Ubiquitious Computing needs to catch up with Ubiquitous Media

Anil Madhavapeddy[1], Nick Ludlam[2]
University of Cambridge Computer Laboratory[1], Interceptor Communications Ltd.[2]
{avsm2@cl.cam.ac.uk[1], nick@intercomms.co.uk[2]}

1 Introduction

The field of Ubicomp is, in a word, "varied". A typical Ubicomp research team is highly multi-disciplinary, consisting of hardware, software and social engineers collaborating to achieve the vision of invisible computing in our day-to-day environments. This variety has led to research in a myriad of fields such as context-awareness, sensor networks, low-power computing, activity inference, and location sensing infrastructure. However, the field currently lacks a unifying factor that will drive this scattered research into real-world deployments.

This position paper points out a vertical market that could hugely benefit from the expertise of the Ubicomp research community—the media industry, currently in the throes of a digital revolution. Analogue broadcast systems such as radio and television are moving to digital formats, and consumers are starting to get familiar with "digital lifestyle appliances" promoted by industry such as the popular iPod[1] and Tivo. We discuss some of the problems created by this transition, and point out a golden opportunity for the Ubicomp research community to get involved in creating an effective platform for the future of ubiquitous media for consumers.

There are traditionally two mechanisms of delivering video content to users: (*i*) wireless or cable broadcast medium (e.g. live television); and (*ii*) physical media such as video cassettes or DVDs. More recently, a third delivery mechanism has been added to this list: the Internet, driven by the rising penetration of broadband services to US homes. With more bandwidth, consumers are able to search for and download specially encoded video files which let them view selected programs "on-demand". Aside from the copyright and legality issues, the sheer volume of activity on this front (e.g. the rise of peer-to-peer software such as Kazaa) demonstrates the strong consumer desire to break out of the "push model" of television into a more interactive and flexible content consumption model.

This paper does not attempt to create a taxonomy of problems facing the media industry. Instead, we highlight some of the more interesting problems that have arisen in recent years and examine some relevant Ubicomp projects in the field that apply to them. Finally, we describe the beginnings of a project in conjunction with the British Broadcasting Corporation (BBC) that will start to tackle some of these problems.

1.1 Television On-Demand

Television has traditionally been a broadcast medium, with the consumer selecting a video stream to watch from a limited set of channels. The advent of cable television has increased the number of channels into the thousands, but the basic broadcast model is still intact. If consumers wanted to watch a video "on-demand", they would normally purchase, rent or borrow a DVD or cassette.

Recently, companies such as Tivo have been selling Personal Video Recorder (PVR) appliances which hook into a television feed and record content based on user selection to a local hard-drive. When consumers wish to watch television, they can simply flick through all the recorded programs and search until they find the content they feel like watching. The Tivo attempts to record content based on user's past preferences, e.g. "I like the Simpsons" resulting in all episodes of the cartoon to be stored. This approach is limited to programs the user actually watches on that Tivo, and breaks down when the user's interest rapidly change in the outside world (e.g. moving house or falling ill).

Problem 1: *Video appliances need to be able to infer user demand and record the content that they will want to watch based on their current activities in the wider world.*

Activity inference has been looked at by a number of Ubicomp projects. Fishkin et al discuss how objects with RFID tags attached could allow interactions with items in the home to be detected automatically [3]. Koile at al use vision to infer "activity zones" of user actions [5]. The PlaceLab project [6] uses commodity wireless hardware to retrieve user location, another useful beacon for context inference. Microsoft's AURA [9] lets users scan barcodes on everyday objects and register them on a central website—this information could be used to trigger video recording in the home (for example, scanning organic health food at the supermarket could allow the predictive recording of news snippets about that topic).

If these systems were integrated with a digital media platform, users would see the real effect of activity inference as television they want to watch is made available to them in the background.

The shift from live broadcast will have a profound impact on the multi-billion dollar advertising industry, who are already finding users capitalising on PVR systems to

[1] In an effort to cut down on references, we do not cite terms which can be easily looked up via an Internet search engine.

filter out advertising during automated recording. However, as Google has demonstrated with their successful AdWords business, users are not averse to advertisements, only irrelevant ones. Providing advertisements to users *peripherally* instead of obtrusively, e.g. when they are engaged in activities other than watching television may win back user interest.

Problem 2: *Context-aware and peripheral advertisements may drive the future of consumer video-on-demand advertising.*

Ubicomp offers a good breath of research in this field. Project Aura's vision of Distraction-free Ubiquitous Computing [1] seeks to minimise the amount of interruption a user suffers. Mankoff et al also offer insights into the use of "ambient displays" to occupy the periphery of user attention [7]. When combined with activity inference, these approaches could result in an advertising experience for the user that is very different from the existing interruptive delivery mechanism.

1.2 Physical Storage

Traditionally, the standards governing rich media have been dictated by the physical storage format being used. For example, the MPEG-2 bit-rate of 3-10Mb/s allows a dual-layered commercial DVD with a capacity of 8.5 GB to store around 3-4 hours of content (enough for a typical Hollywood movie with additional features).

Internet delivery, and the ability of consumers to encode their own content onto digital media such as hard-drives eliminates this requirement. For instance, the Apple iTunes music program allows the creation of "MP3 CDs" or "audio CDs". Audio CDs can be played back by normal CD players, while MP3 CDs use a more advanced codec that packs more music onto the CD, but can only be played by MP3-aware players. Similarly, modern mobile phones are capable of playing video that has been transcoded to the 3GPP format, and the DivX format is popular for the storage of television programs without a large degradation of image quality.

Problem 3: *The proliferation of digital media clients used by consumers has created the need for a flexible way of storing large amounts of high-quality data on their person or at home, and transcoding it on-demand to other devices such as mobile phones or televisions.*

In Ubicomp, the Personal Server project [11] provides an insight into how users can effectively carry around large amounts of data on their person. However, such a device would need high-bandwidth wireless links to the outside world, and more computation power to be able to transcode content on-demand into different formats as the user requests them (or by automatically sensing the class of device requesting the media).

As the user's collection of personal data increases, the consolidation of storage becomes a double-edged sword. The user benefits from centralizing large cassette and DVD collections into a single portable drive, but then has to deal with the long-term backup of unreliable physical media such as hard drives. Dealing with how to securely synchronize rich media across a home network and ensure there are no single points of failure is a difficult problem across heterogeneous networks and devices.

Problem 4: *There is a need for storage devices to infer and indicate how much of the information they contain is "unique"; that is, not backed up or available elsewhere in the user's storage network.*

Many Ubicomp researchers are very familiar with programming low-power embedded devices; storage synchronization is an active research area in the networking community with efforts such as the Co-operative File System [2]. Combining this research with work on tangible interfaces [10] on personal storage gives us an exciting glimpse into how a future personal storage device might look: a wireless device with a physical indicator as to how "unique" the data it contains is. Giving users an immediate notion of the value of a storage device would warn them to purchase more storage to back it up, not put critical data on it, or even just be very careful handling it. Anyone who has accidentally erased the last copy of a picture from a digital camera will appreciate how valuable this is as more of our media storage shifts into digital intangibility.

1.3 Publishing for All

The digital media revolution extends to production equipment as well as playback. The majority of cameras sold in recent years have been digital, and camcorders to record video have dropped to affordable prices. This has led to consumers having an increasing need for video editing, conversion and viewing software. In addition, they also need to easily share it with friends and family.

Problem 5: *There exists no standard mechanism for "hyper-linking" into video based on a variety of factors such as its location, time and content, in order to allow other Ubicomp interaction techniques such as visual tags to be used by consumers to easily share rich media.*

Universal video links might be links to local devices on the wireless network, or if that device is unavailable, a central Internet server for a slower download. One scheme proposed by Kindberg [4], which allows users to access local physical resources as Web resources, could be extended to solve this problem. Once these hyperlinks can be expressed, traditional Ubicomp interaction techniques such as visual tags [8] can be used to make it easy for consumers to print out "link cards" to give to their friends, or just dynamically generate them between two compatible devices such as mobile phones (both of which are equipped with a screen and a camera).

2 One Platform to Rule them All

Open-source software support for video processing is rapidly gaining ground. Projects such as MPlayer and FFmpeg offer backends for performing video conversions to and from many formats such as DivX, MPEG or 3GPP.

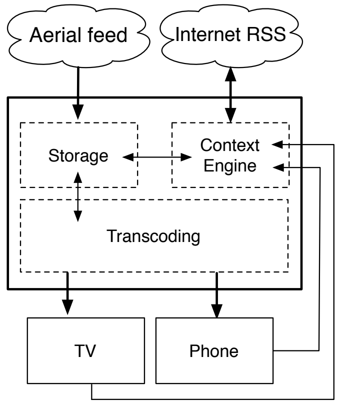

Figure 1: Architecture for our prototype Contextual PVR

MythTV is a PVR project which uses Linux and video hardware to record and manage streams of television data.

However, none of these individual elements mesh together to genuinely move the user experience beyond the industry state-of-the-art. We have begun a joint collaboration with the British Broadcasting Corporation (BBC) to investigate the potential impact that large-scale local storage of digital content provides. We have constructed a platform with enough storage to continuously record 7 days worth of television across all BBC channels broadcast in the UK. This platform will be used to explore the socio-technical challenges introduced to *both* the broadcasters and consumers in the emerging digital market.

As described earlier in this paper, there are numerous challenges to solve: (*i*) allowing users to search the vast amount of recorded television without requiring a keyboard and mouse; (*ii*) analyse how viewing habits and user expectations change as media becomes truly "on-demand" by being continuously recorded and stored; (*iii*) classify videos from sources such as Internet information feeds via RSS or from the XML meta-data held in the broadcasts themselves; and (*iv*) a "context-feed framework" to allow external devices such as televisions or mobile phones to provide feedback about user activity in the real world. Underpinning all of these is the requirement for "invisible content security" via a Digital Rights Mechanism that protects intellectual property while still granting fair use rights and not suppressing spontaneous interaction between individuals.

We urge the Ubicomp research community, consisting as it does of a rich variety of hardware, software and social engineers, to step forward and examine how their research could help improve the handling of digital media for the next generation of home appliances. Ubicomp has so far been a scattered research field, with several discrete (though interesting) areas of research such as low power computing, context-awareness, location-awareness, sensor networks and activity inferencing. Demonstrating our ability to solve real user issues through the unified deployment of our technology into a rapidly transforming digital society would finally initiate Ubiquitous Computing as a viable long-term research area.

3 Acknowledgements

This work was sponsored by Intel Research Cambridge and Interceptor Communications Ltd. We are grateful to Richard Sharp, Eleanor Toye, David Scott and Alastair Beresford for their caffeine-fuelled "ubi-rants" which shaped the thoughts presented here.

References

[1] Project Aura: Toward distraction-free pervasive computing. *IEEE Pervasive Computing*, 1(2):22–31, 2002.

[2] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)*, Chateau Lake Louise, Banff, Canada, October 2001.

[3] K. Fishkin, B. Jiang, M. Philipose, and S. Roy. I sense a disturbance in the force: Long-range detection of interactions with rfid-tagged objects. In *Ubicomp*, September 2004.

[4] T. Kindberg. Implementing physical hyperlinks using ubiquitous identifier resolution. In *WWW '02: Proceedings of the eleventh international conference on World Wide Web*, pages 191–199. ACM Press, 2002.

[5] K. Koile, K. Tollmar, D. Demirdjian, H. E. Shrobe, and T. Darrell. Activity zones for context-aware computing. In *Ubicomp*, pages 90–106, 2003.

[6] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. Schilit. Place Lab: Device positioning using radio becons in the wild. In *Proceedings of Pervasive 2005*, LNCS. Springer-Verlag, May 2005.

[7] J. Mankoff, A. K. Dey, G. Hsieh, J. Kientz, S. Lederer, and M. Ames. Heuristic evaluation of ambient displays. In *CHI '03: Proceedings of the conference on Human factors in computing systems*, pages 169–176. ACM Press, 2003.

[8] D. Scott, R. Sharp, A. Madhavapeddy, and E. Upton. Using visual tags to bypass bluetooth device discovery. *ACM Mobile Computer Communications Review*, January 2005.

[9] M. A. Smith, D. Davenport, H. Hwa, and T. Turner. Object auras: a mobile retail and product annotation system. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 240–241. ACM Press, 2004.

[10] B. Ullmer and H. Ishii. The metadesk: models and prototypes for tangible user interfaces. In *UIST '97: Proceedings of the 10th annual ACM symposium on User interface software and technology*, pages 223–232. ACM Press, 1997.

[11] R. Want, T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. The personal server: Changing the way we think about ubiquitous computing. In *UbiComp '02: Proceedings of the 4th international conference on Ubiquitous Computing*, pages 194–209. Springer-Verlag, 2002.