Deploy or Die: A Choice for Application-Led Ubiquitous Computing Research

Richard Sharp

[1]Intel Research Cambridge,
15 JJ Thomson Avenue, Cambridge, CB3 0FD, UK.
richard.sharp@intel.com

1 Introduction

For the past fifteen years Ubiquitous Computing researchers have been exploring how computing can be pushed beyond the traditional desktop environment and seamlessly *"woven into the fabric of our everyday lives"* [13]. Yet, despite such grand vision, the vast majority of UbiComp research is published and forgotten long before it has any impact on everyday life, let alone woven into its fabric. I believe that the UbiComp community urgently needs to address this issue before the disparity between the now cliched rhetoric and the lack of real world impact drives the field into disrepute.

In this position paper I argue that recent technological developments make it possible for Ubiquitous Computing applications to be deployed on a global scale and that, with a few exceptions, the research community is currently failing to embrace this opportunity for real-world impact (Section 2). I propose a direction for application-led UbiComp research that makes widespread deployment its success criteria (Section 3) and discuss the attitudes and practices within the community which are currently hampering the pursuit of deployable applications (Section 4). Finally, I highlight one pioneering research project that has already demonstrated that building deployable UbiComp applications can generate both global impact and research contributions (Section 5).

2 The Computer for the 21st Century is Here. Where is UbiComp?

In 1991, Weiser stated that the hardware required for ubiquitous computing would come in two parts: (*i*) cheap, low-power computers that include equally convenient displays; and (*ii*) a network that ties them all together [13]. These two requirements have arguably been satisfied for the past five years: programmable cell-phones, PDAs and laptops have become commonplace and a plethora of wireless networking technologies (e.g. Bluetooth, 802.11, 3G) are now standard.

So, given that the necessary infrastructure now exists, where are the context-aware applications [10], augmented homes [5], smart offices [12], geo-annotation systems [4], electronic tourist guides [3] (and dangling pieces of string [14]) that UbiComp promised? By contrasting the claims of UbiComp papers with our everyday experiences it is clear that the former are detached from the latter.

The UbiComp community has survived by continually adapting its research program (application domains chosen in an ad-hoc manner) in order to save its hard core assumption (that computers will become invisible, automatically inferring and catering for our every need). To remain credible, at least some of the technologies, applications, user-interfaces and usage models that the field has predicted must soon be seen in the real world. Like AI before it, UbiComp risks being categorized as what Imre Lakatos terms a "*degenerating research program*" [7].

3 A Deployment-Driven Methodology for UbiComp Research

Given this urgent need for real-world impact, I propose that application-led UbiComp research projects adopt the following deployment-driven methodology: (1) build a well- engineered, robust UbiComp ap-

plication that leverages existing infrastructure (WiFi, PDAs, Laptops, Cell-phones etc.); (2) release this application publicly; (3) build a real- world user-base around the application; and (4) study this user-base and learn from users? experiences.

The vast majority of UbiComp projects that attempt to follow this methodology will inevitably fail to progress to stage 3; many uncontrollable factors contribute to whether a real-world user-base can be established. However, the hope is that, if the community as a whole makes a concerted effort to build deployable applications (and if UbiComp research is indeed relevant to people's everyday lives) then some will attract significant user-bases. These successful applications will be lucky enough to proceed to stages 3 and 4, generating both much-needed impact for the UbiComp community and valuable insights into how or why the applications were adopted by users (on a potentially global scale).

Projects that do not manage to build real-world user-bases should not be regarded as failures. By merely achieving stage 1 they will have encountered and solved interesting research challenges, e.g. how was the application designed and adapted to work on existing infrastructure? What engineering approaches were used to facilitate scalability?

This proposal is hardly revolutionary; indeed, many would argue that most technical computer science research already proceeds in this way. UbiComp, however, is certainly not embracing a deployment-driven methodology. Of the 26 projects presented in the proceedings of UbiComp04, not one describes a publicly released application that users can download and benefit from. A single project, Krumm's *NearMe Wireless Proximity Server* [6], achieves stage 1 but fails at stage 2: as I write this article I can find no way of downloading either his client- or server-software onto my laptop.

4 Attitudes Within UbiComp Hampering Pursuit of Deployment

Given that other Computer Science research disciplines have successfully generated impact by applying the above 4-stage methodology, why has UbiComp not followed suit? I believe that there are two common attitudes in the community that are hampering the development of deployable UbiComp applications. These are discussed below.

4.1 Repeating ParcTab: An Obsession with Building Custom Hardware

The philosophy of using custom hardware to support application-led UbiComp research dates back to the ParcTab Ubiquitous Computing Experiment where Want et. al. argued that it allowed them to "*glimpse into the future* [11]. At this time researchers had no alternative but to build custom hardware in order to explore the potential of Ubiquitous Computing; there was no existing infrastructure that could support UbiComp applications. The success of the pioneering ParcTab project greatly influenced the UbiComp community. In particular, its approach of deploying custom hardware to explore futuristic applications became generally accepted as the de-facto methodology for technical UbiComp research.

Today, when Weiser's requirements for Ubiquitous Computing infrastructure have been met, this approach must be challenged. Building custom hardware infrastructure still allows one to *glimpse the future*. However, it also eliminates all possibility of public release, widespread deployment and, therefore, real-world impact. In a time when these goals are attainable, researchers should think long and hard about the opportunity-cost of building custom infrastructure.

4.2 Hundreds of Small, Disparate Projects; No Community-Wide Efforts

As evidenced by the proceedings of Ubiquitous Computing conferences, the UbiComp community has typically focused on small, disparate projects. Although researchers have sometimes collaborated in lab-sized teams to engineer larger projects, there have been few (if any) community-wide efforts. If a deployment-driven methodology is to be adopted, greater collaboration is essential.

Designing and implementing robust deployable applications is expensive and time-consuming requiring, in many cases, more resources than a single research group can provide. Other CS research disciplines have successfully tackled this problem by using open source development techniques. For example, consider the Operating System community's BSD UNIX [9] and, on a smaller, but still impressive scale, the

Programming Language Community's *Haskell* language and its associated compilers and interpreters [1]. It is vital that the UbiComp community mirror this approach, publicly releasing source code and actively seeking to build on each other's implementations.

5 An Exemplary Deployment-Driven UbiComp Project

Although I have argued that the UbiComp community is largely failing to embrace the opportunities for real-world impact, there is one notable exception: PlaceLab [8].

PlaceLab is a successful UbiComp project that has followed a deployment-driven methodology. The project aims to enable commodity hardware clients like notebooks, PDAs and cell phones to locate themselves by listening for radio beacons that already exist in the environment (e.g. 802.11 access points, GSM cell phone towers, and fixed Bluetooth devices). A robust software implementation has been developed and made available for public download in a variety of formats including Windows XP, Linux, Mac OS X, Windows CE/Pocket PC and Nokia Series 60 Phones. As a result a significant global user-base has now been established.

The PlaceLab project did not invent the idea of using existing radio beacons for location: there is earlier literature on this topic [2]. The contribution of PlaceLab was to take this embryonic idea and produce a well-engineered platform capable of scaling to a global deployment. In this process the PlaceLab researchers encountered and solved problems that initial proponents of radio beacon location systems had not envisaged. Furthermore, they generated impact for the UbiComp community by developing a global, genuinely ubiquitous location infrastructure.

6 Conclusions

For many years UbiComp researchers have imagined an age when low-power computing devices, display technology and wireless networking capability would be truly ubiquitous. That time is now and, by exploiting this existing infrastructure, the community at last has a chance of achieving real-world impact. However, the current culture of lab prototypes and small, disparate research programs is preventing UbiComp from reaching its potential. In this paper I have presented a deployment-driven methodology which attempts to address this issue and have highlighted areas in which community-action is required if the approach is to be successful.

The next few years will determine the success or failure of Ubiquitous Computing research. The choice is simple: we must deploy (thus demonstrating our relevance to 21st century computing and silencing our critics) or die.

References

[1] `http://haskell.org/`.

[2] BAHL, P., AND PADMANABHAN, V. N. RADAR: An in-building RF-based user location and tracking system. In *Proceedings of INFOCOM* (2000), pp. 775–784.

[3] CHEVERST, K., DAVIES, N., MITCHELL, K., AND FRIDAY, A. Experiences of developing and deploying a context-aware tourist guide: the guide project. In *Proceedings of MobiCom '00* (2000), ACM Press, pp. 20–31.

[4] ESPINOZA, F., PERSSON, P., SANDIN, A., NYSTRM, H., CACCIATORE, E., AND BYLUND, M. Geonotes: Social and navigational aspects of location-based information systems. In *Proceedings of UbiComp '01* (2001), Springer-Verlag, pp. 2–17.

[5] KIDD, C. D., ORR, R., ABOWD, G. D., ATKESON, C. G., ESSA, I. A., MACINTYRE, B., MYNATT, E. D., STARNER, T., AND NEWSTETTER, W. The aware home: A living laboratory for ubiquitous computing research. In *CoBuild '99: Proceedings of the Second International Workshop*

on Cooperative Buildings, Integrating Information, Organization, and Architecture (1999), Springer-Verlag, pp. 191–198.

[6] KRUMM, J., AND HINCKLEY, K. The NearMe wireless proximity server. In *Proceedings of Ubicomp 2004*, vol. 3205 of *LNCS*, Springer-Verlag.

[7] LAKATOS, I. History of science and its rational reconstruction. In *PSA* (1971), R.C.Buck and R.S.Cohen, Eds., Reidel.

[8] LAMARCA, A., CHAWATHE, Y., CONSOLVO, S., HIGHTOWER, J., SMITH, I., SCOTT, J., SOHN, T., HOWARD, J., HUGHES, J., POTTER, F., TABERT, J., POWLEDGE, P., BORRIELLO, G., AND SCHILIT., B. Place lab: Device positioning using radio beacons in the wild. To appear, Pervasive 2005.

[9] MCKUSICK, M. A berkeley odyssey. *Unix Review 3*, 1 (January 1985).

[10] SCHILIT, B., ADAMS, N., AND WANT, R. Context-aware computing applications. In *IEEE Workshop on Mobile Computing Systems and Applications* (Santa Cruz, CA, US, 1994).

[11] WANT, R., SCHILIT, B. N., ADAMS, N. I., GOLD, R., PETERSEN, K., GOLDBERG, D., ELLIS, J. R., AND WEISER, M. The ParcTab ubiquitous computing experiment. Tech. Rep. CSL-95-1, Xerox Palo Alto Research Center, March 1995.

[12] WARD, A., JONES, A., AND HOPPER, A. A New Location Technique for the Active Office. *IEEE Personal Communications 4*, 5 (October 1997), 42–47.

[13] WEISER, M. The computer for the 21st century. *Scientific American 265*, 3 (January 1991), 66–75.

[14] WEISER, M., AND BROWN, J. Designing calm technology. *PowerGrid Journal 1*, 1.