IEEE pervasive COMPUTING
MOBILE AND UBIQUITOUS SYSTEMS
www.computer.org/pervasive

The 2005 UbiApp Workshop: What Makes Good Application-Led Research?

Richard Sharp and Kasim Rehman

Vol. 4, No. 3
July–Sept. 2005

Editor: Mary Baker ■ HP Labs ■ mary.baker@hp.com

The 2005 UbiApp Workshop: What Makes Good Application-Led Research?

Richard Sharp and Kasim Rehman

Application-led research programs involve the design, implementation, deployment, and evaluation of specific applications. In ubiquitous computing, such projects are often interdisciplinary efforts involving (among others) social scientists, ethnographers, technologists, and human-computer interaction specialists. This research typically aims to

- help identify users' problems and application requirements,
- provide infrastructure developers with application requirements, and
- validate technology and provide insights into its use.

At Pervasive 2005, the UbiApp Workshop on 11 May addressed the question, what makes for good application-led research in ubiquitous computing? The workshop was organized in response to the concern of some members of the ubiquitous computing community that application-led research needs to make more coherent progress. Their perception was that with few exceptions, such research is neither systematically building on what little new knowledge it has derived so far nor setting specific challenges and benchmarks to guide its progress. In light of these concerns, the workshop's tenet was that ubiquitous computing research could benefit from better metrics for the selection, analysis, and evaluation of applications and common infrastructure. The workshop aimed to identify methodological problems in the way researchers conduct application-led research and to recommend how to address these problems.

The 25 attendees represented a variety of positions across academia and industry. Here, we summarize their discussions and present the outcomes in cases where they reached consensus. More information, including the full text of the attendees' position papers, is at www.cl.cam.ac.uk/ubiappwsweb.

WHAT APPLICATION-LED RESEARCH IS

Workshop participants observed that ubiquitous computing research falls into roughly two categories:

- *Application-led research* is driven by a domain problem and evaluated by deploying a solution and quantifying the benefits brought to the chosen domain.
- *Technology-led research* isn't necessarily motivated by the benefits brought to a particular application domain but is carried out because it's interesting or challenging from a purely technical perspective.

The attendees' consensus was that for the ubiquitous computing community to be effective, it must engage in a combination of technology-led and application-led research. Technology-led research can open up new possibilities for applications, and, conversely, new applications can generate requirements for ubiquitous computing technology.

Some attendees observed that many ubiquitous computing research projects design and build new applications as "proof of concept demonstrators." Such applications aren't motivated by current problems but are built to demonstrate some new technology's particular features. The utility of application development as "proof of concept" was called into question. In his position paper, Nigel Davies (Lancaster University) summarized the views of many attendees:

> *The problem, very often, is that there is no actual concept to be proven. Either the concept has already been proven viable (there really is no need to prove again that we can build a context-aware tour guide), is never in any doubt (we know we can build location-based services) or is not actually proved by the demonstrator (proof is a very strong term!).*[1]

While application development as proof of concept might be a useful tool in technology-led projects, attendees generally agreed that application demonstrators don't constitute application-led research. In contrast to demonstrators, application-led research should be motivated by current problems and evaluated according the benefit it has given to users.

Published by the IEEE CS and IEEE ComSoc ■ 1536-1268/05/$20.00 © 2005 IEEE

HOW TO SELECT APPLICATIONS

For application-led research to be successful, the community must have suitable methods for selecting applications to focus on. This topic received a great deal of discussion at the workshop.

Several participants stressed that application-led ubiquitous computing researchers should make a greater effort to team up with experts in their chosen domains when selecting problems to work on. Richard Beckwith described how he and other Intel researchers applied this approach successfully to a project to augment vineyards with climate-monitoring sensors. They first performed an ethnographic field study to help them understand the problem space, talking to vineyard owners, vineyard managers, wine makers and their assistants, wine marketers, wholesalers, and retailers. Then, once they had chosen applications on the basis of these interviews, they worked with expert plant physiologists to implement the applications. The resulting system could predict variation in pH, titratable acids, berry weight, and boundaries for frost damage and could define areas that would be amenable to growing more valuable crops. Beckwith commented that "we could have done none of this without the input of a scientist working in the field."

Attendees generally felt that too many ubiquitous computing projects focus on applications addressing trivial problems (turning lights on and off remotely, finding others with similar interests at conferences, and so on). In contrast, independent consultant William Newman argued that application-led researchers should be aiming to address "severe and persistent problems." He pointed out that, compared to much of the world's population, ubiquitous computing researchers often enjoy "relatively problem-free lives." So, we should be keen to look beyond our own experiences when choosing application domains. For example, what opportunities exist to address problems in war zones or refugee camps?

HOW TO IMPLEMENT APPLICATIONS

Many attendees observed that ubiquitous computing researchers often reimplement applications from scratch, rather than sharing code and building on each other's work. Developing large applications can require a huge amount of time and resources. (For example, Georgia Tech's Gregory Abowd told of the many person-years of work that went into the design and implementation of the eClass project, which studied ubiquitous computing's effect on education.) The current lack of cooperation within the community is thus limiting the size and scope of applications that the community can feasibly implement. The consensus was that this problem was serious and must be addressed.

Application-led ubiquitous computing researchers should team up with experts in their chosen domains when selecting problems to work on.

Many attendees hypothesized that a reason for the lack of reusable infrastructure between projects is that the research community tends to value novelty over good engineering practice. So, researchers have few incentives to spend time building modular, maintainable systems; instead, they can attain greater credit by publishing several novel ideas, each accompanied only by proof-of-concept implementation. Future conference program committees should aim to address this issue, striving to reward research that constructs open, reusable infrastructure for the wider community's benefit.

Some participants took this argument further and claimed that researchers currently have few incentives to engage in application-led research at all. Not only does application development take a great deal of time, but also the ubiquitous computing community doesn't regard much of this implementation work as research. Attendees agreed that the community must focus on ensuring that researchers are appropriately rewarded for their implementation work.

One project that has bucked the trend is Place Lab, which provides a reusable software platform for location-aware computing. Gaetano Borriello (Univ. of Washington) explained that Place Lab's aim has always been "to make it easier for others to build location-aware applications." This statement has been justified in that independent groups of researchers have built applications on top of Place Lab components. Some attendees suggested that the community should build more applications like Place Lab, creating the technical infrastructure required to advance application-led research.

HOW TO EVALUATE APPLICATIONS

Researchers commonly evaluate ubiquitous computing applications solely in the context of small lab-scale user studies. Several workshop participants argued strongly against this approach. For example, Nigel Davies emphasized in his paper that lab studies aren't a substitute for deployment, commenting that "it is impossible to understand ahead of time the impact of the environment on technology (or indeed, the impact of technology on the environment), and this is often critical to system design."[1] In the panel discussion, Abi Sellen (Microsoft Research, Cambridge) took this point further, arguing that the only way to evaluate an application against the ideals of ubiquitous computing ("everywhere, connected, context aware, invisible") is through "long-term deployment in the wild."

Of course, small-scale lab studies still have a place—everyone agreed that they're very useful in the early stages of user-centered design. The arguments for evaluation through real-world deployment simply stated that once researchers

CONFERENCES

have performed lab-scale trials, they shouldn't be content to stop there. Instead, application-led research programs should use this data to continue to design, deploy, and evaluate similar applications on a larger scale.

Several attendees commented that applications are often evaluated only against themselves (for example, "our participants said that they found this application useful"). They claimed that self-evaluation alone isn't sufficient; that for the community as a whole to make progress, data must be available comparing separate applications developed by independent research groups. Only if these comparative results are available will the community ever be able to draw general conclusions from the plethora of developed applications. Of course, this requires that research groups make their applications, infrastructure, and data traces available publicly—a characteristic that attendees felt the ubiquitous computing community also lacked.

RECOMMENDATIONS

The workshop produced four general recommendations for community action to increase application-led research's effectiveness:

- *Choose applications carefully*. Rather than settling for trivial, low-value problems, aim to build applications that address severe, persistent problems. When conceiving application-led projects, team up with experts in your chosen application domains, during both design and implementation.
- *Share technical infrastructure*. Design applications so that the software and infrastructure you've developed can be reused to the wider community's benefit. Publicly release the code for applications and their associated infrastructure.
- *Evaluate applications in realistic environments*. Real-world deployment is the only way to investigate fully the complex three-way interactions between ubiquitous computing applications, their users, and the environment. Such evaluations are costly and have substantial implications for organization of the research. There's a need to explain this to researchers, funding sources, and the wider community.
- *Perform comparative evaluations*. Publicly release data sets derived from field trials so that other researchers can verify the published results and analyze the data. Design studies should not only validate a single application but also explicitly compare similar applications.

The attendees agreed that the entire ubiquitous computing community must work together to ensure that application-led research reaches its potential. The responsibility for implementing the recommendations lies not only with individual researchers but also with everyone involved in the research process, including conference program committees, educators, and industrial collaborators.

ACKNOWLEDGMENTS

We thank the UbiApp Workshop organizing committee: George Coulouris, Tim Kindberg, Bernt Schiele, Albrecht Schmidt, and Kasim Rehman.

REFERENCE

1. N. Davies, "Proof-of-Concept Demonstrators and Other Evils of Application-Led Research: A Position Statement," 2005; www.cl.cam.ac.uk/ubiappwsweb.

Richard Sharp is a senior researcher at Intel Research Cambridge. Contact him at richard.sharp@intel.com.

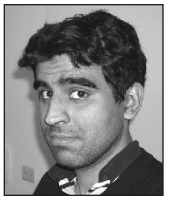

Kasim Rehman is a postdoctoral research associate at the University of Cambridge's Computer Laboratory. Contact him at kr241@cam.ac.uk.