

# A Note On Pervasive Computing

Nishkam Ravi and Liviu Iftode

Department of Computer Science, Rutgers University  
{nravi@cs.rutgers.edu, iftode@cs.rutgers.edu}

Pervasive computing is an application-driven field and as such is ill-defined and unstructured for the lack of definitions, algorithms and architectures. Research is primarily guided by visions and beliefs because it is hard to tell, in a scientific way, which systems/applications will penetrate the market. To realize the diverse set of applications, research has to be multi-disciplinary, which adds to the complexity. The exploratory and unstructured nature of pervasive computing research is, therefore, justified to some extent. Besides, this field is still in a nascent stage. Unlike networks/compilers/operating systems, which exist and are widely used and deployed, pervasive computing exists mostly in labs or in specialized environments or in the form of single isolated applications. As a result, progress is hard to measure and prove. In this position paper, we express our view on what should be the sub-areas of focus, what are the key research challenges and how progress can possibly be evaluated.

As suggested by Mark Weiser [17], pervasive computing world would consist of three sets of entities: user agents (i.e devices carried by users), devices embedded in the environment and the internet. The digital interaction among and between these entities would enable new applications. Instead of trying to realize this vision in one shot, we need to take a prudent iterative approach, and identify the platforms that we can build on top of *today*.

## Early platforms for pervasive computing

Most people spend a significant fraction of their time either at home or at the workplace, where they have access to PCs (especially in developed countries). Most of the remaining time is spent commuting between places/locations either on foot or in the car or using public transport. If we assume that the primary purpose of pervasive computing is to provide ubiquitous access to computing, working with: (a). car PCs, and (b). smart mobile phones may take us reasonably close to this goal *in the near future*.

Mobile phones are steadily evolving and apart from the hardware advances, the ever increasing set of C++ and Java APIs, as available from Symbian and Nokia development forums, is an indication of the market support. What is more important is the fact that certain operating systems (such as Symbian [4]) and development profiles (such as CLDC and MIDP [2]) have become a standard and are being *worked on top of*, unlike most previous pervasive computing initiatives (such as middleware and service discovery protocols), which were never *reused*.

At the same time, several car companies including GM, Toyota and Chrysler are taking serious initiatives to standard-

ize software and bootstrap vehicular computing. Researchers from academia and industry are joining forces to develop real deployable vehicular computing systems [9]. Dedicated Short Range Communications (DSRC) [1] is the first result of this effort. Vehicular computing should, by all means, be considered a core branch of pervasive computing and entertained at top pervasive computing conferences.

To summarize, assuming that people have access to PCs at home and at workplace, putting a PC in every car and a smart phone in every pocket and building on top of this simple infrastructure may take us closer to the pervasive computing goal in a structured fashion. The specific application domains of pervasive computing (such as healthcare or ISR [7]), which are somewhat established and have their own set of issues, should be researched on independently and can also benefit from generic pervasive computing research.

## Application classification

We classify pervasive computing applications in three categories based on interaction modes: (1) user agent to user agent (2) user agent to internet, and (3) user agent to smart environment; where the user agent could either be a smart phone or a car PC. Both (1) and (2) are easy to bootstrap as they do not require any extra infrastructure, while (3) depends on the presence of smart "tagged" devices in the environment. We argue that (3) should be rethought. While RFID tags [3] have taken us closer to (3), it will take a long time before we have an RFID-tag-reader-rich environment. RFID is a promising technology and should be researched on, but alternate mechanisms should be thought of simultaneously.

We suggest that instead of recognizing and interacting with devices through tags or wireless interfaces (which require the environment to be instrumented with receivers/transmitters, thus increasing deployment cost and effort), interaction should be based on information that can be sensed directly, such as video and audio. Embedded microphones and cameras (as present on mobile phones) along with image and audio processing algorithms can go a long way in enabling this mode of interaction [14, 13, 8]. Once this mode of interaction is established and creates a market for applications, it will be easier to invest money into tagging the environment with RFID tags or visual tags [16] on a global scale.

## Primary research challenges

Common to (1), (2) and (3) are four research challenges: (i) minimization of human involvement and simplification of human-computer interaction, (ii) judicious use of limited battery lifetime on mobile devices, (iii) privacy, and (iv) sponta-

neous interaction between devices through wireless interfaces or images/audio.

Use of context-information [15] (such as location, activity, calendar and personal data) can help reduce user involvement. This would require research at different levels of context-awareness, including context sensing, processing, storage and inference. Speech-recognition interfaces and device-display symbiosis [10] can be instrumental in simplifying human-computer interaction and should be researched on more.

Unless ambient energy sources are discovered and exploited, battery lifetime on mobile devices will continue to be a nagging concern. It is important to build mechanisms that optimize energy consumption of applications[6]. At the same time, there is a need to come up with *battery management* techniques [12, 18] that can help in the judicious use of battery lifetime across different applications.

Sharing of personal data gives rise to privacy concerns. Although, this problem does not manifest itself until pervasive computing applications are used and deployed, taking a cue from the database community, it is important to take a pro-active approach and solve this problem before it even arises. Information-flow control models [11] and usable-privacy mechanisms [5] can help alleviate this problem.

Wireless networking protocols (e.g 802.11, Bluetooth and 3G) have and will continue to play an integral role in the development of pervasive computing. A lot more work is needed before wireless networking can enable configuration-free interaction between alien devices. It would be wise to explore image/audio based interaction between devices simultaneously. Independent research on smart and cheap ways of *tagging* devices would be valuable. RFID and visual tags are a good start.

### Research methodology

There are typically two approaches to research— a horizontal layered approach, and a vertical exploratory approach. In the early days of a research area, vertical exploratory approach is better. As the field matures and some solutions get more accepted than the others, it is important to develop on top of them in a layered fashion. In the case of pervasive computing, both have to be done simultaneously. The reason is that while pervasive computing is 15 research-years old, in the real world, it is not even born yet. It is, therefore, important for top pervasive computing conferences to welcome both kinds of work. In other words, it is important to advance the state-of-the-art by encouraging novel controversial ideas that have not been thoroughly tested due to lack of proper evaluation methodologies. At the same time, in order to understand what can be deployed today, it is important to encourage not-so-novel but usable systems/applications that build upon previous work and have been thoroughly tested.

A suggestion would be for top conferences to have two separate tracks, one promoting exploratory research and novel ideas substantiated through early experiments and logical arguments, and the other promoting matured experience papers substantiated through user studies; and these two categories of papers should not be allowed to compete with each other during the review process. The final measure of any pervasive computing system/application is what value it gives to the user (in terms of functionality and usability) in return for the

price that the user has to pay to avail of it (in terms of resources such as user attention, battery lifetime and personal data). This can only be measured through user studies.

Finally, given the lack of good economic models for pervasive computing applications, it is important to understand what users *really* want in contrast to what researchers think users want. User studies can be instrumental in achieving this goal.

### Conclusions

In this position paper, we have argued that: (1). car PCs and mobile phones can serve as early platforms for pervasive computing and should be built on top of, (2). pervasive computing applications can be classified in three categories based on interaction models: user-agent to user-agent, user-agent to environment and user-agent to internet. (3). in addition to trying and instrumenting the environment with tags and transmitters, we should explore *passive* ways of interacting with the environment i.e through audio and visual signals, as humans do (3). the four key research challenges facing us today are: human-computer interaction, limited battery lifetime on mobile devices, data privacy and spontaneous interaction between devices (via tags, wireless interfaces or video/audio) (4). we should promote both exploratory research and thoroughly tested matured research and let them not compete with each other (5). user studies are very important for understanding what users want and are the final measure of progress for pervasive computing.

### References

- [1] Dsrc. <http://grouper.ieee.org/groups/scc32/dsrc/index.html>.
- [2] Midp. <http://java.sun.com/products/midp>.
- [3] Rfid. <http://en.wikipedia.org/wiki/RFID>.
- [4] Symbian os. <http://www.symbian.com>.
- [5] J. C. et al. User-controllable security and privacy for pervasive computing. In *WMCSA '07: Proceedings of the 8th IEEE Workshop on Mobile Computing Systems & Applications*, 2007.
- [6] J. Flinn and M. Satyanarayanan. Energy-aware adaptation for mobile applications. In *SOSP '99: Proceedings of the seventeenth ACM symposium on Operating systems principles*, 1999.
- [7] M. Kozuch and M. Satyanarayanan. Internet suspend/resume. In *WMCSA '02: Proceedings of the Fourth IEEE Workshop on Mobile Computing Systems and Applications*, 2002.
- [8] C. V. Lopes, A. Haghighat, A. Mandal, T. Givargis, and P. Baldi. Localization of off-the-shelf mobile devices using audible sound: architectures, protocols and performance assessment. *SIGMOBILE Mob. Comput. Commun. Rev.*, 10(2), 2006.
- [9] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. Trafficview: traffic data dissemination using car-to-car communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 8(3):6–19, 2004.
- [10] M. Raghunath, C. Narayanaswami, and C. Pinhanetz. Fostering a symbiotic handheld environment. *IEEE Computer*, 36(9):56–65, 2003.
- [11] N. Ravi, M. Gruteser, and L. Iftode. Non-inference: An information flow control model for location-based services. In *Proceedings of the 3rd International Conference on Mobile and Ubiquitous Systems (MobiQuitous)*, 2006.
- [12] N. Ravi, J. Scott, and L. Iftode. Context-aware battery management for mobile phones: A feasibility study. Technical Report DCS-TR-608, Rutgers University, December 2006.

- [13] N. Ravi, P. Shankar, A. Frankel, A. Elgammal, and L. Iftode. Indoor localization using camera phones. In *WMCSA '06: Proceedings of the Seventh IEEE Workshop on Mobile Computing Systems & Applications*, 2006.
- [14] M. Rohs and P. Zweifel. A conceptual framework for camera phone-based interaction techniques. In *Proceedings of the Third International Conference on Pervasive Computing*, 2005.
- [15] D. Salber, A. K. Dey, and G. D. Abowd. The context toolkit: aiding the development of context-enabled applications. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, 1999.
- [16] D. Scott, R. Sharp, A. Madhavapeddy, and E. Upton. Using visual tags to bypass bluetooth device discovery. *SIGMOBILE Mob. Comput. Commun. Rev.*, 9(1), 2005.
- [17] M. Weiser. The computer for the 21st century. *Scientific American*, 265(3):66–75, September 1991.
- [18] H. Zeng, C. S. Ellis, A. R. Lebeck, and A. Vahdat. Ecosystem: managing energy as a first class operating system resource. In *ASPLOS-X: Proceedings of the 10th international conference on Architectural support for programming languages and operating systems*, 2002.