issues of location privacy where several solutions offer privacy-preserving location-based services and location anonymization, but little is understood about users' concerns regarding location privacy or the broad range of meanings for "location",or its uses in pervasive systems.

g. Applications like targeted advertisements. These can be based on patterns of users' behavior in the virtual world (cyberspace), but the use of broader sensing modalities is just beginning to be incorporated into inference-based systems and the implications for security and privacy are unknown.

Vision and Challenges

Existing conceptual gaps: longer term picture, vision, new areas; in no particular order:

The security and privacy of a pervasive system must scale.

a) Huge amounts of data may be collected and aggregated over long time scales (years). We need to understand the privacy implications of such long-term historical records, and develop usable privacy abstractions and interfaces so people are aware of the (evolving) risk and the opportunities for personal choice to manage those risks. On a related topic, sampling of sensed data can reveal identities even if the data has been anonymized. If data is kept for months or years, does it become more or less vulnerable to such threats?

b) We need meaningful behavioral data-mining with privacy support. Although the data-mining community has done a lot of work on privacy-preserving data-mining, much of that work is in the context of databases where each person is represented by a single 'record'; in the pervasive-systems context, the data may represent a time series of observations about the users' behavior in multiple pervasive systems, the information may be complex structured sensor data (rather than attributes), etc.; so the existing methods may not apply.

c) Individuals are not the only entities that need to consider their security, or their privacy. Organizations – schools, corporations, and governments – have a need to secure their systems, and to protect proprietary interests. How do pervasive systems reflect the needs of the organization as well as the preferences of the individuals within an organization (which may sometimes conflict)?

d) Embedded components must be secure and thus must be adaptive to new threats. For example, the emerging area of mHealth raises new risks – which still need definition – and include security threats that can, quite literally, kill you.

e) *Usable* security and privacy for pervasive systems remains a challenge. Usability becomes more difficult along many dimensions of scale: as the number of devices expands in a person's life, the number of interactions is quadratic and the configuration challenge explodes. This *configuration challenge* is a cognitive burden that we know (from the pc world) will lead to security holes and privacy leaks.

f) We need natural interfaces for security and privacy that suits the task at hand – even more important as a pervasive system fades into the background is the need to offer access to its security and privacy aspects. Intuitive abstractions drawn from the non-technology world may help users to express their privacy preferences and be aware of their exposure to outsiders.

g) Mechanisms to allow users to control information across a range of integrated applications and services. Today, each application or service has its own configuration interface, and where they exist, privacy settings.

h) A conceptual framework that will help researchers and developers balance utility, privacy, and social benefits and make those choices well.

i) An answer as to whether privacy management can be automated, and to what extent?

j) A better *definition* of privacy that allows it to be refined from application to application or context to context, or culture to culture? What information needs protection, and how much?

k) An answer to whether *trust* should be embedded into low-level distributed computation and communication technology, or whether it belongs in higher layers. Conversely, is it necessary to have low-level support for trust (e.g., trusted computing platforms) to be able to build high-level trustworthy systems? Sound models are needed for trust in safety-critical pervasive applications, and applications that include actuators.

l) Solutions to enable researchers to conduct large-scale experimentation with real users, for a wide variety of pervasive systems. Not all pervasive systems are apps that run on smart phones.

m) Cyberwarfare: it is conceivable that a large-scale pervasive system, especially one that has ‘disappeared’ because it is so ingrained into daily life as to become invisible, may be the target of (or vector for) cyberwar. An adversary may disrupt such a system as a method of disturbing, misleading, or even terrorizing a large population. Consider an attack on home-heating systems in midwinter, on commerce when all transactions (including point of sale) are conducted via mobile phone, on public health when everyone’s clothing is connected to the Internet.

Interdisciplinary Collaboration

What interdisciplinary collaboration would be critical to address item (2) above? What are major challenges/proposed solutions for interdisciplinary research?

We need stronger relationships with other disciplines including social science (sociology, psychology, law, public policy). These connections may be particularly difficult to form in cases where the key players are not motivated by research funding. There is a language barrier across disciplines; each has its own jargon and methods.

We need better ways to engage industry – handset makers, sensor makers, software makers, platform providers, and cellular providers. Today it is difficult for more than a handful of researchers to build meaningful relationships.

Challenges

Challenges and approaches for supporting large scale experimental research in the areas addressed by your breakout? For example, how do you get academics access to real smart grid facilities (not just simulators)? What are the cost, safety, training, development, support, and other issues? Can you support a “planetlab” of such facilities? Ditto for smart buildings, transportation, etc.

We should build shared testbeds of various kinds of pervasive infrastructure. A shared testbed should have a board (drawn from the research community) that can review proposed research projects, allocate testbed resources, and consider ethical issues (particularly if human or animal subjects are involved).

a) We should develop large, representative user cohorts who are willing to be part of a series of ongoing studies of pervasive-computing technologies. See social-science examples.

b) We need help defining policy (or regulations) for ethical research in pervasive computing, regarding user privacy, and regarding security research in pervasive infrastructures. We need community norms that will guide researchers to ethical approaches to research.

c) We need to educate IRBs about the nature of pervasive computing research; most IRBs are accustomed to medical research and certain types of social-science research but have little experience with information technology.

d) Our community should share IRB experience across research groups to lower the barrier to setting up testbeds and experiments. Examples of successful IRB applications may allow others to ramp up research more quickly.

Mechanisms to Improve Experimental Research

What mechanisms do we need to explore that will improve data availability for experimental research? There could be many aspects here including IP issues, sharing, provenance, infrastructure for making it available, data quality, who manages the data, form of access, etc.?

a) We need to fund research on ways to publish datasets that remain meaningful and usable, but protect privacy of users in data. That is, how do we anonymize traces (collected in pervasive sensor systems) while maintaining utility for research? Large-scale data collections will never be possible without an answer to this question.

b) For pervasive-infrastructure research, the community could work with national labs who have data on such infrastructure (at the national scale), to make data available for research.

c) We should learn from long-term efforts to collect social-science datasets. Those communities have been doing this for decades.

d) We need to develop community norms (see above) – what does our research community believe is ethical to collect, and to share, when human subjects are involved?

e) We need to make data available in ways that social-science researchers can use it. They too can benefit and bring an important (different) perspective to the data.

Educational Opportunities and Challenges

Educational opportunities and challenges, including multidisciplinary education and training of faculty & Postdocs. (This may be mostly focused on education to enable interdisciplinary research).

a) We should support cross-training of postdocs & graduate students, so that technology students learn sociological and psychological research methods, ethics, philosophy of privacy, while non-technology students learn about technology, security, privacy. We need a cohort of researchers who are comfortable working across the CS / Social Science boundary.

b) We as technologists are obligated to educate the public about security & privacy risks – and best practices – in pervasive computing. What are the risks, and the best practices, anyway?

c) We need more research on the ethics of pervasive computing – especially in pervasive systems that include humans. And, are the ethics of society evolving as technology drives social change? Has our cultural sense of privacy evolved? Are we more comfortable with invasion of technology into our personal life?

d) We need more research on the ethics of international cyber conflict, especially when it involves direct attacks on individual citizens.

Recommendations to NSF

Fund several pervasive infrastructure testbeds, of various flavors, designed with security & privacy research in mind. It is often necessary to construct a testbed differently if one wants to conduct security-related research, e.g., because one needs to be able to attack the devices and services within the testbed, and yet not cause negative consequences outside the testbed.

a) Work to develop strong cross-directorate collaboration between CISE & SBE, to encourage more inter-disciplinary funding programs and inter-disciplinary projects.

b) Consider asking the NAS/NAE to bring together researchers & industry experts to clarify ethics of research and products in pervasive computing. The questions are still muddy, and the answers even muddier, and the product developers can do pretty much whatever they want. Where are the boundaries? What are our norms?

c) Seek ways to explore the potential for cyberwarfare via pervasive systems, and the risks therein, in open (non-classified) research programs.

Breakout Session Report:
Smart Objects / Tags / Buildings

David Wetherall and Hari Sundaram

Participants: Mads Haahr, Mario Sznaier, Gang Zhou, Du Li, Gil Zussman, Qinru Qiu, David Liu, Niklas Elmqvist, Brian Kelley, Joshua Smith, Shwetak Patel, Andreas Savvides, Joshua Smith, Vincent Wong, Karthik Dantu, David Wetzlof

Introduction

The breakout group focused on physical objects, including everyday objects, and buildings, enhanced with computational elements. These computational elements include the ability compute, store and communicate. We made the distinction between smart objects and familiar electronic devices deemed smart i.e. smartphones.

Advances in smart objects, when deployed at scale, can profoundly influence our daily lives. These objects can assist the elderly with tasks of daily living, provide critical monitoring of our nation's infrastructure, help with food safety, efficient and safe public spaces, design of proactive and responsive buildings, and advance agriculture.

State of the Art and Existing Conceptual Gaps

Our ability to develop and scale smart objects is rapidly improving: we can now manufacture small and inexpensive sensors. Sensors attached to physical objects can be made small enough they do not alter the object's affordances. It is now possible to embed ambient energy harvesting technologies in smart objects; while there are important constraints, including the smallest physical dimensions at which harvesting yields benefits, this is now a practical solution. The harvesting of energy at the smart object is an important factor to enable scaling. Participants also noted that industry has adopted the ZigBee specification for wireless monitoring devices, and a university-industry collaboration has adapted IPv6 to run on emerging smart objects and sensor networks, although much work still remains on issues of discovery, routing, and transport.

The group focused on several conceptual gaps. We can now develop point solutions — enhance a specific object — but cannot yet create smart object ecosystems. Seamless, scalable integration across devices has proven to be challenging for several reasons. First, while we have the knowledge to build smart objects, a scalable communication architecture for smart objects is less clear. Second, we lack programming languages, tools, and abstractions to work across smart objects in a device independent way. Furthermore, since such such objects are typically energy constrained, we need programming models that view energy capture and usage as afundamental char parameters of their design. . Third, there is a lack of consistent semantics across devices to enable service composition and integration. Finally, we need clarity on applications for smart objects, including how to address privacy concerns; national priorities including health, energy, education, and security may be important drivers for smart objects. The applications will help close the loop with smart object ecosystem design, by providing clarity on the essential smart object constraints.

Grand Challenge Applications

The group brainstormed applications, to support smart object research. A driving theme of our discussion was the idea that smart objects should cause change — either through actuation of other devices or services, or by changing human behavior through notifications. We now list some key applications:

1. *Smart Food*: If each food item, including processed food, vegetables, meat / fish / poultry, had an embedded tag, the tags could record interaction history. Then, customers could examine such tags to determine food origin, and if the food had been handled safely.
2. *Safe Streets*: Embedded tags in vehicles and traffic lights would enable public safety — we can alert people crossing the street to potential threats. Additionally, with embedded tags in cars and in parking infrastructure, we can enable efficient parking.
3. *Responsive Buildings*: In developing smart buildings, including warehouses and hospitals, both sensing and actuation are necessary. Robots can actuate changes to the state of the building based on events detected by embedded tags.
4. *Mobile objects*: We can develop a "smart swarm" of bees, which can help with agriculture, including pollination.
5. *Decentralized physical object search*: If all manufactured objects are enhanced with tags, then decentralized physical object search is possible. You can "ask" the table, about the misplaced book — we can develop a "Google" for the physical world. We can embed relational information, including cyber-physical links, in the tag, which can be queried at a later stage. Search and query mechanisms can be particularly important in the case of people with disabilities and the elderly.
6. *Infrastructure maintenance*: In addition to the design of new, smart buildings, maintaining and monitoring *existing infrastructure* is of critical importance. These include working alongside infrastructure in existing homes, and retrofitting building, bridges, power stations and other physical infrastructure of national importance.

Mechanisms to Improve Experimental Research

Lack of infrastructure reuse is a significant impediment to advancing the state of the art in smart object research. Today there is a lack of shared infrastructure, software tools and data traces. The lack of shared infrastructure is crucial — today, researchers have to develop both the hardware, and a smart object test-bed to support their research. While infrastructure development is valuable in many circumstances, rediscovering and spending time dealing with issues addressed by contemporaneous research, is wasteful. We do not need each and every team to develop a smart-sensor, -power meter, and -building. A process overseen by the NSF that enables researchers to reuse and share infrastructure, including hardware devices, software tools and libraries, and data traces, developed by peer NFS-funded research would be very valuable and accelerate the Nation's return on research investments in this area.

Educational Opportunities and Challenges

We need an interdisciplinary approach to educate PhD students about developing smart object ecosystems. As many participants observed, this is a cross-disciplinary area, including computer systems and software, communications (both lower and the upper network layers), and circuit design. At present, a typical faculty member encourages their PhD students to take courses from different disciplines with the expectation that the students will integrate knowledge across classes.

We need to design inter-disciplinary courses, both at the graduate and undergraduate levels, which are broad in terms of scope, and which focus on integration across the different disciplines. Holistic courses, for example, could require each student to build smart objects, including hardware platforms, software libraries, and web service, thus clarifying the need for integrative understanding.

The challenge is to make the course have depth; one possible solution is to have several inter-disciplinary courses in sequence. Another possible solution is to have multiple "self-contained" courses that increase in depth. The group also focused on energy use as a key concept to transcend pedagogical approaches to smart objects. Class competitions can motivate efficient energy use, with minimization of power as the main goal. Another challenge brought out in the discussion is that the new class needs to fit within the existing programs.

Recommendations to NSF

There are three concrete recommendations to the NSF:

1. The NSF should explicitly support cross-disciplinary, integrative proposals. In particular, it needs to recognize that *there is depth in integration*. This is essential to advancing the state of the art in smart objects. Prior funded research, while advancing the state of the art in a particular layer (e.g. communication), does not in of itself, yield novel outcomes in smart objects. Such layer specific research makes "black box" assumptions about other layers; the assumptions may not hold in practice, while integrating across layers. The challenge in exploring pervasive sensing at scale is a strong focus on the systems perspective – of really getting the whole thing to work – rather than the component perspective which more often is common in academic circles. One keynote said it best: "Nail it before you scale it."
2. The NSF should encourage proposers to develop cross-cutting curricula, to foster the development of graduate students in the area.

3. Infrastructure re-use is critical in advancing the state of the art: there is little need for each project to develop its own infrastructure. The NSF should develop mechanisms to ensure that each new project builds on the outcomes, including devices, testbeds, datasets, from concurrent or prior funded research. This includes support for novel mechanisms to fabricate and share hardware artifacts – something akin to a MOSIS for embedded systems hardware.

Breakout Session Report:
Smart Phones

Geoffrey Challen and Andrew Campbell

Participants: Gaetano Borriello, Andrew Campbell, Roy Campbell, Geoff Challen, David Chu, Sajal Das, Mario Di Francesco, Mads Haahr, Ahmed Helmy, David Kotz, Narayanan Krishnan, Mohan Kumar, Thomas Little, Jie Liu, Mirco Musolesi, Kishore Ramachandran, Mahadev Satyanarayanan, Andreas Savvides, Bill Schilit, David Wetherall, Vincent Wong, Feng Zhao, Gang Zhou.

Introduction

Phones are the first pervasive mobile computing technology. Between 1990 and 2010 the number of mobile phone subscriptions grew by two orders of magnitude. Today's phones are migrating from so-called feature phones—limited to voice and text messaging—to smartphones which integrate powerful processors, multiple communication technologies, ample storage and sensor suites. The ubiquity and increasing capabilities of smartphone devices make them our best option for realizing the pervasive computing vision at scale.

State of the Art and Existing Conceptual Gaps

Today's smartphone is as powerful as larger mobile devices were several years ago. It integrates multiple processors, including some specialized for specific tasks. It can communicate data over 1,000s of meters to cellular towers using 3G or 4G, over 10s of meters to 802.11 access points using Wi-Fi, and over 1s of meters to many other devices using Bluetooth. This array of communication technologies mean that phones may provide last-hop communication to body area, and other deployed sensors, that lack the power required for long-distance communication. Cheap and plentiful storage allows smartphones to 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 like the Apple AppStore and Google Android Market have accelerated smartphone innovation by providing access to millions of deployed iPhone and Android devices.

Grand Challenge Applications: Phone 2020

To frame our discussion of the future of smartphone research, our session outlined a vision of the smartphone in 2020. We imagine the capabilities of Phone 2020 and some exciting future applications below. Working backward, we develop a set of research challenges that must be addressed before Phone 2020 can become reality.

In a distracted world, Phone 2020 will help us deal with the data deluge by offloading much of the current human burden caused by information overload. Phone 2020 is itself continuously capturing large quantities of data about our lives—including location traces, readings from internal and external sensors, and logs of our mobile-based activities—and contributing to the steady increase in data collection. But it will also help analyze and interpret these new data streams to maximize their value. By learning our patterns, Phone 2020 will make suggestions about our daily lives, anticipate our actions, and become woven into the fabric of our existence.

In order to assist us, the future smartphone will interact with everything—other phones, the cloud, nearby sensors and actuators, vehicles and buildings—and display information in ways tailored to each user. It will process the environment and help us discover and navigate the world around us, including visibility into social networks. By better understanding users, the Phone 2020 will manage their attention and know when to interrupt. Through an increase in its own capabilities and by seamlessly inter-operating with powerful cloud resources, Phone 2020 will be starting to make desktop and laptop computers obsolete.

The new capabilities of Phone 2020 will support new applications that open up new markets. Smartphones will define the classroom of the future. They will augment reality to further education, socialization, health care and gaming. They will sense reality to manage cities, workplaces and traffic while continuously recording our digital lives. Smartphones of the future will help us work more efficiently, serving as portable office and personal digital assistant, conserving useful working hours and creating time for leisure and entertainment. We expect future applications to be long-lived—leveraging continued interaction with users over time—and local—exploiting the density of smartphone penetration to augment or replace communications infrastructure, critical in developing countries where such infrastructure may be unreliable or nonexistent. The ubiquity of smartphones and their proximity to their human users will make them a critical component of future approaches to disaster relief and emergency management.

Phones will also continue to be integrated with online social networks. Smartphones are already the quintessential social device. The desire of people to connect with each other drove the adoption of cellular phone technologies. With social networking exploding on the Internet in 2011, Phone 2020 unites the social network with the social device. It will help us further understand the structure of existing social structures, while assisting in the formation of ad-hoc social networks grounded in physical gatherings of people with similar interests. Phone 2020 will also contribute to network science by monitoring user behavior, and supporting applications such as disease tracking.

Mechanisms to Improve Experimental Research

In order to build Phone 2020, we identified a number of challenges that our community must address. These divide into three categories: (1) developing the capabilities of the smartphone and its environment, (2) improving interaction between smartphones and users, and (3) coping with the potential for massive large-scale data collection using smartphone-integrated sensors.

The Phone 2020 vision is predicated on continued improvements to smartphone and smartphone infrastructure performance. Future smartphones must be more powerful, communicate more quickly, store more data, and integrate new interaction technologies. Unfortunately, these goals are at odds with data bandwidth and battery capacities, both of which are scaling slowly. We expect future smartphones 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 limiting device-to-device inter-operation and the potential for Phone 2020 to interact with all the devices it encounters. We also discussed the importance of integrating the smartphone with existing Wi-Fi networks to improve connectivity and network performance. Peer-to-peer architectures were suggested as a potential way to improve performance, particularly when infrastructure is lacking.

Another property that is not scaling over time is human attention. We already pay too much attention to our smartphones to believe that we have achieved the invisibility captured by early visions of ubiquitous computing, and this problem is worsening. To better optimize our attention future smartphones must deploy interfaces allowing more nuanced interaction with users and capable of processing emotional cues. To improve the interaction between humans and their devices, new algorithms must be developed enabling behavior-based modeling, computing, and testing. In addition, user interfaces need to be reconsidered, including those that, while unsuitable for larger devices, may work well on smartphones. Phone 2020, with its ability to interact seamlessly with objects around it, will be able to leverage “found” interface elements in the environment to enable much richer interaction modalities than those possible on the smartphone itself.

Smartphones hold the potential both to contribute to and to alleviate the growing data deluge. Large-scale deployment of sensor suites on smartphones combined with cheap bandwidth and storage will lead to a growing amount of data produced by the smartphones of the future. Securing this information—much of it sensitive and personal—will be a major challenge. Designed as a personal device, smartphones are increasingly interacting with each other and the environment, creating new opportunities to steal and misuse information. Developing security and privacy models that users can understand and adapt to their needs is a critical challenge to the continued advance of this technology.

Interpreting and processing the collected data will also be difficult. There are opportunities for harnessing the distributed power of large numbers of smartphones through collaborative computation. These capabilities, if developed, might complement the continued aggregation of computation in the cloud. Fundamentally, however, the smartphone of the future will be a portal to the intelligent processing and management of data in order to reduce user distraction and allow users to focus their attention elsewhere.

Educational Opportunities and Challenges

One germane direction for the academic community to explore is in the use of smartphones to enhance education and learning. The future smartphone may enter the classroom and help put lessons in context, as well as extending the reach of learning beyond the classroom.

We also believe that continued growth and competitiveness in the smartphone market depends on educating the next generation of computer scientists on smartphone development. Given the centrality of the smartphone and the cloud to future computing, we must train engineers that can help integrate these two technologies in ways that harness the properties and capabilities of both. We recommend support for the continued development of courses in smartphone programming, application development, and smartphone-cloud interaction.

Recommendations to NSF

Our recommendations highlight areas where the research community can make significant and distinct contributions. Industry is already very active in this space and has many advantages, particularly when working at scale. However, there remain many opportunities for the academic community to develop the future smartphone in directions complementary to those being pursued by industry.

We recommend that the NSF develop research programs addressing the key challenges to realizing the Phone 2020 vision outlined above:

1. We need to continue the development of smartphone and infrastructure capabilities to support demanding new applications.
2. We must tear down the walls that divide devices from each other and limit the ability of the smartphone to fully understand its environment.
3. We need better interfaces allowing the future smartphone to conserve human attention.
4. We need smartphones to help users cope with the ever increasing amount of data accessible to and collected about them.
5. We need security and privacy models that users can understand and adapt to match their expectations and the current context—the highly dynamic pool of surrounding devices and communications channels, the social setting, and the user's activity.
6. We believe it is important to understand and document our continued co-evolution with our mobile devices: how we are changing them, how they are changing us.
7. We believe to develop new tools and methodology for determining ground truth when pervasive human and context sensing applications are deployed in the large.

To enable academics to succeed at complementing industry, the NSF should provide them with resources and infrastructure facilitating experimentation at scale. NSF can also take a role in partnering with industry to gain access to large numbers of smartphones, air time, call logs or other large data sets. Further partnerships with industry might also allow us to do citizen-driven science in other areas that leverage the smartphone as a pervasive computing platform.

Application distribution channels like the Android Market and Google AppStore also provide academics with the opportunity to deploy research systems at scale by leveraging channels established by industry. We can release our own code on the AppStore, perhaps piggybacking on top of other popular applications. Users worldwide might be willing to participate in a large-scale virtual laboratory. At sufficient scale such a laboratory could provide built-in guarantees to researchers allowing academic research to reach large numbers of deployed smartphones.

Breakout Session Report:

Social Networks and Modeling

James Rehg and Paul Lukowicz

Introduction

Pervasive mobile devices provide a new capability for measuring and modeling social networks. While social webs such as Facebook provide a means for estimating the structure and strength of social connections, to a large extent these on-line sites serve to 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 measure and gauge the strength of these interactions directly, through pervasive sensing of social interactions under naturalistic conditions. This can be viewed as the 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

While a significant amount of work has been done on analyzing an adult social network, relatively little attention has been focused on the social networks constructed by children, particularly children at a young age. In this context the ability of pervasive computing platforms to sense social interactions is vital, since these children do not participate in traditional on-line social media. The ability to model and analyze children's social networks would be profoundly useful in a variety of psychology and education contexts. For example, it is well-known that children who may be at risk for Autism Spectrum Disorder will respond differently from neurotypical children to a social bid, i.e. a request for social engagement. There is great interest in psychology in understanding the patterns of social interactions among children in natural settings. For example, there is interest in characterizing the behavioral phenotype for autism, i.e. the manner in which the syndrome is expressed as multiple categories of responses to social bids and patterns of interaction, or the avoidance of interaction. The ability to measure and analyze social interactions could be an important element in large scale approaches to screening for risk of ASD, for example based on naturally-occurring behaviors measured in a daycare environment. In addition to its potential utility for diagnosis and treatment of behavioral and developmental disorders, pervasive sensing of social interactions under naturalistic conditions can be valuable for education, making it possible to understand patterns of interaction that act positively or negatively to impact learning, including an increased capability to measure socialization and identify bullying.

Another exciting possibility at the intersection of social networking and pervasive mobile devices is the ability to influence behavior at community scales. The use of social media to organize large groups of people has been widely-observed, but with sufficiently powerful mobile devices it should be possible to go beyond simple communication functions and provide context-aware services that directly influence behavior. In the field of preventive medicine and health, for example, it has been shown that if you can provide people with information that directly relates to their behavior at the moment when they are making behavioral choices, then the opportunity to impact behavior is maximized.

Challenges and Opportunities

The dynamic evolution of social networks over time necessitates an approach to PeCS in which data is gathered continuously and analyzed using models which capture the dynamics of evolving patterns of interaction. This is a significant challenge for traditional machine learning techniques which heavily leverage the iid assumption, i.e. that all data elements are independent and identically distributed. In practice, social interaction data will be coupled in time and will come from stochastic processes which may not be stationary. This implies a need to research data modeling techniques which go well beyond the standard models such as HMM which leverage the Markov assumption. In this context the study of semi-Markov models of stochastic phenomena which can capture more complex temporal dependencies is to be encouraged. One example in which dynamic phenomena come to the forefront is in the formation of social groups, for example when students meet for the first time at the start of the school year. The study of these emergent socialization phenomena is of great interest in psychology and sociology and could be enabled by PeCS at an unprecedented scale.

Another area with significant research challenges is the use of PeCS to influence social behavior, both collectively and at the individual level. This leads to the notion of developing socially adaptive systems, which can be viewed as a logical extension of the more common idea of context-aware computing. These adaptive systems need to be informed by the evolving social context in which their users are living out their lives.

Applications

The applications can be divided into three main categories. The first category has to do with understanding the way information and opinions build up and spread through social networks. One concrete example in the area of advertising is the identification of opinion leaders, which can provide the basis for viral advertising and marketing strategies. Currently such leaders can only be identified within fairly large-scale closed social systems such as Facebook or Amazon reviews. However there are very likely a much larger number of opinion leaders who exert influence on a much smaller scale which could be identified from the fine-grained interaction data available through PeCS. Other examples include the understanding of the build up or radical opinions, dissemination of emergency information, and political and social science.

The second application class is more research oriented and relates to using social networks as a tool for large scale data collection in PeSC systems. The idea is to leverage social networks to recruit volunteers for data collection and use them to collect and combine data to relate information from different users and to enrich sensor data with semantic information contained in the social network. In addition social networks can be leveraged as a means of crowd sourcing data annotation, and also as data sources for mining activity and context descriptions.

The third application class builds on the two described above. It is related to combining information from social networks with sensor information from mobile devices of many users to monitor and recognize collective phenomena and trends within whole communities. As a simple example, a correlation of long term GPS traces with locations of different retailers can be used to get information about consumer confidence and changes in it (do people go to WallMart or to Gucci). Using the social network this can be differentiated by social group and other related characteristics. Other examples include mobility patterns, health related life style attitudes, demographics, information relevant for urban planning or demographic developments. Applications are also easy to imagine in homeland security and disaster management. In a way we can consider the combination of PeCS with social networks leading to a virtual ‘’nervous system’’ of society. The general idea builds on the concept of Reality Mining as Proposed by Pentland et. al., going further with respect to emphasis on collective phenomena and considering complex interactions over different temporal and spatial scales.

All three application classes described above can lead to an entirely new type of social networking service where the boundaries between the real and the virtual worlds become increasing blurry. Understanding these sorts of applications and their implications is a highly relevant research topic.

Data Issues

The lack of availability of rich datasets is a significant barrier to research in this area. Most of the existing large-scale networks are closed, and in addition these networks are experiencing phenomenal rates of growth, to the point where the networks are growing faster than they can be crawled. In making an analogy to the speech recognition community, there is the potential for significant impact if funding agencies could contribute to the development of datasets and infrastructure for data collection. As an example, the development of a standard platform for measuring social interactions could facilitate the collection of data on a broader scale. Another productive direction is the collection and identification of best practices for data collection based on the experiences of some existing projects. Sharing best practices could help bootstrap the nascent research community in this area.

Scalability is a significant concern, based simply on the large quantities of data under discussion. In current practices, research data is often hosted on individual servers at research organizations, and must be copied across the network to support collaboration and data sharing. There would be significant benefit in coming up with a standardized system for hosting large scale data repositories on cloud computing sites so that they could be accessed by research teams directly, using services such as Amazon's, thereby saving the need for copying and replication. A research infrastructure program that facilitated this kind of large-scale infrastructure for data sharing would be of great benefit to the research community.

Privacy concerns are a significant issue that must be addressed in any discussion of social networking and the use of technology to measure social interactions. These concerns arise at a variety of levels. Large scale social network data make it possible to infer the actions of motives of individuals and are thus a significant privacy concern. Careful de-identification strategies and policies must be developed, and the risk of loss of anonymity is always present. In many data modeling applications the final fitted models are often based upon aggregations of data, e.g. the cluster means and covariences in fitting a mixture of Gaussians. It may be possible to do some initial data aggregation prior to model fitting, for the purpose of pooling individual records together and reducing the risk of loss of anonymity. In general, research into creating machine learning techniques which combine privacy preservation with the more traditional concerns of accuracy and generalization would seem to be a fruitful area of future research investment.

Breakout Session Report:
HCI

Christine Julien and James Landay

Participants: Peter Bajcsy, Roy Campbell, Niklas Elmqvist, James Fogarty, Livai Iftode, Santosh Kumar, Joshua Smith, Mario Sznaier, Svetha Venkatesh, David Wentzloff, Roy Want, Vincent Wong, Qiang Yang, and Pei Zhang

Introduction

Because pervasive computing applications are intimately integrated with our physical spaces, human-computer interaction (HCI) is an essential concern. This becomes increasingly true as pervasive computing systems increase in scale to include more devices, more capabilities, more humans, and more data. In this breakout session, we discussed the state of the art in HCI for pervasive computing systems, conceptual gaps that exist in considering HCI issues as pervasive computing systems scale, and potential research directions. Most of the issues discussed focused on evaluating the HCI aspects of these emerging systems since understanding the interaction between the human and the system comes directly from these evaluations. In this report, we review the salient aspects of our discussion, including recommendations to the NSF.

State of the Art

Within HCI, commonly used techniques are instilled with interdisciplinary connections from a variety of domains including psychology and anthropology. Techniques coming out of psychology are generally tailored for laboratory settings, and they do not scale well to situations outside of the lab. These methods do not scale over time, they do not scale to multiple devices, they do not scale to multiple locations, and they do not scale to high volumes of data. These aspects characterize pervasive computing applications, and the HCI aspects of these applications must be evaluated in situ.

Alternative techniques for investigating and evaluating HCI aspects of pervasive computing systems have built on anthropological techniques. These approaches do allow our experiments and evaluations to get out of the lab and into the real world. These techniques, such as ethnographic techniques, do provide a more direct measure of the interactions of humans and the systems under evaluation. For example, surveys on mobile phones have made some progress in evaluating pervasive computing systems in situ. By relying on these end-user devices, monitoring and sensing have become more pervasive, but we are not yet really using them to investigate HCI issues. In the end, these existing techniques still do not scale to the degree necessary for solid evaluation of HCI issues.

Existing Conceptual Gaps

We identified several conceptual gaps that provide opportunities for advancements in HCI as it relates to pervasive computing at scale. First, as described above, field studies provide a significant challenge in understanding good (and bad) HCI for pervasive computing. Clearly, direct observation approaches are more reliable than indirect ones (like surveys), but indirect approaches may be more scalable. In these indirect approaches that provide a reasonable starting point, you tend to lose quite a bit of context about the user, his or her intentions, and the context of the interaction. This gap identifies the need for new field study techniques that bring together aspects of both direct and indirect approaches.

Another issue in HCI research for pervasive computing systems 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 to use. At the same time, other aspects of scalability may (as is intuitive) make pervasive computing system less usable. This second gap demonstrates the need to identify and formalize these differences and incorporate them into the design of our emerging pervasive computing systems.

The third conceptual gap we identified addresses the need to be able to design, rapidly prototype, and deploy pervasive computing artifacts to get reliable feedback on the HCI issues of the applications. Specifically, we must determine how to go from pervasive computing application "pilot studies" to large-scale, meaningful studies. These challenges are further complicated by the common use of specialized hardware for these pilot studies because the hardware can be unreliable and hard to use, and these challenges are hard to mask with software. At the same time, we generally want to automate our evaluations to get a large amount of data with less effort.

Directions

In our discussions, we made several observations that open new research directions. Specifically, we believe that HCI is still an important aspect of pervasive computing, but it should be, ideally, an invisible aspect. That is, users should notice, as little as possible that they are interacting with a novel system. This is even more difficult than making pervasive systems easy to use; instead the idea is that pervasive computing systems should become a natural part of the environment. At any given time, even though a user is not explicitly interacting with a computer in the traditional sense, the user's implicit interactions with resources, devices, and capabilities in the pervasive computing system are naturally HCI. We do want the technologies to be invisible, but sometimes we want the data to become visible again; that is, we want the data in pervasive computing systems to be able to impact user behavior.

With respect to understanding and defining what good HCI is in pervasive computing at scale, two competing options exist: we can either train every user to use the same interface or we can enable interfaces to tailor themselves to each user's preferences. The latter is related to the field of adaptive user interfaces; challenges that have been identified in this emerging field include the fact that these adaptive interfaces have the potential to lose the consistency of the interface that users tend to naturally expect.

Finally, we also discussed the fact that HCI in pervasive computing systems at scale and the crowdsourcing that is becoming popular in pervasive computing applications may have significant relationships. Crowdsourcing may provide techniques to help applications and their users deal with the deluge of data. At the same time, crowdsourcing introduces new HCI issues, including introducing questions of how to deal with cultural differences among users.

Interdisciplinary Collaboration Potential

There are obvious opportunities for interdisciplinary collaboration, including traditional avenues for user studies like psychology and anthropology. At the same time, better integrating with application domains and interactions with experts from those domains could lead to interesting new interactions and a better understanding of the application issues. One example discussed was the collaboration with experts from the health domain. A major challenge of these collaborations is that these domain experts tend to expect real deployable systems before they are ready to interact on evaluation studies.

Experimental Opportunities

As HCI challenges often entail experimental studies, much of our discussion revolved around supporting large-scale experimental research. Specifically, we must understand how to move from our common "pilot" studies to large scale studies. We need to identify ways to automate evaluations so that we can get large amounts of data without excessive effort requirements. We must also be able to generate meaningful sets of test subjects that can be widely accessible to researchers. This requires identifying meaningful, unbiased motivators for these test subjects. To make sure the results from the test subjects are useful, the research (i.e., the system) must be reliable enough that people will use it. We not only need to be able to support these studies in terms of supplying resources, but we must also develop methodologies that make studies easier to perform in the first place.

Recommendations to the NSF

We identified three recommendations to the NSF that could support the breakout group's vision of HCI in pervasive computing at scale:

1. *Programs to support immediate stage scaling of studies.* At this moment, we are capable of small scale studies. If we have a reliable system, we have partners to do evaluations, but we need to first perform medium-scale studies to build the needed reliable systems. Therefore, we need support for this scaling process.
2. *Recognition of the type of funding needed to perform this work.* These scaling studies are more traditionally considered "development efforts" that require engineers for system development. These should be recognized as part of the research process, given the fundamentally different nature of HCI research. Joint programs with industry may provide a potential avenue to achieve access to data and development resources.
3. *Support of interdisciplinary pervasive computing HCI projects*. It is essential that we broaden the number of researchers involved in and aware of HCI issues at scale.

Breakout Session Report:
Energy Analysis, Harvesting, and Storage

Shwetak Patel and Vijay Raghunathan

Participants: Karthik Dantu, Prabal Dutta, Brian Kelley, David Liu, Vinod Namboodiri, David Wentzloff, and Gil Zussman

1. Introduction

One of the biggest challenges of realizing Mark Weiser's vision of pervasive computing as a near-invisible technology that does not constantly require human intervention is the problem of powering the thousands of pervasive computing devices that are expected to be embedded in our surroundings and everyday objects. Most of today's pervasive computers are battery powered. Next-generation pervasive systems will be expected to operate for several months to years without the need for battery replacement, because frequent battery replacement for hundreds of devices is not only infeasible, but also means that the devices are not "invisible." Given that battery technology is improving at a slow rate, the limited battery capacity of these systems will pose a significant challenge in the viability of these systems. A promising and viable alternative for powering next-generation pervasive micro-systems is to scavenge energy from ambient sources such as solar radiation, vibrations, radio frequency transmissions, or thermal gradients (we refer to this as 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. Unlike the slow trends in battery capacity, the rapid advancements in electronics, embedded systems, and IC design has enabled the new possibility of practical power harvesting.

2. Current Status and Challenges

While the notion of energy harvesting has been extensively explored in the context of large systems such as solar farms and windmills, micro-scale energy harvesting, as a systematic discipline, is not as mature. Current state-of-the-art in micro-scale energy harvesting is limited to various research prototypes, which represent point solutions rather than generalized designs. Often the resulting point solutions are inherent in the very nature of the design of power harvesting systems, because they are often tied to specific physical phenomena and applications.

Realizing highly efficient micro-scale energy harvesting systems is challenging due to three main constraints. 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 1V. For example, miniature photovoltaic cells and thermo-electric generators produce voltages in the range of 0.2-0.6V. 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. It is, therefore, 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. Further, energy availability can be intermittent in nature. Pervasive computing systems that are powered by these micro-scale energy harvesters should be able to adapt to such vagaries using intelligent resource management techniques. We have to re think the design of embedded systems to address many of these new challenges, which we taken for granted in the design of traditional battery-powered solutions.

3. Key Research Directions

Overcoming the challenges described above requires a concerted research effort involving all layers of the design hierarchy, ranging from devices, circuits, and architectures to power management algorithms, design exploration frameworks, and new networking protocols. In particular, we believe that the following research directions (and questions) are key to the analysis, design, and management of environmentally powered micro-scale systems:

1. 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. It is important to then consider the question of the right metric 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, output power, etc. Currently, there is no way to compare various solutions.

2. Efficient power extraction from transducers and energy storage: A variety of interesting energy harvesting transducers has recently become available, such as thermoelectric generators, piezo, and photovoltaics. It is important 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. There is also a need for energy transducer models (circuit level models, higher-level models that are parametrizable and can be simulated) to abstract away the transducer devices while designing higher layers of an energy harvesting system. A key aspect to developing these models is to decide on what key information needs to be captured by these models (dynamic range of available power, temporal and spatial dynamics, etc.) 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. Key research directions here include exploring new energy storage architectures that synergistically combine heterogeneous energy storage elements (e.g., thin film batteries and ultra-capacitors) to minimize losses during energy storage.

3. Efficient HW/SW Systems, Algorithms, and Resource Management: This research direction involves the design of efficient hardware and software systems that are "harvesting aware." In addition to pushing the limits on ultra-low power design (through techniques such as sub-threshold design, low-leakage memory, etc.), 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 exposes various control points (e.g., the amount of energy currently available from the transducer) to software. At the network-level, key questions that need to be answered include how to build large networks out of intermittently available devices? How do these devices bootstrap and join a network? Are concepts from delay tolerant networking applicable and useful here? This brings up new questions about how interconnected devices are synchronized in this kind of environment. Can we exploit the inherent asymmetry between different devices in the network due to spatial variations in energy availability? In other words, can we leverage the fact that certain devices will have more available power than others and they coordinate between each other to accomplish a task? How can we efficiently and accurately predict future energy availability and use that information for resource planning?

4. Enabling systematic design space exploration: This research direction involves the creation of design methodologies and tools that enable systematic design space exploration of micro-scale energy harvesting systems. This thrust is essential to transforming the study of micro-scale energy harvesting systems from an art that relies on designer intuition into a systematic science. The first step in doing this is to develop simulation models of various system components (e.g., various energy transducers, power converters). These models can then be used to create simulation tools that allow designers to quickly evaluate the impact of various design choices and parameters while architecting micro-scale energy harvesting systems. We also need to develop hardware/software building blocks that allow researchers to quickly prototype power harvesting applications. Many of the interesting insights will come from actual deployments and the use of the technology. In addition, algorithm developers will need platforms to conduct their research. These platforms could come in the form of hardware modules or development kits.

5. Educating the next generation of CS/EE researchers and engineers: From an education perspective, a crucial requirement is to teach students to consider energy availability as a first class design metric and citizen, along with other metrics such as computation cycles or memory. As one of the participants in the breakout session at the workshop succinctly put it, we need to train the next generation of students to also consider and deal with “energy underflow” along with “stack or buffer overflow.” We are not proposing a radical new curriculum or new classes, but rather infusing these concepts in the existing classes.

Summary

It is expected that significant advances in the research directions described above will allow us to fully realize the potential of micro-scale energy harvesting and greatly reduce the reliance on traditional batteries in next-generation pervasive computing systems, removing one of the biggest showstoppers to their large-scale adoption. Successful completion of the proposed research vision will also achieve significant environmental impact by greatly reducing the large number of batteries that are discarded every year. In addition, the inherent optimization innovation required for this form of computing will provide significant insights into other areas of computing and electrical engineering.

Breakout Session Report:
Intelligent Transportation / Vehicular Networks / Aerial Networks

Liviu Iftode and Mohan Trivedi

Participants: Mohan Kumar, Thomas Little, Vassilis Kostakos, Vinod Namboodiri, Doug Terry, Mads Haahr, Wang-Chien Lee, Kishore Ramachandran, Roy Campbell, Niklas Elmqvist, Joshua Smith, Mohan Trivedi, Pei Zhang

Introduction

Intelligent transportation, vehicular and aerial networks are defining components of the pervasive computing vision. Research in these fields is driven by five factors: mobility, safety, environments, infotainment and convenience. Vehicular networks have the potential to become the 2nd largest pervasive computing sector after the smartphones. Intelligent transportation cannot fulfill its potential without ubiquitous vehicular networks. Aerial networks are also emerging as very dynamic mobile ad-hoc networks with specific challenges and opportunities.

State of the Art and Existing Conceptual Gaps

Vehicular networks

Significant amounts of research have been conducted so far, mostly from a networking perspective. Less covered are aspects situated at the border with other fields such as computer vision, human factor, privacy and social aspects. Research has been dominated by industry, which shares little with academia because of the fierce technological competition. Academic research must identify its strengths and define its research agenda accordingly in order to stay relevant. Research in Europe and Asia is ahead. V2V, V2I and cellular are the major communication models in vehicular networking. Wireless networking covers both inter-vehicle as well as for intra-vehicle communication.

Intelligent transportation

Intelligent transportation aims at reducing congestion, fuel consumption and, to a much limited extend so far, the emissions. Multi-modal transportation models are emerging. Economics can provide powerful methods to understand a driver's behavior. So far, the human aspects have not been sufficiently addressed.

Aerial networks

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 unprecedented size and scale. They will enable new classes of applications including commercial pollination, search and rescue, surveillance, environmental monitoring etc.

Challenges and Approaches

Human-centric vehicular computing

Pervasive computing technologies hold major promise to enhance automotive safety by introducing a new range of “human-centered” driver assistance systems. Human centric, pervasive computing environments with integrated sensing, processing, networking, and displays provide an appropriate framework to develop effective driver assistance systems.

One of the key requirements in the design of an active safety system is the ability to accurately, reliably and very quickly identify the conditions that would lead to an accident and to induce corrective actions so that the accident can be prevented. Therefore, research in human factors is crucial.

Opportunity to study pervasive computing at scale

Vehicular networks have the potential to become a real-world pervasive computing testbed at very-large scale, allowing research to investigate interesting human and social aspects related to the adoption of pervasive computing. Including smartphones in the vehicular networking and intelligent transportation infrastructure can accelerate the penetration of the latter, and reduce the dependence of the research on the automakers. Intelligent driver support systems may provide an ideal application domain for addressing some of the challenging multidisciplinary research problems in pervasive computing.

Closing the loop between sensing and control

A long term vision is to use modern wireless technology, environment monitoring, and urban traffic management to “close the loop” between urban sensing and vehicle route control with the aim of simultaneously reducing congestion, pollution, and traveler delays. A systematic approach to traffic management requires a solution that must rely on real-time collection of traffic density and air pollution data, and it must feature real-time communication mechanisms for fine-grained traffic.

Today, there is a broad spectrum of largely disconnected solutions to alleviate traffic congestion. For instance, traffic lights, on-board navigators, and city traffic center do not talk to each other. The challenge is to connect existing solutions via state-of-the-art communications and networking to provide efficient, coordinated real-time traffic and air quality control. The closing of the loop between traffic and air quality data sensing and vehicle routing will enable an urban traffic management that can adjust to the rapidly changing traffic and air quality conditions typical of large cities.

To realize these goals, effective and efficient techniques for gathering data from the transportation system needs to be developed allowing anywhere, anytime use, and fusing it into usable information for pedestrians, vehicles and drivers.. The collection and processing of ubiquitous traffic information is essential for implementing proactive control strategies and getting feedback on their effect. This vision raises multiple challenges in the areas of urban sensing, emission models, traffic and pollution simulation, traffic management, control and enforcement, wireless networking, usability and human-computer interfaces, and security and privacy.

Data driven research

As in other pervasive computing domains, there is a need to shift from computing to data driven research. Sensing traffic, air quality and driver's behavior produces huge data-sets, which must be scientifically analyzed and aggregated in order to be meaningfully and usefully presented to the driver, and to the traffic control systems. Systematic efforts are required to understand and characterize driver behavior, situation criticalities, interactivity patterns from real-world, distributed, and multi-modal massive sensory datasets. Learning, classification and prediction can be applied to driver behavior, activity and intentions as well as for vehicle-human interactivity and vehicle trajectory patterns. At the same time, gathering correlated traffic, driver and environment datasets, along with their processing may enable informed formulation of new traffic management policies and regulations for the era of vehicular networking and intelligent transportation.

Multidisciplinary approach

Research in vehicular/transportation/ aerial networks is particularly inter-disciplinary within the field of computer science as well as multi-disciplinary. Within the computer science, future vehicular computing and intelligent transportation solutions will require knowledge and research in networking, data processing, real-time distributed computing, security and machine learning. Multiple disciplines other than computer science, such as psychology, cognitive sciences, transportation research, atmospheric sciences, social sciences and economics, can contribute to the complex understanding of the transportation problems and solution space. Covering all these aspects require large collaborative efforts, which are difficult to assembly, manage and fund. Education and the job market do not sufficiently address and value multi-disciplinarily skills.

Pollution modeling

Pollution modeling is a major challenge. Pollution is non-uniformly distributed, being on average higher in close proximity to roadways, but vary widely on time scales of hours, which can change exposure levels dramatically. Pollution levels near roadways depend on traffic density, vehicle speeds, congestion, and local wind speeds and direction. More precisely, air pollution can vary on length scales of tens of meters for some pollutants, but the distribution of pollutants on this scale is poorly characterized due to lack of spatially resolved measurements. The spatial heterogeneity arises from the interplay between the complex topography, the variable atmospheric mixing and the highly non-homogeneous emissions. Thus, the potential for mapping of pollutants with high spatial resolution via sensors integrated into a smart traffic sensing system is largely untapped, and will likely produce insights beyond those currently available. This will require developing more accurate models, which will be simulated in real time to provide input for traffic control.

Multi-modal transportation

Complex transportation problems cannot be solved within the isolation of a single transportation mean. Urban transportation is essentially multi-modal and future intelligent transportation solutions must address traffic and pollution issues assuming multiple cooperative and competitive transportation means. Lack of access to data in particular of synchronized data across multiple transportation platforms, is a serious obstacle for academic research. Sustainable metropolitan transportation is one of the stringent goals of multi-modal transportation exploration.

Social aspects

As long as humans are still in control of their own transportation decisions, any relevant research in intelligent transportation must include human and social aspects and must identify the right methods to advance research in these directions. Smartphones might be the right interface to capture human-to-human interaction and influence with respect to transportation. Vehicular social networks will likely emerge as ad-hoc communities of regular participants and can benefit from information sharing.

Smarter roads and active highways

“Smarter” connected cars alone are not sufficient for realizing the grant vision of \ intelligent transportation. The roads (i.e., transportation infrastructure) also need to become smarter by equipping themselves with timely pervasive transportation, gathering information from smart cars and their occupants.. For instance, given the knowledge of an on-going congestion up front on its route to the destination, a smart car, enabled by the en-route driver information service, may inform its driver about the situation and recommend her to make a detour.

A long-term, highway concept will evolve from a transportation infrastructure to a cyber-physical system that will shift from global traffic management to individual vehicle routing, similar to air traffic control. In this sense, highways will become *active* managers of their own traffic similar to air traffic control. Future highways and future vehicles will communicate with one another, making the highway system aware of the drivers’ travel plans and allowing it to cooperate with, and actively instruct, the driver to achieve them. In particular, *active highways may* allow drivers to reserve slots in special high-priority intelligent lanes. This fine-grained traffic management model will guarantee travel time bounds, handle exceptions and enforce global community and environmental policies using real-time information from vehicle- and infrastructure-based sensors.

Research in active highways will be helped by autonomous vehicles, too. To this end, research must also include architectures for cooperative active safety systems utilizing vehicle-to-vehicle, and vehicle-to-infrastructure communication channels which can support large scale real world traffic conditions.

Participatory sensing for pervasive transportation services

Efficient collection of fresh and timely pervasive transportation heartbeats, transforming them into useful information that reflects the real-time road conditions will be very important. Roadside sensors, and probing vehicles, are two different approaches typically adopted for collecting data. With recent development of participatory sensing activities, participatory sensing via volunteered probe vehicles has a great potential to provide complementary coverage to the roadside sensor approach. Research effort is needed to incorporate participatory vehicle sensing, and address various technical issues that arise. We envisage the need for developing a pervasive transportation services framework, which consists of three components corresponding to data collection, data processing/mining, and data dissemination.

Aerial networks challenges

Aerial networks present distinct challenges in computing at scale. Unlike traditional sensor networks, such swarms not only collect data but need to have the ability to make decisions using the data collected in real time. Network connectivity is much more dynamic making sharing data harder. Actuation (flight) is much more expensive in terms of energy requiring a rethinking of trade-offs to be made in terms of using communication to make better actuation decisions. Actuation also causes a lot more uncertainty than sensing, requiring design of robust algorithms that work both at scale and uncertainty. Lastly, some of the classic sensor network problems still persist like tasking a set of nodes to perform a given task.

Unique research directions that need to be addressed for this field:

1. Revisiting the energy tradeoff: Conventional sensor networks research has assumed that communication is the bottleneck for energy efficient operation. This has led to design of mac/routing protocols that minimize communication overhead. However, in such systems, actuation is an order of magnitude more expensive than communication. Therefore, it might be more beneficial to communicate more information in the hope of making better actuation decisions.

2. Information dissemination: Most applications that utlizesuch systems are in space. Most applications either have a map of this space or explore the space to discover features of interest and work on these features of interest. Given this paradigm, it is evident that such discovered information needs to be rapidly shared among the swarm for efficient operation. However, given the unreliable nature of communication between aerial vehicles, it is very important to determine the information that is useful before dissemination along with intelligent aggregation. Research in multi-robot systems has mostly used probabilistic frameworks, and formal methods, to fuse/propagate information. We need research to make such algorithms work at scale.

3. Collaborative sensing. Aerial networks are essentially a collaborative sensing and distributed control problem applied to an ultra-light, miniature flying sensor network. Given the cost and energy constraints, truly lightweight distributed protocols must be designed. This vision explores the new dimensions of collaborative-mobility, multi- level distributed processing, and collaborative sensing with minimal sensor capabilities. An in-depth investigation into limited capability mobile sensors is necessary, and will have a lasting impact that extends beyond today’s sensor systems.

Mechanisms to Improve Data Availability for Experimental Research

Large-scale testbeds, simulators and relevant datasets are critical for influential research in intelligent transportation. Smartphones, that already have considerable market penetration can be used as a stop-gap platform to create large-scale vehicular networks. Crowdsourcing may also serve as a palliative method of research for data collection.

Educational Opportunities and Challenges

Intelligent transportation is an opportunity for inter-disciplinary education, which has not been sufficiently encouraged. Addressing job opportunities in order to make the field attractive to students is essential.

Recommendations to NSF

Three major recommendations to the NSF have been made:

1. A cross-cutting program in vehicular/intelligent transportation issues

2. Facilitate collaboration with industry, as well as international collaboration

3. Coordinate with other national research programs, such as the Transportation Research Board of the National Academies, Strategic Highway Research Program (SHRP 2), the Research and Innovative Technology Administration of the U.S. Department of Transportation, and the U. S. Department of Transportation Intelligent Transportation Systems, Joint Program Office.

Breakout Session Report:
Smart Health

Diane Cook and Santosh Kumar

Participants: Michael Anderson, Oliver Brdiczka, Andrew Campbell, David Chu, Sajal Das, James Fogarty, Mario De Francesco, Ahmed Helmy, Christine Julien, David Kotz, Narayanan Krishnan, James Landay, Mirco Musolesi, Wendi Nilsen, Vijay Raghunathan, James Rehg, Mahadev Satyanarayanan, Andreas Savvides, Mario Sznaier, Roy Want, David Wentzloff, and Feng Zhao

Introduction

As acknowledged in the "Smart Health and Wellbeing (SHB)" cross-cutting program at NSF, information and communication technologies are poised to transform our access to health information and participation in our own healthcare and wellbeing. The long term vision is that smart health will simultaneously reduce the cost of health care, and improve our health by encouraging healthy behaviors, reducing susceptibility to diseases, cutting down visits to health care providers, and intervening in real-time as a result of mobile monitoring.

Current Status

Significant progress has been made in the last half decade in smart health. These advances include increasing adoption of electronic health records and personal health records; availability of web-based monitoring of personal health; commercial availability of wearable sensors and their integration with smart phones; monitoring and management of physical activity in the mobile environment; and increasing support for aging in place, to name a few. These initiatives have fostered new collaboration among the computing, medical science, behavioral and social sciences, and healthcare provider communities. These encouraging developments, however, 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, and among healthcare provider and patients, all of which is accelerated by support and encouragement directly from the President of the U.S.

Vision and Challenges

Smart health can help improve the health and wellbeing of society in several meaningful ways. Some specific vision ideas are listed below together with a discussion of the challenges that need to be overcome in realizing these visions.

1. Ubiquitous accessibility to caregiver and availability of health information. Health 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 where individual health information is readily available to that person (or their designee) anytime and anywhere, and which can be readily shared with a provider who may be located anywhere. Electronic health record (EHR) and personal health record (PHR) are enabling technologies for such a vision, but several major challenges such as security, privacy, standardization of interfaces, data processing and visualization, etc. others need to be addressed to realize this vision.

2. Attention to mental health. There is now a growing awareness about prevalence of mental health issues in our society. These issues include depression, autism, post traumatic stress disorder (PTSD), chronic stress, and cognitive decline, to name a few. Smart health technologies can enable screening and treatment for mental health problems. 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, preserving the privacy of participants during treatment, and evaluating the efficacy of treatments.

3. Delivering timely intervention. Smart health offers a unique opportunity to deliver intervention when and where it is needed, especially if it can be delivered via a cell phone. In addition, the need for intervention may also be detected automatically using smart health technologies. Further, technology can be used to monitor and encourage physical activity, which may help reduce the trend toward obesity in this country. 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 much needed sensors, algorithms, models, and user interfaces.

4. 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.

5. Post-Disaster support. Smart health could also be applied to provide timely support to victims of a disaster. This may include on-the-spot training and instruction to willing volunteers to turn them into caregivers, and monitoring of vital signs using readily available technologies embedded in smart phones (e.g., camera), etc.

In summary, we envision a future where smart health technologies embed themselves in the infrastructure, in our environment, in the fabrics we wear, and in mobile devices we carry, thus becoming so pervasive that they essentially disappear from our explicit cognition. At the same time that we become less aware of them, they become more attentive to our health needs anytime and anywhere.

Recommendations to NSF

There are two unique aspects of smart health that distinguish it from other technologies - every person needs health care irrespective of their educational background, and it directly affects our health. The recommendations we make in the following section are designed to support the realization of the vision outlined above while being cognizant of the unique role smart health plays in human society.

Research Support

1. Scaling of smart health technology. Smart health technology needs to scale so it can be adopted widely by both general public and by scientists in scientific studies of health issues. Doing so requires supporting research that will reduce the cost, minimize user burden, improve usability, extend energy efficiency, and improve robustness for more reliable measurements in the noisy mobile environment.

2. Validation of smart health tools. Since smart health tools affect our health, it is critical that they do not provide erroneous information that could lead to faulty diagnosis or treatment. Therefore, it is critical to support research that seeks to not only develop smart health tools but also establish their validity in a variety of real-life environments, and on diverse groups of people.

3. Behavior modification. Since smart health offers an opportunity to deliver interventions in real time on cellular phones, new research needs to be supported to develop, evaluate, and validate novel behavioral interventions to a range of health challenges and conditions, such as stress, depression, obesity, social anxiety, addictive behavior, and other health issues.

Infrastructure Support

Research on smart health is expensive because it has to be applicable and validated for real-life adoption, while ensuring it does not affect health adversely. Infrastructures need to be developed to foster research on smart health from the technology community. We recommend support for the following three types of infrastructures.

1. Open extensible platforms for smart health. The sensor network community greatly benefited from the open hardware mote platforms and the associated TinyOS platform. These platforms made it possible to experiment with new sensors and test various innovations in real-world conditions, sometimes even to learn more about the real-world conditions. Smart health needs similar open and extensible platforms for both wearable sensors, and a mobile phone software framework, that is affordable and accessible to the larger scientific community.

2. Testbed - A cohort of subjects. A testbed for smart health cannot simply be a collection of devices, as is the case for contemporary computing test-beds such as Planet Lab. Smart health is about data collection on large groups of people to study individual differences within populations. Any experiment with human subjects requires the time, effort, and expense associated with IRB approval and the consent process. This is a significant roadblock to development of smart health technologies. We recommend that cohort of subjects be established to address these issues. A cohort is a proven method for large scale experimentation with human subjects in behavioral and medical disciplines where various issues for recruitment, retention, and management of cohorts have been worked out over the past several decades. Usually a board is formed to approve any requests for access to a study cohort. There are numerous benefits to having a cohort, such as reduced cost of additional studies, availability of rich history of subjects in the cohort, and subjects are known to be compliant.

We acknowledge in this discussion that some fairly established smart home and sensor-rich test-beds do exist. These include the Georgia Tech Aware Home, the Gator Tech Smart Home, the MavHome, the CASAS test-beds, and the iDorm. A few of these do make datasets available to the public but a much larger set, and diversity of test sets, is needed.

3. Access to health datasets. Datasets on real people (such as the MIT arrythmia dataset and the PhysioNet dataset) are critical to the development and evaluation of various algorithms and models in smart health. We believe that access to collections of data sets on smart health – available to researchers while preserving the privacy of data contributors – is critical to advancing the field of smart health. Hence, we recommend supporting collection, hosting, and management of such data sets.

Educational Opportunities

Smart health offers us a new opportunity to transform health education. The technologies that already exist, and those that will be developed, can be used to make health education more accessible and hands-on in K-12+ educational settings. We recommend supporting such educational endeavors.

Smart health research is encouraging collaboration among multiple new disciplines, but we lack a common vocabulary. Supporting multidisciplinary courses that orient practitioners, domain scientists, and engineers to work together and develop a common vocabulary will facilitate the work of the next generation.

Joint Programs with NIH

Supporting health care research has traditionally been the domain of NIH. However, smart health falls on the boundary of basic research and its application to health care. NIH usually supports use of technologies that have been validated for use in applied research, but has not been receptive, traditionally, to the development of technology itself. An exception to this is the Genes Environment Initiative (GEI) program, which was a time-limited technology development program. This creates a gap between technology development and its eventual adoption for health care research and practice. We recommend development of joint programs between NIH and NSF to help support the development and validation of smart health technology so it can readily be transitioned to large scale adoption in health care research and practice. Validation of health technologies is usually expensive due to human subject involvement, and hence we believe there is a need to pool resources across both agencies to support this critical stage in the advancement of smart health.

Breakout Session Report:
Sustainability and Energy Management

Vincent Wong and Brian Kelley

Participants: Geoffrey Challen, Prabal Dutta, David Liu, Jie Liu, Klara Nahrstedt, Vinod Namboodiri, and Hari Sundaram

Introduction

The research in the area of sustainability and energy management aims to reduce the energy usage and to improve the energy efficiency of the system by monitoring and control. Continuous monitoring often occurs remotely, far from the central facilities, via networks of low power sensors, advanced metering infrastructure (AMI) that interacts with smart appliances and home area networks (HANs), and integrated wireless communication infrastructure and services. Sustainable systems typically rely upon heavy penetrations of renewable energy sources with intermittent power generation characteristics. It is an inter-disciplinary area and covers different topics including home area networks, smart grids, vehicle electrification, and low power wireless sensor networks. In addition, some complex problems only emerge at scale, with emulation test benches.

State of the Art and Existing Conceptual Gaps

We classify the state of the art into the following complementary areas:

Wireless Sensor Networks

Over the past decade, there have been significant advances in our understanding of wireless sensor networks. This includes the design and implementation of operating systems, data aggregation techniques, and scheduling algorithms. One of the challenges is to make use of the results of wireless sensor networks for energy management. In addition, the waking-up from sleep mode, and scanning for signals, is effectively reducing the need for sensor device power consumption.

Home Area Networks

A home area network enables different devices and appliances to communicate with each other. Various standards (e.g., Zigbee) have also been approved. For energy management, there are also different products available in the market which allows home users to monitor their electrical usage. With the use of AMI, there is ample opportunity for autonomous closed-loop demand-response systems.

Metropolitan Area Smart Grid Systems

The development of the smart grid is still in an early phase. A variety of technological innovations from different fields and disciplines such as pervasive computing, power engineering, system control, and communications are required to enable the smart grid. Future smart grids will continuously monitor aggregated streams of power information, along with dynamic load profiles; and manage vehicle electrification and battery systems, monitoring several generations of renewable energy sources at multiple sensor scales (macrocells, microcells, and picocells); and control power flows from increasingly distributed renewable energy sources. Future smart grids will also directly and indirectly control loads in homes, buildings, and across metropolitan areas via closed-loop demand-response mechanisms. Finally, the integration of large wireless sensor networks in closed-loop demand-response systems will require new wireless sensor network control models in order to scale.

Data Centre

The energy cost of a data centre is much higher than the energy cost of a typical commercial building. Thus, it is important to design energy efficient components and systems for a data centre (e.g., with better cooling systems, less energy consuming processors). Data Center themselves will also host future pervasive energy services, new management protocols, and will offer enhanced energy optimization services.

Challenges and Approaches

In this section, we outline the major challenges of developing protocols and algorithms for effective energy management.

Energy and Power Consumption of a City

For energy management, one of the challenges is the use of pervasive computing and handheld devices to measure and monitor the power consumption (e.g., carbon footprint) of an individual as they move and interact with the environment. This is particularly true of environments where groups of individuals inherently share energy resources. This usage should be combined with the individual's home usage as measured by AMI. A proper metric is required to provide a meaningful quantitative comparison. By applying statistical techniques (e.g., sampling/polling techniques), one can then infer the power consumption of the population in a city. If the data of the power consumption of an individual is available to his/her friends (e.g., through a social network such as Facebook), it may provide an incentive for people to reduce their own power consumption. If the data of the power consumption of similar or neighboring cities is also available, the results can provide recommendations for adopting techniques to reduce energy usage. Human adaptation based upon recommendations is likely to lead to altered human behaviors, which may not be easy to capture in models.

Closed-Loop Demand-Response Management

With the use of an AMI meter, demand response management enables individual users to control and schedule their devices and appliances based on real-time pricing and forecast. This allows the users to shift high-load delay-tolerant applications (e.g., charging an electric vehicle) to off-peak hours in order to reduce the peak-to-average ratio load demand. A grand challenge posed by this breakout group is to monitor and control the closed loop demand response of an entire city to within 5% of a target reference power with significant dynamic power variability due to high penetration rates of renewable sources.

Reuse of Surplus Energy

Some households have begun to deploy renewable energy sources (e.g., solar panels). By installing an array of wind turbines, a community can also set up a power microgrid. Various microgrids can be connected to each other, or via the smart grid. Pervasive computing and smart grid can be used to propagate the signaling information and allow users to aggregate the load, and to trade or transfer power between each other.

Aggregation of Vehicles' Batteries at Scale

It is envisioned that the penetration of electric vehicles will be increased in future. During the daytime, when there is a large number of electric vehicles parked, their combined battery power can become a time-varying energy source. With proper control and pricing mechanism, this energy source can support frequency regulation in the electric grid by providing a regulation service, up or down, to meet the demand. A challenge here is to develop a pervasive computing systems that can aggregate and control electric car battery storage, and other customer or utility storage systems over a metropolitan area, to behave as a large virtual electricity storage device for the future smart grid.

Microgrid Systems

It is desirable to achieve a local power generation capability and services which are detachable and independent from power tied to the grid.. Microgrids can act as an island with a set of loads, energy storage, and dynamic adjustable capability, which can follow the intermittent power profile constraints of local renewable energy sources. The challenge includes sensor monitoring and autonomous control of distributed clusters of microgrids. Future smart grids will consist of many distributed microgrid systems, each containing a significant component of renewable energy, electricity storage, sensors, and AMI meters attached to each building. Grid-tied power conduits should be minimized.

Wireless Network Control of Metropolitan Area Power

Macro-cell area dynamic monitoring, and optimization of smart grid power, can be based upon large distributed sensor network load measurements. When very high penetration rates of renewable energy systems create sizable fluctuations of bus power, or transmission line failures occur, pervasive computing systems can be used to determine the optimum network control in order to alter dynamic power flows across spatially large distribution systems. ata from multiple sensor points can usefully be gathered to control these systems even when power distribution is connected in complex topologies , and each microgrid resulting in variable dynamic power generation. The goals include dynamic stabilization of the smart grid, peak power shifting, power cost minimization, and carbon based metric optimizations.

Mechanisms to Improve Data Availability for Experimental Research

Large-scale Testbed: It would be beneficial for the researchers if NSF can collaborate with power utility companies to create a large-scale national testbed (e.g., interconnection of microgrids). The scale can be similar to the PlanetLab project.

Simulation Model: Although there are simulators available for different disciplines (e.g., networking, power systems), there is no simulator that can simulate the behavior in both the transmission lines and communications infrastructure.

Analytical Model: Besides simulation models, stochastic models with standard performance metrics, and standard sets of profiles for energy efficiency are also crucial for performance evaluation and comparison.

Centers of Excellence Involving Utility-University Partnership: Utilities would partner with one or more university partners to form Centers of excellence in energy management for sustainable energy technology, aggregated HANs, power information systems, sensors, and dynamic control for energy management.

Educational Opportunities and Challenges

New course in "Sustainable Power Networks": Currently, most of the undergraduate students who are in computer science, electrical engineering, or computer engineering do not need to take courses in power systems (and vice versa). It is important to design a new undergraduate course, which covers the fundamentals in power systems, communications, protocols, security, wireless networks, and energy efficient smart grid design principles. This will allow the future system developers to understand, consider, and apply power and energy behavior in the design of the power/communications system.

Increase the Breadth in a Different Area: Researchers and graduate students with a background in computer science and computer engineering should either attend courses in power systems or work closely with power engineers so that they can acquire a better understanding of the requirements in power systems.

Recommendations to NSF

Partnership with Utilities to Enable a Large-scale Testbed: NSF can actively collaborate with utility companies to create a large-scale national test-bed (e.g., interconnection of microgrids) and Utility-University combined Centers of Excellence.

Access to data from utility companies: NSF can discuss with Utility companies and Utility-University Centers of Excellence group terms for generating anonymous data on pricing information, load demand curve, etc. The data is important to evaluate the practicality of different demand response algorithms.

Workshops: NSF can organize an exchange program, or workshop, between academics and engineers working in power utility industries. This will provide an opportunity to exchange ideas and increase collaboration.

Breakout Session Report: Theoretical Foundations

Justin Shi and Gil Zussman

Participants: Ioannis Stavrakakis, Gustavo de Veciana, Svetha Venkatesh, Bill Schilit, Gang Zhou

Introduction

This breakout session focuses on theoretical foundations for pervasive computing at scale.

State of the Art

The session attendees agreed that Pervasive Computing at Scale (PeCS) is not yet a mature theoretical field and that at this stage there is no strong theoretical foundation for PeCS.[1] Namely, there are almost no specific metrics (e.g., similar to scaling of network capacity in networking) or analytical tools that are tailored for PeCS and that can effectively deal with the scalability issues. Moreover, existing theoretical elements in related areas have not yet been adapted for Pervasive Computing (e.g., Fitt's law, used in UI design, is a type of analysis that does not easily carry over to mobile devices and new interfaces). In general, PeCS theory builds on theoretical contributions in other areas such as HCI, networking, machine learning, control theory, mobile computing, and social networks, but has yet to be solidified into a unified theory.

Conceptual Gaps

Several gaps have been identified, including:

- In relation to the state of the art mentioned above, lack of a unifying theory has been identified. Such a theory should be able to deal with scalability to very large numbers and to a variety of devices, applications, and interaction methods (as a few examples). In particular, since PeCS necessarily exacerbates the current scalability challenges in system architectures, there is a need for insights regarding the scalability as a function of the number of nodes, number of users, and the amount of data. Moreover, performance measures other than speed should be considered.

- Related to the item above, the lack of theory supporting power consumption optimization and energy efficient operation has been identified as a significant gap. It has been stated

[1] A specific example that demonstrates that the field is relatively application-oriented is that while IEEE has transactions and magazines on various topics (networking, wireless communications, control, etc.), there is only one magazine for Pervasive Computing.

that having a measure for power consumption similar to the complexity measures (e.g order, $O(\cdot)$) would be useful and can provide directives to efficient power aware programming.

- There is a need for better understanding of interactions among components, resources, and humans. Namely, there is a need for theory that would support the understanding of emergent behavior. Since the human interactions with the devices are interleaved with interactions among the devices and among humans, there is a need for theoretical tools that will take the users, and their interactions, into account.
- There is a need for new adaptable machine learning techniques that can operate on large datasets without supervision. The underlying theory for real-time useful information collection is yet to be developed. Moreover, tools to exploit temporal dimensions for controlling noise and unnecessary data acquistion (without relevance) are needed.
- There is still a need to understand the motivational forces for PeCS (health, economics, politics, religion, personal preference, etc.). In particular, it was mentioned that there are probably larger forces than Mark Weiser's vision. There is a need to understand what these forces are.
- There is a need for measures of Mark Weiser's "disappearing effect".
- There is a need for extensions of Fitt's Law for Pervasive Computing environments and emerging interfaces.

Potential Collaboration

The natural collaborations are among the disciplines of computer science, engineering, and the physical sciences. The session participants emphasized the two latter disciplines, since many of the technological innovations that require the development of new theoretical foundations stem from research performed outside the computer science community (e.g., the development of touch screens, advanced wireless communication techniques, advanced sensing techniques, and battery and energy consumption optimization techniques). Within the computer science field it has been emphasized that attention should be given to issues such as privacy and security, due to the ubiquity of the devices. The aspects of privacy and security should probably be studied from a legal point of view.

In addition, collaborations with social science disciplines such as psychology, sociology, anthropology, and economics should be encouraged, since, as mentioned above, the users and their interactions have to be taken into account in the development of appropriate theory. The session participants felt that social scientists can help the community to ask the right questions that would lead to important theoretical contributions.

To conclude, the session participants believe that PeCS has an inherent broad appeal. Therefore, wide collaborations (consortium-like) engaging computer scientists and engineers as well as social scientists, legal professionals, regulatory bodies, and government agencies would help to bring together concepts and ideas, and to identify the important theoretical concepts.

Educational Activities

Along with the development of PeCS theory, there will be a need to enable the students to understand the motivational forces beyond the technical developments. This can lead to changes from "teaching" to "facilitation" such that the students will not only learn how to do things but also understand the underlying theory. In particular, foundational content in the area of PeCS could be "engineered" into introductory courses, especially in mobile application programming, sensor programming, networking, and system design.

Recommendations to NSF:

The attendees suggested the following:

- Foster collaborative relationships between multiple mature theoretical areas that provide the basis for PeCS theory (e.g., HCI, networking, and machine learning). This may require programs that cross NSF divisional structure (CISE, Engineering, Social Sciences, etc.). It was noted that the Current NSF divisional structure encourages "depth-first" thinking, and that the cross-cutting programs may not cut deep enough to expose fundamental research issues.

- More emphasis should be put on research focused on theoretical foundation and collaborations between foundational projects and systems/experimental projects.

- The use and dissemination of tools built on sound theoretical foundations should be encouraged.

It was also noted that PeCS broad nature calls for cross-agency collaborations, especially with legal, financial and legislative branches.

Breakout Session Report:
Clouds and Crowds

Klara Nahrstedt, Jim Kurose, and Deepak Ganesan

The Need for Metrics

A key component of Mark Weiser's vision was to minimize the human distraction of technology with ubiquitous computing. However, there are no clear metrics to measure distraction. The group discussed the challenges in identifying such metrics. New metrics need to be considered. Unless we can measure distraction, we are not going to improve the pervasive systems. A number of approaches could be considered – throughput vs utility functions vs MOS for user experience. Possible metrics could consider 'easy of use' measurements, but might need some utility metrics for 'distraction' and 'invisibility' as mentioned in Weiser's vision. It is also not clear what distraction is – it has a personal aspect that may be age specific.

Distraction could be measured with respect to tasks. For example: when children multi-task, what do they give up? In case of fighter pilots, multi-tasking is an illusion since pilots do certain tasks in certain order (multiplexing between different tasks). Degradation of tasks should be quantified. We need to conduct user studies, observe users and measure difficulty of tasks. Another concept that needs to be measured is attention to tasks. Measurement of tasks in the field – consider a phone that asks if a user pays attention to the give task. We need to identify protocols that can evaluate tasks and hence how much attention a user pays to specific tasks.

As a result there is a need to study and identify a common framework for measurements of pervasive computing at scale (PeCS). Such a framework could explore machine learning and take data-driven approaches and benchmarks.

Benchmarking

Many other areas have benchmarks (high performance computing, databases, economics). We discussed the need for benchmarking which raised a number of challenges. What are the benchmarks to compare pervasive systems? How do we show that one system is better integrated than the other? What benchmark numbers do we need to show? (e.g., in computer architecture a benchmark is the performance of sorting algorithms on certain datasets). Is usability testing the response to the benchmarking?

Common Data Sets

The breakout also discussed the need for empirical datasets to help evaluate pervasive systems possibly gained from individual test-beds. It is important for the community to agree upon metadata for experimental datasets.

Economic Models

We discussed the issue of economics and business models for pervasive computing at scale. What are the economic models for deploying devices at scale? Our research usually captures cost such as delay, message overhead, but we need to factor in the management cost to the research agenda.

Cloud vs Crowd

There was a discussion of the architectural tradeoff of the cloud and crowd. Metrics in cloud are better understood that crowd, for example, Google cloud measures the number of clicks to understand the demand of websites. The performance of cloud computing matters because the latency of the cloud and its response to mobile devices is critical.

A major issue for the cloud in pervasive computing is locality of the cloud with respect to devices and the crowd. It could be that cloud is placed one hop away from devices via WiFi, or further away. One hop clouds represent "cloud-lets" around the network – a form of distributed clouds. Cloudlets could be efficient, supporting high-bandwidth communication because of their locality, and secure, leaving no state behind after their use. If cloudlets are close to devices and use generally available infrastructure (e.g., home computer, car computer, transition PC) then the economics of the cloudlet could be quite different from commercial cloud services today. There is a need to fully understand the role of clouds, their placement with respect to mobile phones. An example of smart rooms and smart offices was discussed. These are typically costly manage and require expertise significant expertise.. The development of cloudlets should attempt to avoid these problems, and be cost-effective.

We discussed if the concept of crowd/could represent a cloud storing data for neighboring nodes. One can imagine a crowd as the mobile cloud where mobile devices are treated as first class entities. Considering crowd as a mobile cloud could be useful to solve the latency problem, since one could get information quickly from/to mobile devices also accessing the cloud in aggregate.

A number of issues were raised during the discussion:

- The crowd could sense who is in the neighborhood. Are we heading towards SETI@mobile? The problem of using the crowd as a data cloud could impact energy mobile consumption. There are clear tradeoffs between computation and communication when considering the crowd as a cloud. If one does computing on the cloud (data warehouse) versus on the cloud (as a set of phones), then the energy usage is much more efficient.

- As we increase the scale of phones, and processing gets more distributed, there might be advantages to use phones as temporary clouds. The placement of cloud with respect to mobile phones is guided by economics. It is not clear what the business model is to have the cloud close or remote to mobile phones. Is this a question for NSF/academic community or industry?
- Measuring the crowd's means of organizing crowds and understanding what metrics matter. The cloudlets, crowds as mobile clouds, or local phones themselves could help balance the dynamics between the peak and average case of computing on the cloud (as a data warehouse). Phones are getting computationally powerful with multi-core architectures. Using this type of load offloading represents an agile way to move between different modes of computation depending on the situation. The whole load balancing between local phones, neighboring crowds or cloud-lets and remote clouds will succeed when the user does not know or manage how this dynamic behaviour is happening.

Crowd-sourcing

Crowd-sourcing is an important concept in pervasive computing, and could be used to identify metrics. A number of issues were discussed. Can crowd-sourcing be used for user interface design? It might be helpful to get feedback from the crowd to determine whether design of the application UI is good or bad. Crowd-sourcing can be also used to identify traffic conditiond, and local hotspots.

Test-beds

We need test-beds to experimentally validate the various algorithms, protocols, metrics, as well as collect datasets and benchmark and compare various pervasive systems. It might be difficult to setup a Planetlab-like uniform test-bed, rather there could be different forms of "testbeds" and sharing. Several forms of "test-beds" and sharing were discussed:

High-performance /supercomputing centers were mentioned as an example. It might be useful to establish a high-quality pervasive computing hardware environment, and then allow various researchers to conduct experiments. These specialized test-beds could have engineering staff to maintain the test-beds and work on tools to allow scientists to experiment with different pervasive hardware and software. This might be appropriate for environmental sensing, camera surveillance, etc. The advantage of this form of test-bed is that the engineers can worry about heterogeneity of devices, fast technological advances, deployment of devices and adjustment to the physical environment, etc.

Another form of test-bed was discussed for mobile sensors and people. Researchers, who have test-beds in their institution, could develop software and make it available to the community. Open-source software could be installed, and instantiated at other campuses. Hardware availability might be the barrier to this type of test-bed, butthis approach has proved to be successful in other disciplines.

Distribution of software, devices and experiences can be successful where the community contributes new software and it can also lead to standardization of devices and software. For example, XScale was successful – there was a source website and then people contributed to the source. Other examples were discussed from other disciplines. For example, a network of nano-technology was established at one university where researchers developed a software platform. The researchers publicized the software platform widely, and other institutions took the software platform for their nano-technology hardware platform.

Another form of test-bed could be based on a "collaborative exchange" of experiments. In this case, if a researcher runs an experiment on a testbed at his/her campus, he/she can run some additional experiments for another researcher. Then theseresearchers could run other experiments on the collaborationg campus collecting additional data, hence enhancing the scale of the experiment. This approach might fit nicely with an experimental test-bed of people. Humans are important part of pervasive computing

Another form of test-bed is an educational test-bed in our universities. Issues mainly concern how to recruit people as part of the test-bed. IRB issues emerge in these types of human test-beds. We should learn from social scientists, and other sciences that do experiments with human subjects..

Education Opportunities

It will be important to teach students to work with pervasive computing hardware and software. Establishing a common platform in universities would make sense. We discussed the idea that if one capstone class develops software and applications, this could be distributed and used by other capstone classes. One would need a common curriculum and robust software for teaching.

NSF Recommendations

We discussed the role of NSF and our recommendations are as follows:

- It is important that the academic community and NSF explore/fund disruptive technologies, what trends are coming, how technology scales, and changing use paradigms.
- Assist in establishing test-beds, but these test-beds might look different for pervasive computing at scale:
    - NSF should fund research test-beds, but test-beds that help us validate what cannot be done today.
    - How large scale should the test-beds be? What scales do we care about?

    - Examples of test-beds were discussed, such as the oceanography test-bed putting sensors on the floor of the ocean. This kind of testbed should always be extensible.
    - NSF should also partner with industry to assist in building testbeds (e.g., Intel)
    - For test-beds to be successful, there is a need for support from engineers to maintain the core system.
- There is a need to establish a national repository for data traces. NSF could help support this important task, ensuring the data is made available to the research community.

Appendix D - Plenary Session Report

Facilitators: Mahadev Satyanarayanan and Roy Want (Reported by Diane Cook)

Theme: Twenty Years after Mark Weiser's Vision for Ubiquitous Computing – What's Next?

In this discussion Mark Weiser's vision of the "disappearing computer" is revisited in this discussion. Reviewing the history of the field over the last twenty years, we see that smaller mobile devices of all sizes and types are becoming increasingly prevalent. The field has certainly changed over the last twenty years. The following questions were thus posed to workshop attendees:

1. Why is the vision so powerful?
2. How might it be derailed?
3. What things were not in the vision?

Why is the vision so powerful?

The vision is appealing in theory and has had an impact on the population in practice. One possible reason why the vision is so powerful is that is reaches much larger masses than technology has in the past. As was stated in the movie The Social Network, developing countries like Bosnia lack roads but "they have Facebook". Mobile phones are accessible and reach around the world. Cost of the devices steadily decreases while access to the devices and software and diversity of the applications steadily increases.

Another stated reason is that the technology appeals to the fundamental human trait of being lazy, in much the same way as Facebook appeals to a fundamental need to be social. There was some disagreement of whether technology truly appeals to a sense of laziness or whether pervasive computing has simply lowered the barrier of entry.

Current society exhibits an information-heavy lifestyle. When technology disappears then it also has the effect of simplifying our lives. Pervasive computing is not only technology that is easier to access, but it is also more mobile technology. A computer can be frustrating because it is stuck in one place and users need to access it in one physical location and posture. As a result, use of technology in that context is limited to the same position and posture and we tend to change our lives to use and accommodate these restrictions. The vision frees us from some of these constraints – now users can share information in a device independent, place independent fashion.

Some argue that reality has moved in the opposite direction of the vision and become more difficult, not less. However, this may be an issue of "fractal complexity". The technology at first appears hard then is simplified for ease of use. A user looks closely at how it works and it once again appears hard.

Finally, the question was raised of whether the vision claimed that technology will literally disappear, or whether it will become so natural that individuals do not think about it. For example, a book read on a kindle is such a natural activity that while the technology has obviously not disappeared, the user does not view the technology any differently than viewing a written book. In fact, perhaps one could say that the functionality of computational devices is growing while the devices themselves get simpler so the vision is not ever achieved; it just stays on a constant level of appearance.

Things that might derail the vision.

While the original goal for the field was quiet functionality, the reality of pervasive computing and associated costs (such as pop-up advertisements) are prohibitively noisy. Is the example of information and advertisement bombardment found in the movie "Minority Report" an inescapable future for the field?

One response to this concern is to push the field and its commercialization practices towards click-based ads rather than forced impressions. One can also view this issue as a decision theory problem. Developers need to make sure that the benefits of forced information / advertisements outweigh the problems. We can view a mobile device as an auctioneer which is attempting to buy a user's attention. Google is an example of a company that is already adopting this click-based approach to directing a user's attention toward advertisements.

Another issue that has arisen is the failure to resolve the many conflicting standards that have emerged. While the vision of pervasive computing seems to have become a reality, there are also now islands of connectivity. In practice most devices are not able to communicate with each other because of this lack of compatibility.

The last issue that was discussed in this context was the threats that viruses, malware, and even legitimate uses of the technology pose to privacy and security. An individual's context-aware data, which is captured in order to meet their daily needs, is also being sold all over the Internet. Advertisers can thus "read your mind", a state of technology which is eerie, annoying and unethical.

Workshop attendees agree that people do not like the idea of "big brother" watching them, but recognize at the same time this is necessary for the technology to be effective. However, the question was raised of whether this dislike of information sharing is a generational issue.

The point was also raised that users do not necessarily understand the technology and how it works. Most agree that technology designers and developers need to make the technology transparent enough so that users understand the true privacy risks. There was also a concern expressed that we need to educate researchers, students, companies, and the general public about these possible concerns, about how technology works, about ourselves, and to try to understand why we like the technology so much.

Things that were not prophesied

Although aspects of the original vision have become reality, there are certainly technologies, applications, and impacts emerging that were not foreseen. Examples of these include the world-wide web, crowd sourcing, and social networking. The question was raised of how the vision should be modulated to take these capabilities into account.

The point was raised during this discussion that the vision of Charles Babbage does not align with the pervasive computing vision. In pervasive computing we need to not just consider a single user and device, but also need to plan for and reason about communities, social networks, and networks of systems.

There was a lengthy and controversial discussion of the immediacy of information through pervasive computing and the impact this has on humans. Some in the group argue that pervasive computing has encouraged shorter attention spans and shallower thinking. An analogy was made to the way in which humans transitioned from collecting information on printed paper to getting it from radio and television sets. While the intent was to free our time to do more interesting things, one could argue that technology has done the opposite. Because we are bombarded with information we do not spend time thinking through the information deeply.

While many agreed with the fundamental issue that was raised, there was some discussion of whether this change was necessarily good or necessarily bad, or just change. As was pointed out, the next generation may view spending hours reading a book as a bad use of time. Spending large amount of time collecting and sifting through information may eventually become as archaic a way of life as riding horses and sword fighting is today. As one person pointed out, deep thinking about one paper and point of view may not be needed because we can now quickly look at more than one thought, one argument, one paper, and can thus quickly generate a more comprehensive view.

Workshop attendees agreed that methods of learning have changed. Students do not learn primarily from a teacher anymore, nor primarily from a book. Instead they are learning from collaboration. We need to adapt to this not-prophesied change by adjust our methods of educating and communicating with students. There was also a question raised of what the difference is between individual thought and distributed collaboration.

The discussion raised a number of points regarding the history of the field, the future of the field, and the potential dangers. The discussion concluded with the point that technology may not disappear because a generation gets used to a current look and feel. There is no question, though, that the vision has had an impact on the research community and on the world's population. The group felt that these issues warranted further reflection as we shape new goals for the field.

Appendix E – Workshop Survey Results

The complete set of survey responses is available at http://sensorlab.cs.dartmouth.edu/NSFPervasiveComputingAtScale/survey.html.