Pers Ubiquit Comput
DOI 10.1007/s00779-006-0116-1

ORIGINAL ARTICLE

Ubiquitous computing in the real world: lessons learnt from large scale RFID deployments

Shin'ichi Konomi · George Roussos

Received: 15 January 2006 / Accepted: 1 June 2006

Abstract Ubiquitous computing technologies are slowly finding their way into commercial information systems, which are often constructed at considerably larger scale compared to what is possible in research demonstrators. Furthermore, lengthy and costly preparation or upgrade of existing infrastructures, training of employees and users in the new ways of working, controlled introduction of new functionality, features and services to manage risk, unexpected behaviors due to the wider variety of possible real-world situations, incremental approach to systems development so as to better identify successful aspects, regard for the economics of systems as a core requirement, and selection of open or closed systems are all issues that are mostly outside the scope of current ubiquitous computing research but play a critical role in industrial deployments. In this paper we review two case studies of fully operational Radio Frequency Identification-based systems: the Oyster card ticketing system used at the London Underground in the UK, and retail applications deployed at the Mitsukoshi departmental stores in Tokyo, Japan. We examine each case in terms of technologies, user interactions, and their business and organizational context and make several observations in each case. We conclude by drawing general lessons related to ubiquitous computing in the real world and identify challenges for future ubiquitous computing research.

1 Introduction

The core enabler for ubiquitous computing is technology that makes the physical and digital worlds interlinked and thus intimately related. Every object in the world we live in has a digital representation that follows the situation of its real self and vice versa. This unique linking of bits and atoms opens up numerous possibilities for new computing interactions which are currently explored by ubiquitous computing research. One of the main implications of this interlinking is that people, places and things acquire unique machine readable identities within systems of very large scale that must be accommodated within inflexible physical constraints and constantly changing usage context. To be sure, auto-identification capability opens up membership to ubiquitous computing systems for numerous entities and potentially results in massive increases to the number of constituent elements and system complexity. Understanding the issues raised by this increased complexity and exploring solutions can be hard to carry out in lab-based studies or case studies of limited scale. Such understanding often requires that experiments are carried out at scale, a fact that invariably implies high costs that are almost certainly prohibitive in a research context. Nor is it possible to identify and address such issues using large scale simulations, as these are limited by the simplified assumptions involved, and cannot take into account the emerging behaviors caused by real users. In this paper,

S. Konomi (✉)
Center for Spatial Information Science,
The University of Tokyo, Tokyo, Japan
e-mail: konomi@csis.u-tokyo.ac.jp
URL: http://www.iis.u-tokyo.ac.jp/~konomi/

G. Roussos
Birkbeck College, University of London, London, UK
e-mail: g.roussos@birkbeck.ac.uk
http://www.dcs.bbk.ac.uk/~gr

we attempt to record early experiences with commercial uses of ubiquitous computing technologies at scale with a view to extracting useful lessons that can only be gained in real world deployments.

Both systems discussed in this paper depend on Radio Frequency Identification (RFID), one of the earliest technologies to effect this physical/digital coupling and to support transparent automatic identification. Although there are ubiquitous computing research efforts currently underway to explore alternative auto-identification technologies that may provide considerable advantages using more powerful processing capabilities and holding capacity, RFID is readily available today at relatively low cost and has already found numerous applications. In particular, there are several applications employing RFID as a core component, which regularly support large numbers of users—two such systems are the focus of this paper. In both cases, the real world imposes a number of conditions that are not often addressed when research is conducted with reduced scale models and we put particular emphasis on stresses and constraints placed on these systems by practical, commercial, regulatory and market considerations.

Full-scale commercial implementations of ubiquitous computing technologies can provide considerable insights into their use since—rather than making assumptions about the structure of their operational context and supporting infrastructure—such systems must consider and adhere to the requirements and actual conditions of real operational environments. One such requirement of considerable impact is that new ubiquitous computing system components must be introduced so as to interoperate with existing and legacy (physical and information) systems already in place and must measure and regulate their impact on existing infrastructure so as not to catastrophically disrupt the capability of the organization to provide services. Other considerations include the possible need for organizational changes required to accommodate and indeed to benefit from the new capabilities for service provision, and of course how to balance cost against business benefits and the technological capability of the company and its suppliers. This view of ubiquitous computing situated within a market environment can offer a more realistic context for research in ubiquitous computing, in particular it can help frame future work so as to better meet the needs of the real world.

The two case studies that we consider in this paper are the Prestige project run by Transport for London which developed, deployed and operates RFID-based ticketing across the metropolitan public transport network in London, UK; and the consumer-facing retail implementation of item-level RFID for high value apparel products at the Mitsukoshi department stores in Tokyo, Japan. First we examine the particulars of each case and then we attempt to draw common lessons from both cases, which have implications for ubiquitous computing in a more general context.

2 The Prestige project

Prestige is the name used by Transport for London to refer to their ongoing activities to develop new ticketing and service infrastructures across the metropolitan public transport network in London, UK. The project was initiated in 1998 and has succeeded in designing and implementing a comprehensive overhaul of the legacy system which has been in place since the 1980s across all modes of public transport within the commuter area of the city of London. Interestingly for the current discussion, Prestige identified the use of Radio Frequency Identification (RFID) as the key technology on which to develop its new solution, the so-called Oyster card, which is used to provide access and collect payment for the use of the Metro system. As a result of this decision, Transport for London currently operates the largest implementation of RFID, which is used daily by more than 5 million commuters in their trips.

In the following sections we first discuss the motivation for this project and in particular the service and business drivers behind the choice of RFID. Then we examine the rationale behind the technical and operational solutions selected for RFID deployment and how they achieve the required balance between business and performance requirements. We conclude with a discussion of the lessons learnt and how they relate to current research in ubiquitous computing.

2.1 Business and service considerations

Until the early 1980s commuters had to pay separately for each mode of transport that they used as the bus network, the Tube and commuter railways employed separate systems. This situation inevitably produced considerable inconvenience as it introduced significant delays due to the requirement to use multiple travel passes when exchanging means of transport. An integrated ticket, commonly called the Travelcard, was introduced for the first time in 1981 to cover the bus and the Underground, and was extended in 1985 to include mainline railways. Travelcards are still in use today but have several core limitations since they are based on magnetic strip technology.

In 2000, the Greater London Authority integrated all public transport services including the buses, the

Underground, the Docklands Light Railway, the Croydon Tramlink and the London River Services, into a single organization named Transport for London (TfL). This was the first step towards a unified, tightly coordinated and efficient metropolitan transport system. The role of the TfL is exactly the management of transport services across the city and its main priority is to build the infrastructure that can support the city's rapid expansion in the coming decades. As a measure of the complexity of the task it should be noted that management of the London Underground services alone involves approximately 500 trains at peak times, 253 stations owned (275 served), over 12,000 staff and numerous engineering assets. Moreover, the London Underground is in continuous operation since 1863 and as such it has inherited a variety of assets that have been created using legacy technologies but cannot be completely or concurrently overhauled due to the disruption this would cause to the operation of the system

Public transport is a critical component of the urban infrastructure of any city and its uninterrupted running depends heavily on it [4]. In London in particular, public transport accounts for 36% of the 27 million journeys made on an average day and grows steadily at a 15% annual rate. This growth compares favorably against private transport, which increases at the much slower rate of just over 1% (since 2000), especially taking into account the fact that the population growth is just below 3% during the same time. Travel by London residents is primarily within greater London as nearly half of all trips taken are within this area and two-thirds either started or ended there. The importance of public transport becomes even more evident if we look at the number of people entering central London: on an average weekday (in 2004) 1 million people entered this area during the morning peak time (7 am–10 am). Of these, 88% used some form of public transport, with bus and Underground traffic on the increase and a slight decrease in rail compared against previous years. At the same time, road traffic remains almost constant since 1999, with a small reduction in total traffic for major roads (motorways and principal roads) and an increase on minor roads.

TfL was also given the task to work closely with the private sector using a new model for collaboration, the so-called Public Private Partnerships (PPP). This approach was (and still is) viewed as an efficient way to mitigate critical service and operational inefficiencies that exist within the transport system. These shortcomings have been the result of more than two decades of underinvestment and limited attempts to modernization, which have significantly deteriorated the quality of services offered. The perceived advantage of PPP is in the use of private sector companies that are seen as better positioned to deliver improved performance through access to specialist technical resources, greater financial certainty and more efficient use of financial resources, project management skills, and knowledge of good industry practice in whole-life asset stewardship. PPPs are implemented as Private Finance Initiative (PFI) contracts that define the responsibilities and roles of the private and public sector partners. Prestige is one of three such PFI partnerships developed by TfL and is focused specifically on improvements in the operation of the London commuter system and to implement infrastructures that can guarantee higher service levels that can cope with the expected growth in commuter numbers.

One area where potential efficiency improvements can be achieved is in streamlining ticketing, thus reducing the time required to access trains. Improvements can be made in self-service ticketing (for example in the operation of credit card processing facilities by passengers without the involvement of staff) and payment, faster entry and exit and better transportation between the gates and the platform. These improvements are the core of the Prestige project with a number of additional objectives identified. Reduction of ticket counterfeiting and ''fare dodging'' in particular, costs operators 5–6 million British pounds on average each year.

Upgrades in all aspects of the entry and transfer subsystems are expected to cost over 250 million British Pounds over a 30 year period of which only about 10% is associated with the implementation of RFID technology. Clearly, the cost of the technology itself in this case can be easily facilitated and indeed it could be self-financed to some extent by the gains due to the reduction in fraud.

2.2 Balancing requirements and capabilities

Travelcard ticketing is based on magnetic strips which are printed on branded heavy paper cards carrying the TfL watermarks. Ticket vending machines accept payment in exchange for magnetic strip tickets that store the appropriate information representing allowed trips on the network. Tickets must then be inserted into dedicated slots at the gates at the beginning and end of a trip, to allow the read-write units to confirm that the passenger is entitled to travel. In both cases, Travelcard tickets must be read and written using suitable electromechanical devices installed in ticket vending machines and at gates or turnstiles at each station.

The Travelcard can decrease the amount of time spent by passengers in transactions, but the rate at which passengers pass through the turnstiles is somewhat lower than when using manual ticketing methods. The reason for this is that magnetic strip tickets require specific placement of the card so that proper contact with the reader head is made [5]. This implies in turn that in the vast majority of turnstile designs the card must be inserted into a read-write unit and then recovered after a short but noticeable time period. The aggregate time required to insert and retrieve the ticket results in a slower process to gain entry.

Furthermore, the tickets themselves are printed on flexible material with limited durability and are thus prone to faults. This is particularly so when tickets are stored without a protective cover, which is common practice for single journey or daily passes. A variety of external factors (creases due to unintended folding, water, oily residues and so forth) can easily prevent the reader unit from accurately retrieving the stored data in which case manual intervention from staff is required either to operate the gate after visual verification of the ticket or for the provision of replacements.

Magnetic strip ticketing has another drawback in that—depending on the particular design—it is relatively easy to produce counterfeits. This is less of an issue for read-only cards or for systems that operate at increased coercivity levels but adoption of either of these options leads to much less flexible system. Another problem with magnetic strip technology is that read-write units require the use of mechanical rotating parts (to swipe the magnetic strip) and are thus subject to relatively frequent failures and require frequent maintenance.

2.2.1 *User interaction*

The Prestige project elected to replace Travelcard-type ticketing with RFID which was viewed as the technology that can deliver the required performance improvements. RFID systems only require that the card and the read-write unit be in close proximity thus eliminating the need for the card to be inserted into a unit for precise positioning. A suitable reader-writer unit can be placed in any of a number of convenient locations on the ticket vending machines or the gates, for example on the top of the entry doors at easy reach for the majority of the passengers as seen Fig. 1. Moreover, the card itself can be made out of durable strong plastic which reduces wear and tear of the card and significantly reduces reading errors. This fact results in considerably shorter times required to operate the turnstiles.

The collocation of both magnetic and RFID cards readers on the same gating equipment and the fact that both magnetic and RFID cards have the same size gave rise initially to some interesting behavior. One of the pre-public release findings of the Prestige project was that passengers used to Travelcard tickets frequently inserted their new Oyster cards into the slot of the magnetic strip reader, which compressed and damaged the card. Subsequently a provision was made so that the plastic enclosure of the RFID chip be able to withstand at least 100 such events.

Fig. 1 Location of RFID units on entry gates: RFID readers have circular yellow reading areas. The entry slot for magnetic cards is also visible on the vertical surface of the gate

Although RFID cards can be operated at a distance and thus do not require contact between the reader and the tag antenna, in practice two considerations also had to be factored into an appropriate design and the associated selection of reading range. First, it must always be clear which card is presented to which gate by which passenger to ensure that correct charges are applied, even in locations where gate density is high resulting in a large number of individuals using the system concurrently and at close proximity to each other. To achieve this effect the range of the system was restricted to only a few centimeters, a choice which provided the desired effect. Second, cards must be read and updated with very high accuracy which implies good placement of the card antenna within the field created by the reader. A well know limitation of RFID systems is that when the plane defined by the coils of the tag antenna is approximately perpendicular to the field created by the reader, then the coupling effect is weak and leads to failure to charge the RFID chip. In pre-production trials with the system, such erroneous use of the card was observed far more frequently than expected likely due to the location of the Oyster card reader on the gates, leading to indeterminate transactions. As a result usage guidance provided by TfL suggests that the card is ‘‘touched in and out’’ with the card placed flat on the reader (cf. Fig. 1). Despite the fact that contact is not required, touching the card guarantees that the orientation of the antenna is appropriate and a successful read or write very likely. Even so, one of the non-operational hardest problems that still remain is verification that the card is fully up-to-date, an issue discussed further below.

A clear advantage of the new RFID-based approach is the considerable reduction in time required from the point of entry to the station until the arrival on the platform. The new ticket validation process has led to reductions of approximately 3–4 s which has in turn led to cancellation of the gate bottleneck and thus higher throughput and less congestion as passengers are distributed more evenly along the system. This has been one of the major successes of the system in terms of service improvement. The time benefits of this technology are even more evident in buses: in the past it has been impractical to operate magnetic strip read-write equipment on buses and as a consequence ticket issue and validation of passes was carried out manually by the driver. The small form factor and the lack of any mechanical parts in the reader have made viable the use of the Oyster card in this case and have accelerated the ticketing process which is now conducted automatically. In the longer term, the introduction of the current system is expected to have even more significant implications for busses in particular as it allows the possibility to operate cashless carriers with security and safety gains.

From a usability perspective, the biggest success of the Oyster card seems to be what Brian Dobson of the Prestige project calls the *fumble factor*: ‘‘before, with magnetic, you’re taking it out of your wallet, you’re putting it through the gate, you’re picking it up and you just don’t have to do that anymore. So the fumble factor was a major service offering and probably the biggest reason people use it, because it’s so easy to use and you can re-use it.’’ Indeed, together with the financial attraction these two factors seem to be the most useful aspects of the card for passengers, that it need not be taken out of a wallet and that it can be used without wearing off for long periods of time.

Another system design challenge common between Prestige and other documented ubiquitous computing technologies that provide some degree of ‘‘transparent’’ operation, is how to best provide feedback of successful or unsuccessful operations to the passenger. Again, the emphasis in Prestige has been on support of the highest possible throughput which in this case implies that feedback should be simple and unambiguous. The Oyster system uses sound and simple red or green LEDs to provide confirmation of successful authorization as well as embedded displays in the gates and on ticketing machines for more complex interactions. In this way, access authorization is separated from general travel management and payment issues: the former is optimized for speed and is carried out using simple feedback and the later is designed with ease of use and adequate information provision in mind.

2.2.2 *Information management and networking*

Another advantage of RFID cards is that they have significantly higher information holding and processing capacity. This extra capability makes it possible in addition to storing and retrieving the code that describes the class of ticket (and thus allowed trips) to uniquely identify each individual card and its owner, the exact value stored on the card and to maintain the record of the most recent travel history on the card itself. The last feature in particular can be used to improve system robustness.

The availability of unique identifiers that can pinpoint a specific card also allows preventing fare evasion. Unique identifiers are used in combination with suitable cryptographic protocols to authenticate a card to the system, and to provide a unique key for each card and thus provide protection of the communication between the card and the reader thus making the construction of counterfeits very hard overall. Moreover, unique card codes also imply that specific cards can be singularly recognized and associated with activities that violate the rules of acceptable use. Such cards can be subsequently hotlisted thus preventing the passengers that carry then from entering the system. Management of hotlists is an interesting aspect of the system: lists are downloaded to gates and bus ticketing equipment which carry out locally the verification of the access credentials presented by the cards. The lists are refreshed daily to include new card IDs that have been found in violation and are reviewed regularly to remove cards that have appeared inactive for longer periods. This strategy does not completely guarantee the elimination of fare evasion by well organized and determined commuters who are committed to perform complex usage patterns employing multiple cards. Nevertheless, it does provide a good compromise in the context of overall risk management approach adopted by the system and fits well with the system design philosophy which gives precedence to robustness and failure tolerance.

Indeed, system design driven by appropriate risk management considerations is a tactic employed consistently by the Prestige project. Notably, a core design tradeoff relates to balancing system survivability against its ability to authenticate users accurately and provide access to transport services even when the card management back-end becomes unavailable. In this case, the transport system as a whole will continue to operate for at least seven days without any noticeable disruption to passengers at the minor additional risk of increased unauthorized access. This design decision has had great success until now and in practice the system has been operating without any high-profile or large scale disruptions for several years.

The main ingredient that has guaranteed such undisrupted operation is the heavy use of caching and replication architectures at all system levels. The Oyster card itself employs its extra storage capacity—in addition to the unique identifiers and in some cases some personal identification data—to hold records of the most recent read-write events observed and the active value of the ticket. Even in cases when read-write errors lead to indeterminate transactions, the back-end system can resolve the situation and instruct gating equipment to restore the card into a valid state at subsequent interactions. While this opens up a small window of opportunity for fare dodgers the financial impact would be localized and still compares favorably against the magnetic strip system. The success of this approach is also highlighted by the fact that since the introduction of the Oyster, revenue under similar traffic conditions has been increased by 10% and at the same time ticket inspectors have been reassigned to alternative control points.

A complication that in smaller installations of RFID is a mere annoyance or can be addressed effectively through longer user training (which is not an option here given the scale of this system) is that RFID cards despite being successful in authenticating passengers who gain access to the system, may fail to confirm the successful update of the transaction history maintained on the card itself with the new trip details due to synchronization errors. This happens with some regularity and results in an increased degree of uncertainty regarding the status of the system. This problem has been addressed by differing the processing of indeterminate transactions to the back-end that can make a better job at establishing the true status of the system as it can view the totality of the sequence of recorded transaction and the local status of the card is discarded in subsequent writes.

2.2.3 *Fraud*

One of the main business drivers for the introduction of the Oyster card has been the reduction in fraud. This requirement must be balanced against performance since a system that would completely guarantee that fraud is not possible would inevitably incur a very high computational overhead. This overhead can easily outweigh the performance gained from the introduction of RFID. The combination of multi-level caching (at the gate, station, back-end and analytics layers of the system) and hotlisting at entry and exit points provides a compromise that both provides good performance and limits the opportunities for fraud within affordable levels—in practice no significant violations have been identified in the years of operation of the RFID system. At the same time, offline activities including cross-checking of all transactions and verification at the back-end ensures delayed, albeit accurate, identification of deviant situations.

2.2.4 *System scale*

At its current state the Prestige system has created the largest RFID-based systems in operation today incorporating over 13,000 readers in 273 stations and 8,000 buses with over 4 million cards used on a regular basis as of the time of writing. Another 2,600 RFID readers are installed in external retail points, primarily newsagents. The successful implementation of a system of this size required a long period of planning and preparation that lasted over 6 years, with distinct incremental phases that built on top of each other. This prolonged period of preparation has been critical for the success of the project since it allowed for the management of the impact of the new technology on existing systems, particularly sufficient time to build up capacity and thus prevent consumer demand from overflowing the system. Internal corporate processes also benefited by allowing sufficient time for lessons learnt in earlier phases of the project to be applied to later stages. This phased approach also allowed staff and customers to become first familiar and then confident in using the new system. Specifically, large scale London-wide problems that would inevitably become high profile and damage passenger confidence have

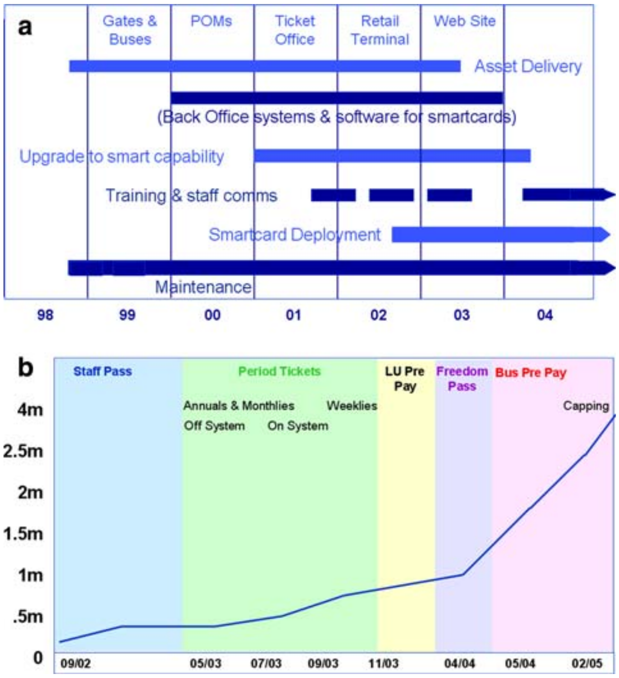

Fig. 2 **a** Phased introduction of the ticketing, gate and back-end technologies for the Prestige project (Dobson S6). **b** Number of Oyster cards in circulation growth (S10). Courtesy of Brian Dobson at TfL

been avoided further building public confidence in the system. Finally, as no comparable RFID system in terms of scale has been built before, the phased introduction of functionality allowed for the identification of errors or system limitations that became evident only at full system scale testing but do not appear in limited testing phases of the project. Figure 2 shows the timescales of the different project phases and the growth of the number of cards in circulation.

2.3 Lessons learnt

Having discussed the context and some of the features of Prestige in the previous section, we now turn our attention to lessons learnt from this project. In particular, we examine issues related to privacy protection and organizational changes that have been required.

2.3.1 Privacy protection

Unlike other studies with RFID-based consumer systems [3, 9] the Oyster project has been neither criticized by privacy groups nor phased significant consumer reactions. Quite the contrary, it appears that the system has been adopted easily and is trusted across most segments of the user population. Compared against privacy-related user requirements for similar systems recorded in past studies, it appears that all elements of common user concerns have been addressed effectively. Firstly, different ticketing schemes allow for differing amounts of information to be collected so as groups of consumers with different levels of sensitivity in issues of privacy can use different types of tickets at no extra cost. For example, the pre-pay version of the Oyster card can be used completely anonymously or at least as anonymous as cash payments. The only downside in anonymous use is that when it is lost the remaining credit on the card cannot be recovered in the same way that lost cash cannot be retrieved.

Secondly, TfL has a clear organizational identity which is well understood and trusted by the general public. Since only TfL-operated machines can read and write the cards it is always clear who is the transacting party in all consumer interactions, and there is full clarity about the purpose and the use of the collected data. More importantly, there is no ambiguity regarding situations when the cards are actually read and data recorded for further processing. TfL has put considerable effort to take advice on appropriate safeguards of complaint use of personal data, to clarify the use of data collected and more importantly to implement appropriate measures to process the data in accordance with the Data Protection directive.

Finally, the factor that is most commonly identified as the most critical for the success of a system that employs personal data, namely the balance between loss of privacy and perceived benefit, seems to have been addressed and the solution selected appears to be proportional and acceptable to commuters. Indeed, the system offers considerable advantages in terms of convenience in buying tickets and gaining access to the system (Dobson's fudge factor) at a restricted loss of personal data which is directly comparable to that of credit card use. While many researches have and are debating the relative merits of different strategies and models for dealing with the economics of privacy [1] it seems that TfL with Oyster has already identified a practical and viable tradeoff within its niche that is widely acceptable to the public.

2.3.2 Organizational issues

The second lesson learnt from the Oyster system is that while technology can be developed and deployed relatively speedily, the organizational and stakeholder issues require much longer time frames for successful strategies to emerge. To this end, a guiding principle that proved successful in Prestige is to set significant and realistic targets that require a limited number of changes to occur at any one time. When engaging the organization internally it is necessary to identify all the roles and processes affected and allow for the inevitable learning curve to catch up. In particular, operational staff training and internal communications are critical and the associated tasks increase in complexity with the size of the organization.

Though it may be common sense that systems that have millions of users on a daily basis should be durable and easy to use and thus designed accordingly and heavily tested, it is nevertheless unlikely that the transition will be seamless. For this reason, it is necessary to take into account that even the simplest system when operated at large scale will cause some degree of confusion that must be dealt with. Consequently, supportive measures have been in place well ahead and cater for the issues raised during the user learning curve. While this is probably an observation that is common to the deployment of any information system and might not refer particularly to ubiquitous computing systems it is necessary to make two relevant observations: first, that success depends on the details and second, that scale and transparency are competing requirements that require further attention.

2.3.3 Open and closed systems

One of the discussions within the Prestige project has been in terms of the relative merits of open and closed systems. We will not go into the discussion of these as it is well outside the scope of this paper, but we should note that practical experience has shown that a truly open system in this case, while it would offer desirable features and certainly allow additional flexibility in selecting potential suppliers, it would nevertheless add considerably to the complexity of the required infrastructure thus increase the overall risks to the project. A project with similar scope by the Department for Trade and Industry for general use on the rail network developed using open standards has yet to produce any usable results and is now planned to become part of the Oyster project. The solution chosen by TfL in Prestige appears to work in the current context but it is questionable whether it would be applicable in other situations. Yet, the solution implemented by the PFI suppliers is not completely closed: indeed at lower network layers the selected technologies are ISO compatible and would work well with a variety of equipment from different suppliers.

Nevertheless, at higher levels this is not the case: the universally unique identifier and addressing scheme which is the key for successful operation of the system is proprietary and remains inaccessible and undocumented. At the present time this may not raise significant issues, in the longer run it may place significant limitations in the wider usability or the interoperability of the Oyster card with other fare management systems. Seen in the light of the current inequalities caused by the IP address allocation scheme it is reasonable to assume that this could potentially become a challenge in the longer term.

2.3.4 Other issues

A final lesson which relates to the business case of RFID is that the investment required is counterbalanced from the expected returns. In Prestige, the total cost of the project can be accommodated within the scope of business modernization and thus allows for a reasonable business case to be developed. It is hard to imagine that such a system would be selected in the absence of a strong business case.

Finally, it is worth noting that one of the ingredients for success has been the defensive and cautious approach to the introduction of new functionality. While the implementation of RFID opened up a variety of opportunities for consumer services, the actual functionality of the system has been restricted to the essential features that fulfilled the requirements of the business case on which the system was developed. These choices allowed for technical and organizational issues to be resolved early on and in most cases before systems were made available to the public which led to a successful implementation. That is not to say that Prestige has not been aware of these opportunities or that it avoided to address them. It is primarily the case that such experimentation has been delayed until the core system has matured and after the initial risky phases have been completed. Future developments will explore system management opportunities for example system real-time load forecasting, based on the detailed data captured at the gates rather than simulations. On the consumer front, the investigation of the use of the Oyster card as a potential mobile e-payment method will be explored.

3 Mitsukoshi case study

The second real world RFID system we discuss in this paper is developed in the context of retailing by Mitsukoshi, Ltd. at their flagship department store in Tokyo, Japan. Mitsukoshi is a historical retailer of some size and is responsible for the introduction of the concept of the department store in Japan approximately 100 years ago. In 2005, they became one of the first companies to implement item-level RFID tracking across their apparel sales floors. Unlike current research in the use of RFID in retail, this case refers to a fully commercial implementation driven by a specific business rationale. In what follows, we first discuss solutions to technical problems including auto identification, effective RFID reading range, system throughput and reliability, ease of use, and privacy. We conclude with a review of lessons learnt from this case study.

3.1 Business and service considerations

Mituskoshi was founded in 1673 under the name Echigoya, as a Tokyo-based drapery firm. In 1904 the company opened Japan's first department store and currently operates 18 stores in Japan and 24 in other countries including China, Germany, the US, England, France, Italy, and Spain [8]. Mitsukoshi's flagship store in Nihonbashi, Tokyo is Japan's largest department store in terms of revenues.

Japanese department stores have unique strengths and service offerings. While revenue per square foot is higher compared to similar stores in other countries, it depends on expensive customer service employing a high number of sales staff. To be sure, decreasing the cost of service by reducing the number of salespeople without affecting revenues is highly desirable and would result in a considerable competitive advantage. In RFID, Mitsukoshi discovered a potential technological solution to achieve this goal and as a result instigated an investigation into this technology in 2002. This initial interest in RFID eventually led to two major pilot implementations for women's shoes between October 2004 and February 2005, and a subsequent full-scale deployment in April 2005. Following the success of this work, Mitsukoshi is extending the use of RFID to different types of services and apparel items. In addition to the cost cutting implications of this project, there has been an additional motivating factor for the success of such a project in that 2004 was the centenary of the company's first department store and as a result business units were encouraged to do ''something new'' to re-invent retail shopping.

To this end, one of the initiatives undertaken has been the end-to-end implementation of Electronic Data Interchange (EDI) and Supply Chain Management (SCM) systems using barcode with a view to ''streamline the work processes from placing orders to buying in, [which] only reduces costs and does not increase revenues'' according to Masakazu Nishida, head of Mitsukoshi's RFID project. To extend the benefits of this infrastructure and translate it into sales, RFID in all items in the sales floors was seen as the core enabling technology. Indeed, department stores sell a large variety of different items but in small quantities each, which implies in frequent out-of-stock situations and thus loss of sales. While barcode-based systems improve this situation and help manage the supply chain, out of stock conditions cannot be easily avoided because the information of items with low stock is input to the system relatively late in the replenishment process. Mitsukoshi estimated a loss of 4–8% of sales for women's apparel products and 16% for women's shoes due to regular out of stock conditions.

Moreover, sales staff spend a considerable proportion of their time outside the sales area as they often needed to visit the stockroom to check for product availability. This is often the case for shoes in particular as only a small number of sizes and colors can be stored in the physically limited retail space. So, when a customer would request a size not available on the floor, the salesperson would need to walk to the stockroom and check manually while the customer waited idle. This process could be iterated several times until the customer would make a decision and either proceeds with a purchase or not. The inherent inefficiency in this time consuming process combined with frequent out of stock conditions inevitably leads to reduced customer throughput and potentially frustrating experience. Customers may also feel less inclined to make multiple attempts to find a suitable pair of shoes so as not to force sales staff to visit the stockroom more than two or three times at most. This psychological factor places a clear constraint on the negotiation between customer and salesperson.

Last but not least, Mitsukoshi's efforts must be seen within the context of a project of national reach driven by Ministry of Economy Trade and Industry (METI) and initiated in 2003 with a view to support a variety of RFID pilot projects including retailing [6, 7]. Mitsukoshi in association with other retailers, sector associations and information system suppliers participated in this work via a number of joint RFID pilots.

3.2 Balancing requirements and capabilities

Two rounds of prototyping and testing preceded the introduction in April 2005 of full RFID-based retail management at the women shoes floor in the flagship Mitsukoshi store at Nihonbashi, Tokyo. The first pilot run for 2 months starting October 2004 and the second for 2 weeks during February 2005, both aiming to establish and quantify the impact of the system on customer service levels and salesperson efficiency. The full implementation followed soon after the end of the second pilot, making it one of the first major consumer-facing deployments of item-level RFID tagging in a department store. At about the same time, Hankyu Department Stores, Inc., joined the competition and deployed a similar system at a women shoes shop in their main store in Osaka, Japan.

The actual deployed system includes three applications targeting consumers, sales staff and store managers respectively. The first application allows customers to easily check the availability of a particular product using a kiosk terminal (cf. Fig. 3). The salesperson application is used to support efficient checks of the inventory and to locate particular items using a portable terminal. Finally, a store management application allows inventory maintenance and accurate recording of sales. Passive RFID tags operating in the 13.56 MHz (HF) band are attached to each shoe box by the wholesaler storing Electronic Product Codes following the EPCglobal standard [2] so as to uniquely identify each individual pair of shoes. When a particular pair is removed from its box for display in the storefront, the RFID tag is also removed from the packaging and affixed directly onto the shoes. The six salespeople normally working in the storefront use personal digital assistants (PDAs) and a kiosk terminal, which are fitted with RFID readers. Network based EPC address resolution services, for example EPCglobal's Object Naming System (ONS), are not used in this implementation but may be incorporated in a future system upgrade.

The kiosk terminal incorporates a computer and an RFID reader, hidden inside its cabinet (cf. Fig. 3). The PDA used by sales staff has wireless and RFID reader capabilities. A database server is used to maintain RFID data and inventory information and is accessible to both types of devices over the network, which is separate from the network used to connect Point of Sale (POS) equipment and transfer EDI data. This separation is required by information security considerations although it introduces a degree of inefficiency. This system design has been deployed in six stores in separate systems that do not employ a distributed database. The friction between network security and ease of use has also been a central issue in further pilots of a similar RFID-based system at the women denim floor in February 2006 as will be further discussed below.

Fig. 3 Using a kiosk terminal in the shoe floor for women: When a shoe is placed on the table, the system automatically reads the RFID tag that is attached to it and shows relevant stock levels

The early business cases for RFID have been developed on the premise that the technology would enable completely automatic stockroom monitoring. However, the effective reading range of the RFID tags in practice has been only a few centimeters, which makes it unlikely the efficient implementation of such automation as storerooms occupy substantial physical space which requires a very high number of readers. Furthermore, it has not been possible to employ RFID tags with longer range because of regulatory and cultural constraints: RFID tags with larger antennae could provide the required range but it is important for shoppers that tags are small, certainly smaller than the existing price tags. On the regulatory front, at the time that this system was built passive RFID tags that operate in the 900 Hz (UHF) band and provide a suitable range were not permitted under Japanese Radio Law. This limitation was subsequently lifted and Mitsukoshi will likely consider the use of UHF tags in the future, especially Gen2 EPC compliant tags.

In order to palliate the effects of the limited reading range and still address the requirements of the business case, Mitsukoshi opted for a modified business processes that involved some degree of manual intervention: RFID tags attached on shoe boxes are read at the time they are shipped by the wholesaler. When received, shoe boxes are placed in the stockroom and the system bulk updated with the data supplier by the wholesaler. When a pair of shoes is sold at a later time, the salesperson involved in the transaction presents the tag to a reader at the POS which updates the system with the sale event. Inventory maintenance is then based on the data collected in this way at the POS.

Another implication of the limited range of the tags is that first they have to be attached on the outside of the shoe box since the box is much larger than the tags' reading range. Then, when a pair of shoes is removed from their box and displayed, the RFID tag should also be removed from the box and re-attached on the shoes. To facilitate this manual process, RFID tags are placed into plastic envelopes that are attached onto shoe boxes. RFID tags also have hooks so that they can be easily attached to shoes using the string of the price label. In practice, this process has to be carried out in a few cases only as the pairs of shoes that are displayed at the storefront are generally sold last. When such a pair is sold, sales staff simply remove the RFID tag from the shoes, present it to the reader, and retain the tags for reuse.

Given the actual limitations of the technology it may seem that it offers only few extra benefits compared to barcodes that could also be used to support a similar scenario. Nevertheless, even the relatively less significant advantage of RFID to not require line of sight appears to have positive consequences and reduce the burden on sales staff and customers alike. When customers want to check the inventory of certain shoe, they only need to place it on the kiosk cabinet. Nevertheless, achieving this apparently spontaneous reaction requires careful system tuning and a highly reactive system that integrates visual and auditory feedback.

Use of RFID can also help avoid some common sources of inventory inaccuracy. For example, RFID-based systems can easily detect double (or multiple) reads as each tag has a unique identifier and the system will not confuse duplicate reads of the same identification number as two separate boxes. In contrast, barcodes only encode information about manufacturers and product types only and cannot distinguish individual items.

3.3 Lessons learnt

Two pilots conducted in store with customers gave early indications that the system would not actually be used in the way intended initially. Testing uncovered mismatches between system features and business practices, and also played an important role in the system adoption process. During the period between the pilot end and full deployment, the system was further modified so as to better meet the observed behaviors and use. This deployment led to a substantial increase in sales which was attributed directly to the new system. As a result, the RFID approach was deemed a success and work is underway to extend its reach across all apparel sectors.

3.3.1 User interaction

Early on in the pilots it was observed that both the PDA and the kiosk terminal applications were not used as intended. In particular, PDA-based applications designed for sales staff were rarely used in practice, and the kiosk which was thought of as a self-service point for customers did not attract any interest at all. Moreover, sales staff considered the kiosk of particular benefit to them and so they abandoned the mobile terminals.

The PDA-based application had a number of usability problems and did not facilitate the type of tasks a salesperson must perform. The devices are too big to fit in the suit pocket (sales staff are rather formally dressed, cf. Fig. 3) and require use of both hands. The limited battery life also created problems as the salesperson was required to set the device in standby mode often to prevent it from expiring. The PDAs employed require 11 s to wake up from standby and successfully establish a wireless connection, which is too long to wait when a customer is engaged and caused an awkward situation. PDAs were also too sensitive for use on the sales floor and indeed during trial two of the seven devices were involved in accidents and stopped functioning. The option to use ruggedized devices is not appropriate as they are too heavy and cumbersome in this context and do not conform to the sense of elegance that must be conveyed by sales staff. The kiosk provided an alternative for checking the stock and soon sales staff came to realize the benefits of using it, especially during customer interaction and a sales pitch.

While the kiosk application was very well developed and easy to use, consumers nevertheless avoided using it themselves. This does not have to do with narrow considerations of usability as this was clearly not a problem, but rather had to do with their perception of the type of activity they were involved in: ‘‘All in all, they come to buy products. Using a personal computer at a place surrounded by products is rather... Well, [customers are] not in such a state.’’ Instead, customers became aware of the new capability and would ask more frequently the sales staff about availability of items and sizes. Using the kiosk, sales staff are able to check inventory 3.1 times per customer on average which is a considerable increase compared against the 1.7 times per customer on average that they would visit the stockroom before the introduction of the system. This increase has direct impact on improving the ‘‘hit rate’’ and thus more sales. The length of a session with a customer however has remained the same, approximately 12 min. As a result of the improved rate of closed sales, staff developed positive attitudes and encouraged the full deployment of the RFID system. This is despite the fact that at the beginning of the trials consultation with staff indicated that they did not expect the system to have a particular positive impact in their work: “Before the introduction [of the system], sales staff thought it was impossible to serve customers using such a system. They thought running to the stockroom makes customer service pleasant.”

They also expected customers to ask them to check the stockroom anyway, something that was not observed at all. This expectation was based on the fact that the normal operating mode at the time was for sales staff to use a pen and paper system to track of out-of-stock items. This system was very prone to inaccuracies and so customers, based on their experience, tended to doubt the information held there and frequently request that stock was double-checked anyway. Therefore, sales staff expected that the new system would still cause the same response to consumers and would be unable to remove the overhead of traveling to the stockroom. Early on during the pilots, it became evident that customers trust the information on the computer screen more than that recorded in the notebook and they did not ask for further checks. Moreover, because of the speed of inventory checks it became possible to embed interactive information retrieval sessions during customer sessions held in front of the kiosk. Sales staff came to believe that such interactivity better facilitates and indeed accelerates customer interactions. Thus, sales staff embraced the new system with enthusiasm and supported its deployment with urgency.

To be sure, initial system designs lacked features that subsequently were seen as core functionality to support effectively a variety of business processes and different customer profiles. For example, an early version of the system could not support layaway and back-ordered items that are both significant components of the department store service offering. Marking a product as layaway or back-ordered implies that this particular item must not be sold to another customer. The interim solution developed by sales staff was to record “dummy” sales for the requested items, a fact that was confusing to the system and others. Later versions provided mechanisms for handling such actions explicitly. A second limitation discovered through trials was that the system did not support a particular kind of search that is very important for a specific type of customer, namely those looking for very small or very large sizes. Such products are rarely available for the majority of shoe types, and so it becomes particularly difficult to assist such customers and even then sessions can be particularly protracted. Another mechanism incorporated before full system deployment allows these customers to specify their size and retrieve all available shoes.

The introduction of RFID had a clearly positive effect on the performance of the store as revenues increased by between 10 and 20%, a fact that can be directly attributed to the new system and the efficiencies it facilitates both in terms of service provision and product replenishment. Based on this experience, Mitsukoshi is expanding RFID deployment across stores and services. In February 2006, they concluded a pilot for designer denim which is similar to the shoe case in that products are relatively expensive and have a wide variety of sizes and designs. This pilot did not only replicate the experience of the shoes shop but introduce a number of new applications including smart shelves that display real-time inventory updates (Fig. 4); intelligent fitting rooms providing handheld RFID readers that allow in situ inventory checks; and so-called concierge services that employ active RFID tags to identify customers and to provide personalized customer services.

Sales staff were initially similarly skeptical about the utility of the system since contrary to the shoe shop customers do not ask for stock checks as frequently. Moreover, the installation of electronic displays on the shelves raised concerns about negative effects to the store ambience which is a critical factor in attracting customers. The pilot changed these perceptions and sales staff now consider the positive aspects to outweigh any potentially negatives. This change in attitude was primarily due to two factors: considerable improvements in the efficiency of replenishing shelves and second, because the smart shelves allow staff to

Fig. 4 Smart shelf showing the stock status on a color electronic paper display. RFID readers are embedded in the shelf to monitor the presence of jeans. *White numbers* indicate sizes that are available on the shelf. *Green numbers* indicate sizes that are not available on the shelf but available in the stockroom. Sizes that are unavailable are not shown

easily establish what products are available either on the self or the stockroom and thus allowed them to recommend alternatives with very little effort.

Concierge services have been the most challenging area of experimentation and have received mixed reviews. In pilots, active RFID tags were used to track 50 customers in their visits to the store. Staff used 13 stations installed at various locations within the department store to track customers. Also, when customers required assistance they could push a button on their tag to notify a salesperson who would respond. Tags were also used to identify the location of the customer and send SMS messages to their mobile phone with related information. Nevertheless, this use of RFID introduces complex privacy issues which are hard to deal with. This experiment should be seen as an early attempt to explore such ideas in a limited way. Consumers involved volunteered to use the service which was designed to adhere to the Japanese RFID privacy guidelines [7] and the tags were returned when customers left the store. This pilot produced mixed results and it is expected that considerable further development will be required to assess the feasibility of this approach.

4 Lessons from the real world

We conclude the discussion of real-world deployments of ubiquitous computing by attempting to identify common threads and considerations between the systems discussed. We find that there are certain similarities in the development and experience of the two projects which have significant implications for the success of ubiquitous computing in the real world. Nevertheless, we must question whether these issues are uniquely related to the characteristics of ubiquitous computing or whether they are common considerations for any information technology when facing the realities of the world at large.

The first commonality between the two projects is the fact that considerable organizational changes had to be put into place before it was possible to even consider the use of RFID. Indeed, RFID implementation leads to significant increases in recorded information that becomes available within the business. While at first sight this appears to be a positive development which can be turned into competitive advantage, it also places considerable burden on the business information systems which have to be modified and extended so as to cater for the requirements of the new data streams that must be managed. Appropriate preparations including significant investments in new infrastructure is then a necessary (and costly) first step which has to be completed well before any actions are taken towards RFID implementation. Failure to do so, cancels out any potential benefit, as recorded data will not be turned into useful information that can provide concrete business benefits. In fact, preparing the “traditional” IT infrastructure for RFID can represent the majority of the cost associated with deployment of RFID.

However, where organizational constraints appear to play an even more central role is in preparing staff to deal with the new situation created by the introduction of RFID. Indeed, the availability of RFID forces a fundamental transformation on the way service offerings are presented and supported, a fact that demands changes in business processes across all levels of the organization. Clearly the most visible impact of this is in the client facing aspects of service provision where RFID and associated embedded devices are directly involved and clearly visible. In both cases examined, the design of client service space (underground stations and retail space correspondingly) had to be changed and new equipment introduced, and as a result these modifications are hard not to notice. Nevertheless, perhaps the more complex changes were those that affect supporting business processes. As discussed in detail in previous sections these primarily relate to organizational learning and exploiting the new data for better planning and increased operational efficiency.

Another lesson that appears to be shared between the case studies is that of controlled introduction of new features and innovations. While from a purely technical perspective there are no significant challenges in completely reinventing the user experience, in practice this would imply a considerable risk for business as clients would be confronted by an unfamiliar situation. Instead, in both cases the new systems have been first developed as optional additions to current practice, then acquired features that require support by staff, and finally developed into independent client activities (or are expected to follow this route in the case of Mitsukoshi). This *slow growth* approach allows for the management of the unexpected behaviors and practices that could not have been predicted in a laboratory setting and required piloting in the actual operational environment. This is because such behaviors come as a consequence of user context the expression of which is clearly restricted in artificially set up situations. Indeed, since ubiquitous computing affects the everyday, it is reasonable to expect that a considerably larger variety of behaviors would appear due to the freedom allowed in the use of the system. Dealing with the consequences of this diversity requires smaller or more extensive changes which can potentially overwhelm the system designers if they are to be implemented concurrently. A phased approach allows for such changes to be addressed gracefully.

Observing unexpected or unintended user behavior is of course not a unique feature of large scale ubiquitous computing installations. Of course, every new computing technology creates such unintended uses but the situated nature of ubiquitous computing opens up many more opportunities to do so. This fact has been increasingly recognized in research but the experience of the projects discusses here raises further questions regarding the potential expectations from small-scale studies which are prevalent in the field. Generalizing findings related to attitudes and observed behaviors when they are based on a few prototyping iterations, an artificial setting and a limited user population, do not appear to provide adequate descriptions of the situation at hand and are unlikely to provide any long term solutions on note. Indeed, both cases indicate that attitudes and behaviors evolve in time and while living with such systems, and user understanding and attitudes adapt as the system becomes part of the everyday. Clearly this is a challenge for research as such work has to be accommodated within shorter timescales that may not be able to capture the characteristics of the evolving relationship between users and system.

Such a cautious process of iterative refinement with incremental adaptations seems to be more successful in practice. This approach can be seen as to somewhat temper visions of the disruptive nature of ubiquitous computing with respect to previous paradigms. Instead, ubiquitous computing is seen to depend on and to build upon previous generations and to evolve naturally—effectively transforming all computing into an element of ubiquitous computing. As the novelty factor wears off, ubiquitous computing systems appear to be appraised on the basis of their functionality and actual benefits rather than on the basis of excitement towards what is new and different.

Of course, real world applications have to deal with the realities and the mundane details of practical matters: one such issue that stands out is the economic and business feasibility of any such system. What has been a welcome finding is that the costs associated not only with the actual technology investment but also all related organizational and business costs can already make such early ubiquitous computing systems viable. In fact, the cost of RFID deployment itself seems to be only a small proportion of the overall technology cost, with the best part of the overhead relating to the modifications to physical places and objects that have to be made to accommodate its use. It is not the RFID technology itself that costs but rather embedding it into objects and artifacts.

Nevertheless, when looking closer at the financing of such projects one issue that appears to stand out and have considerable implications for research is the argument between open and proprietary systems. Open systems are seen by many as the enabler of shared, global infrastructures which can trigger the shift to wider popularity for ubiquitous computing, much in the spirit of the changes brought about by the internet. If ubiquitous computing is to be open to all for participation and to gain global reach, then clearly, membership must be based on such open specifications that are part of a well regulated and equitable management system.

To be sure, the two projects reviewed here have made different choices in this respect. These choices are not arbitrary nor they represent a prejudiced attitude of the organizations involved, but rather reflect the economic realities and perceived benefits and costs of the available solutions and their operating context. In the case of Mitsukoshi, the emerging EPCglobal standards are selected as the foundation for their system. These are open specifications, managed and developed by a long-established industry association with global reach, which furthermore is a collaborative venture between market sectors and nations. A perceived critical benefit of the use of this system is its support by many different vendors but most importantly the need to automate the exchange of business information across trading partners along the supply chain. This later issue is a critical one and has also played a central role in the development of the internet. Indeed, the cost saving or business incentives gained from sharing or exchanging information appears to be as important for ubiquitous computing as it has been for the internet and e-business.

On the other hand, closed and proprietary systems have considerable potential performance and cost advantages that can make them attractive candidates in many situations. For example, if there is no need for coordinating or federating trust domains it makes sense to use a closed system using symmetric cryptography rather than deploying a managing a full public key infrastructure. In addition to having a simpler architecture and management overhead this approach would also have considerably lower computational requirements which can be a critical advantage taking into account the resource constraints of RFID tags. Since there is no requirement for exchanging or sharing information with external partners, there are few if any negative aspects in such a choice.

Clearly this situation has significant repercussions for the future development of ubiquitous computing as in case the later approach becomes dominant the ubiquitous computing reality will be one of isolated islands of functionality rather than the seamless calm computing vision advocated by academic research.

5 Conclusions

Ubiquitous computing in the real world has concerns that are rarely dealt with in research. Lengthy and costly preparation or upgrade of existing infrastructures; training of employees and users in the new ways of working; controlled introduction of new functionality; features and services to manage risk; unexpected behaviors due to the wider variety of possible real world situations; incremental approach to systems development so as to better identify successful aspects; regard for the economics of systems as a core requirement; and selection of open or closed systems, are all issues that are mostly outside the scope of current ubiquitous computing research, but seem to play a critical role in both case studies we consider here. On the other hand, current state of the art in research does not appear to have been taken into account even in areas that it could have been of considerable value. Indeed, this discussion clearly identifies a gap between practice and the state of the art in research.

To be sure, respecting established practice and avoiding disruption in service provision is paramount for real world deployment. Nevertheless, the question of how to make the best of innovations in either development and design methodologies or core infrastructure technologies should be of critically importance considering the complexity of the challenges involved. For example, as regards the question of identifying unexpected behaviors early on in the design cycle specifically, it should be noted that there are currently several design approaches proposed in research—such as end-user development/design and participatory design—which seem to hold considerable promise. But these do not seem to have been employed in the cases considered here at all. This can potentially have adverse results as, if incremental design is the only approach used to deal with ubiquitous computing systems design in the real world, it would limit the opportunities provided by the technology and potentially fail to deliver the promise of the ubiquitous computing vision.

Therein lies the challenge for ubiquitous computing research: how not only to learn about the concerns of those developing systems in the real world but more importantly, how to translate principles, guidelines and models discovered in the context of research into useful tools for building ubiquitous computing systems in the real world. Bringing the two communities closer together and communicating lessons learnt in ubiquitous computing research so as to inform practical system design and development can have profound implications for the success or the failure of the ubiquitous computing vision.

Acknowledgments We would like to thank Brian Dobson of the Prestige project at Transport for London and Masakazu Nishida of Mitsukoshi, Ltd. for providing access to the projects and for several discussions that provided a wealth of information. We would also like to thank Mikako Ogawa, Masaki Umejima, and Jiro Kokuryo of Keio University for providing useful information and generous assistance in setting up the Mitsukoshi case study.

References

1. Alessandro A (2004) Security of personal information and privacy: technological solutions and economic incentives. In: Camp J, Lewis R (eds) The economics of information security, Kluwer, Dordrecht
2. EPCglobal (2006) EPCglobal Inc. Home Page. http://www.epcglobalinc.org/
3. Günther O, Spiekermann S (2005) RFID and the perception of control: the consumer's view. Commun ACM 48(9):73–76
4. Group Transport Planning and Policy (2005) London Travel Report 2005, Transport for London
5. Lingo L, Orrick P, Smith L (2005) Fare payment technologies, California Center for Innovative Transportation UC Berkeley
6. METI (2006) Overview of METI RFID Tag popularization policy. http://www.meti.go.jp/english/information/data/IT-policy/ic-tag.htm
7. MIC and METI (2004) Guidelines for privacy protection with regard to RFID tags. http://www.meti.go.jp/english/information/data/IT-policy/ic-tag.htm
8. Mitsukoshi (2005) Mitsukoshi home page. http://www.mitsukoshi.co.jp/
9. Roussos G, Moussouri T (2004) Consumer perceptions of privacy, security and trust in ubiquitous commerce. Personal Ubiquitous Comput 8(6):416–429