Ongoing Support for Deployments in the Wild

Nick Taylor
Culture Lab, School of Computing Science
Newcastle University, UK
nick.taylor@newcastle.ac.uk

Keith Cheverst
School of Computing and Communications
Lancaster University, UK
k.cheverst@lancaster.ac.uk

ABSTRACT
As research increasingly turns to deployments 'in the wild' to evaluate technologies under real-world conditions, little consideration has been given to what happens when research ends. In many cases, users are encouraged to integrate these technologies into their lives before they are withdrawn, while in some cases technologies are being left in place after research concludes. In this position paper, we revisit a past deployment and explore events following the conclusion of the project. Using this example, we motivate the need for greater exploration of how handovers of research prototypes can be executed and supported.

Author Keywords
Community; research in the wild; longitudinal.

INTRODUCTION
Research around novel technologies is increasingly coming to recognize the value of deployments 'in the wild', where prototype systems can by trialled and evaluated over long periods of time through actual day-to-day use. This allows us to gain a much clearer understanding of how the technology can be integrated into people's everyday routines and the real-world issues that arise, as opposed to lab-based studies that take users out of their normal setting into a tightly controlled environment, typically for only short periods of time.

We have increasingly come to recognise that our interactions with technology are *situated* within particular "material and social circumstances" [8] that lead us to interpret our experiences differently. Digital technologies, more so than other artefacts, are situated within the broader *communicative ecology* in which the technology is used [2]. As the technology we use in our research has become increasingly robust, compact and mobile, it has become more feasible to deploy technology in real homes, workplaces and communities. These locations have become 'living labs', and various methods have emerged to support this approach. Technology probes [3], a method we have employed extensively in our research, involves the deployment of a functional prototype to inspire feedback while at the same time learning about the deployment environment. Likewise, work by Gaver *et al.* [4] has deployed intentionally vague technologies that the user is encouraged to experiment with over long periods, allowing purpose and meaning to emerge through this process.

When researching in the wild, we often also ask not just how users respond to prototypes in the short-term, but how they can be integrated into people's everyday routines and make a meaningful impact on their lives in the long-term, after novelty and other distorting factors have settled. Deployment over long periods allows natural usage to develop over time, taking into account the rhythms of everyday life.

These qualities raise questions that have not been well-addressed in existing literature. If we are deploying interventions intended to be integrated into users' lives and have some positive impact, what happens at the end of the study? Often these technologies are simply taken back to the lab or redeployed with new users, leaving study participants without a technology they may have come to value and that isn't available to purchase. For example, one of the families who trialled HomeNote, a family communication device deployed in several homes for several months, expressed a desire to keep it after the end of the trial [7].

Some research projects do indeed leave prototypes in the field permanently, but this raises problems of its own: research prototypes are not finished products and their developers are not capable of providing long-term technical support, meaning users are very much on their own once a project concludes. Consequently, the loss of a technology may occur despite the researchers' best intentions. The Hermes office door display system [1] was trialled over a period of many years, and several members of the department's staff found them to be valuable additions to their work routines. Although the displays remain in place, a number suffered hardware failures over time, leading one academic who particularly missed the door display's functionality to fashion his own replacement system.

To complicate matters further, participants are often involved in projects as co-designers and co-researchers, implying a considerable amount of effort on their part beyond simply using prototypes. Participants may spend many hours in focus groups and design sessions, or put considerable effort into completing probes and diaries. At the end of the study, participants may have little to show for

Submitted to "Research in the Wild: understanding 'in the wild' approaches to design and development" at *DIS 2012.*

this effort beyond a fleeting experience with cutting edge technologies and an insight into what might be possible in the future.

Our experience has exposed us to both the benefits and potential pitfalls of research in the wild, but one question in particular continues to cause unease: given that we are encouraging research participants to integrate our technologies into their lives, what happens when the research ends? In this position paper, we intend to outline the need for more exploration of the practicalities of technology handovers and continued support for research prototypes.

WRAY AS A LIVING LAB

Between 2006 and 2010, as part of a funded research project, we investigated the use of public displays to support rural communities in Wray, a small village in North West England. Having previously been used as a test site for a wireless mesh network, providing broadband Internet to the village for the first time and generating considerable goodwill, Wray began to act as a living lab for university projects. Over a period of four years, researchers worked closely with residents, including one 'champion' who acted as a primary point of contact, to develop a public display that supported the local community [9]. The long-term nature of this project was a key characteristic of the research: we investigated how use of the display emerged over time and how real experience with relevant technologies could help community members to engage in a participatory design process.

Over this time, we deployed a series of public display prototypes in the village, beginning with an extremely simple photo display, which comprised a touchscreen display powered by a compact PC (Figure 1). This was iteratively developed into WrayDisplay, which displayed photographs, upcoming events and advertisements, all of which were posted by residents themselves via a web application. During the final year of the project, the deployment was extended to include a second display installed in a local café. These displays had become integral in residents' photo sharing behaviours, with residents sharing over 1,500 photos by the end of the project, while the event listings and advertisements were able to augment existing methods of sharing this content.

At the start of the project, the villagers that provided initial input were told that any hardware deployments would remain in the village, an assurance that was repeated as the project neared is conclusion. This assurance arose from our appreciation of the large contribution made by participants during the study and the way in which the displays had become instrumental in sharing the community's photos. However, the technology could only be provided 'as is' without any guarantees—although participants were assured technical support would be provided on a voluntary basis to whatever extent was possible, it was also made clear that this could not be prioritised and might therefore be limited and delayed in practice.

Figure 1. WrayDisplay deployed in the post office.

In the face of possible hardware failures, residents proved to be more concerned about the display's content, particularly the photographs. Despite some of the photos now being available on the Internet, either on large photo sharing sites or on the village website, it was described as being "all over the place"—the display had brought together disparate sources of community content and one resident noted that recreating this collection would be "a lot of work". If (or when) the equipment failed, funds might not be available for replacements and the content might be lost. For this reason, the display's administrator was provided with a back-up tool that synchronised the display's content with her PC.

HANDOVER AND SUPPORT ISSUES

After WrayDisplay was left permanently in the village, a number of issues occurred that had not been problematic during the study itself. Technical problems were the most prominent, but we also saw a considerable drop in the volume of content and naturally faced an inability to continue the iterative cycle of development that had been carried out through the project.

Technical Issues

Given that the devices comprised of off-the-shelf hardware not designed to operate constantly in a public space for periods of several years, it was almost inevitable that hardware failure would be the first major problem to occur. Perhaps predictably, this occurred only a few months after the project's conclusion, when the PC running the shop display suffered a hard drive failure. Fortunately, the system architecture meant no content was lost: the second display continued to function and content remained accessible via the website and backup tool. Our champion contacted one of the researchers for advice, who provided a package of software and instructions on how a new PC could be set up to replace the display.

Theoretically, reinstalling the display software on a new machine should have been simple. However, difficulties arose due to a number of factors. Firstly, the researcher had moved to a different institution in a different part of the country. Secondly, the software, being a research prototype, was not as easy to install as commercial software, due to the configuration required and several dependencies (e.g. a database installation). One of the two touchscreens deployed also required specialised drivers that were not available to download. Finally, funding and spare PCs are not always in abundant supply. Our champion was able to donate an old laptop, but this was somewhat unsatisfactory by comparison to the very compact PC that had been deployed throughout the study.

Eventually, with support from several researchers at different times, the display was brought back online. The donated laptop was set up through a combination of email support, a site visit by the remote researcher and remote access. The broken PC was collected with a view to replacing the hard drive, which was later replaced with a new machine donated by one of the researchers as a goodwill gesture. However, from end to end, this process had taken some six months.

For the most part, these issues were simply a matter of logistics and communications, and a number of steps could have been taken at the point of handover to attempt to mitigate them. Most obviously, a package of software, documentation and tools should have been provided in advance to allow the software to be reinstalled. The precautions taken had instead centred on the backup tool with the assumption that there would be no attempt to repair the display. This proved not to be the case, particularly given the closeness of the failure to the end of the project. Training could also have been provided to ensure that key individuals in the village were confident to reinstall the software. We might even have budgeted for and provided a spare display setup. These are, of course, all common practice in the software industry, but rarely take place in an academic environment and are typically not factored into budgets and time plans. This raises questions around funding models for research in and wild, which may need to more explicitly address these problems and provide more lightweight means of funding small, ongoing activities relating to research projects.

For the time-being, the displays remain stable in the village, and valuable experience has been gained in quickly rectifying problems that arise. However, there are still a number of points of failure that would benefit from further consideration. For example, the displays still rely on university-hosted servers, meaning control of the system is still not entirely within the community's control.

Loss of Momentum

Although problematic, the hardware failure was relatively easy to remedy. More complex to address is the apparent loss of momentum and content seen after active research stopped. For example, only 173 images were uploaded during 2011, compared to 418 in 2010 and 327 in 2009. After such a long deployment period, we are confident that any novelty effect had been eliminated, and we saw steady levels of use over the length of the study.

Advertisements and events posted to the display have also seen a decline, and we have often seen these pages devoid of any current content. This poses a greater problem than photo content: unlike photos, events and adverts were time sensitive and expired, and we had quickly come to recognise the importance of having current content on the display to encourage use by others and signify appropriate content. Having the display regularly show no content is only likely to discourage subsequent usage.

This raises numerous questions about research in the wild. Although there is little doubt that the display had indeed played an important role in village life, the presence of researchers in the village and regular focus groups appeared to have created something resembling a Hawthorne effect, despite researchers not taking any steps to encourage use of the display. This is by no means to suggest that the display has been a failure outside of the project: though slower than previously, the level of content generation seems to have been enough to sustain it, even in the face of technical failures. It is, however, an important consideration for future studies.

It also invites us to consider how we can design for sustainable content generation over long periods of time. Certainly the reliance on one key champion runs the risk of this individual 'burning out' [6], leading to the collapse of the system. While it is certainly preferable to have content that is generated by community members, there may be alternatives when user-generated content is in short supply. For example, technologies might draw upon known sources of information, such as RSS feeds or social networking sites that already have a critical mass of users and are unlikely to suffer from a dearth of content.

Inability to Continue Development

One difficulty with the open-ended iterative approach that we utilised in Wray was knowing when to stop development. By the end of the project, the display was stable and had met most of the most common requirements that had emerged through the design process. However, although residents involved in the project might normally see computer software either as something immutable or as something that changed in response to factors beyond their control (and perhaps against their wishes in some cases), as participants, they because accustomed to an agile development process. Our champion in particular continued to suggest features that might be implemented, and we received expressions of interest from neighbouring villages who wanted their own displays. Unfortunately, with no development expertise in the village, further modifications

were not possible. Additionally, this meant that the display was not future-proofed and unable to adapt to changes in circumstance, both technical and social.

This highlights the importance of providing methods of exporting data from the system, such as the backup tool. Although the system itself cannot be entirely future-proofed, we can at least ensure that the data, which has considerable value to the community, can be exported in formats that are widely readable and can be transferred onto mainstream services.

Another possible approach to this problem is attempting to create the expertise required. Merkel *et al.* [5]'s approach to community technologies did not involve developing technologies *for* communities, but exploring their needs and then providing them with the skills to address these needs themselves. However, it is not immediately clear how this might extend to the more experimental technologies often seen in research projects. The display may well have already gone some way towards encouraging residents' enthusiasm for technology, but while several community members could have (and already did) maintain websites and other simple technologies, a network of public displays is a more complicated undertaking.

SUMMARY

Developing technologies that can be deployed in the wild is a delicate balance between utilising all the resources available to us as researchers, while ensuring that these resources do not become a cornerstone that, when taken away, causes the technology to collapse. It is easy for us to look back on our work and, with knowledge of what occurred subsequently, identify things we should have done differently. However, the intention of this paper is not to chastise ourselves, but to highlight the need for greater exploration of how technologies can be handed over the participants and supported. If this can be achieved, there is great potential for technologies developed as part of research activities in the wild to have a meaningful impact not just during the lifetime of the project, but for many years afterwards.

ACKNOWLEDGMENTS

The original WrayDisplay research was supported by a Microsoft European PhD Scholarship and the EPSRC-funded CASIDE project (EP/C005589).

REFERENCES

1. Cheverst, K. Fitton, D. and Dix, A. Exploring the evolution of office door displays. In *Public and Situated Displays: Social and Interactional Aspects of Shared Display Technologies*, Kluwer (2003), 141–169.
2. Foth, M. and Hearn, G.N. Networked individualism of urban residents: discovering the communicative ecology in inner-city apartment buildings. *Information, Communication & Society 10*, 5 (2007), 749–772.
3. Hutchinson, H., Mackay, W., Westerlund, B.B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H., Roussel, N., Eiderbäck, B., Lindquist, S. and Sundblad, Y. Technology probes: inspiring design for and with families. In *Proc. CHI 2003*, ACM (2003), 17–24.
4. Gaver, W.W., Beaver, J. and Benford, S. Ambiguity as a resource for design. In *Proc. CHI 2003*, ACM (2003), 233–240.
5. Merkel, C.B., Xiao, L., Farooq, U., Ganoe, C.H., Lee, R., Carroll, J.M. and Rosson, M.B.. Participatory design in community computing contexts. In *Proc. PDC 2004*, ACM (2004), 1–10.
6. Redhead, F. and Brereton, M. A qualitative analysis of local community communications. In *Proc. OZCHI 2006*, ACM (2006), 361–364.
7. Sellen, A., Harper, R., Eardley, R., Izadi, S., Regan, T., Taylor, A.S. and Wood, K.R. HomeNote: supporting situated messaging in the home. In Proc. *CSCW 2006*, ACM (2006), 383–392.
8. Suchman, L. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press, New York, NY, USA, 1987.
9. Taylor, N. and Cheverst, K. Supporting community awareness with interactive displays. *IEEE Computer 45*, 5 (2012), 26–32.