

Contextual software

Part 1: Track Record

The ubiquitous computing (ubicmp) group at Glasgow has a six-year track record of innovative research in mobile computing, human-computer interaction (HCI) and computer supported collaborative work (CSCW). It has built up an international reputation for innovation through a holistic approach spanning theory, studies of systems in use, and systems design that advances both infrastructure and interaction issues. It shows how ubicmp technologies can be designed for and used in real settings. This group was founded as part of the Equator interdisciplinary research collaboration, and is set within the Computing Science department's GIST group, one of the UK's strongest HCI groups. The ubicmp group's work has pioneered systems for cultural tourism and city visiting [1, 2], fitness and sport [3], and mobile multiplayer games [4, 5, 6]. We treat these as application areas in themselves, and also as vehicles for new design concepts and technology, i.e. rules, goals and interactions rely on features of our new technology and of users' context and environment. Systems research has covered mobile ad hoc networks, peer-to-peer data sharing, recommender systems and information visualisation. Recently, our infrastructure focus has been on *Domino* [6], using mobile ad hoc networks to distribute usage histories which describe users' activity and support adaptation of assemblies of software components. Visualisation work has focused on *Replayer* [8], a collaborative distributed system that lets one or more people combine video, audio, field notes and system logs when analysing system trials. The group's conceptual work has fed into and been fed by user experience and system work. It deepened the understanding of the importance of 'conversational resources' in online social interactions [2, 9] and how ubicmp systems can become a part of everyday life [5]. It advanced *seamful design* [9,10], in which limits and gaps in infrastructure are made into system design resources, and the treatment of history as part of context [12]. We showed that ubicmp's design ideal of *invisibility* is unsustainable without occasional periods of visibility [10,11].

Matthew Chalmers, principal investigator, is a Reader and leads the group. He has been involved in ubicmp since 1990, working at Xerox' labs on information visualisation and the first ubicmp systems: active badges. He founded and ran a visualisation group at UBS Ubilab before returning to ubicmp as PI in Equator and an eSocial Science project on visualisations for analysing ubicmp systems' use. In Equator, he led the *City* project, and jointly led the *Seamful Games* project and Equator's theoretical work. He is PI on the EPSRC *Designing for the Augmented Stadium* project (Oct 2007-Mar 2010) using MANETs for crowd interaction at sporting events. He has served on the committees of CHI, Pervasive, Ubicmp and PerCom, is on the editorial board of Pervasive and Mobile Computing, was a VIP Guest at the Microsoft TechFest, and recently gave the closing keynote for the francophone ubicmp conference, Ubimob. He co-authored the UKCRC Ubicmp Grand Challenge, and is funded to promote it in an EPSRC *Collaborating for Success through People* grant (EP/F013442/1, Oct 2007-Mar 2009). His interdisciplinary work has involved collaboration with sociologists, artists, art theorists, philosophers and museum specialists. He was a Projektleiter in the Swiss *Perform Space* project, exploring the theory and practice of performance art, city spaces and new media, and was a judge at the VIPER festival for film, video and new media, in Basel. He maintains a leading role in data visualisation, serving on the editorial boards for the Information Visualization journal and Springer's Information Visualization book series.

Philip Gray, co-investigator, is a Senior Lecturer and has been actively engaged since 1984 in research into models, notations, software architectures and tools for user interface development. Recently, he has focussed on the description and engineering of interaction techniques for mobile and ubiquitous systems, addressing data of variable quality [15]. In this context he is currently PI at Glasgow for the MATCH Project (SFC, 2005-2009) with responsibility for adaptive interaction techniques for homecare systems, and also Glasgow PI and overall Technical Manager for the OpenInterface Project (EU FP6, 2006-2009) developing tool support for the prototyping and evaluation of multimodal user interfaces. Previously, he has been investigator on more than ten externally funded research projects. He is a member of IFIP Technical Committee 13 (Human Computer Interaction) and vice-chair of IFIP Working Group 2.7/13.4 (User Interface Engineering). He has been conference chair for HCI '94, the annual UK HCI conference and for IHM-HCI 2001, the first joint Anglo-French conference in HCI. He has about 80 publications,.

Alistair Morrison is a Research Assistant in Computing Science, and a specialist in information visualisation. In his PhD he developed some of the fastest and most accurate non-linear multidimensional scaling algorithms [7]. He has recently worked on the *Replayer* toolkit, visualising ubicmp log data [8]. Morrison developed dynamic visualisations that let evaluators find video clips of selected participants automatically, using logged locations of trial participants and evaluators. He has served as a reviewer for IEEE Information Visualization, Information Visualization, IEEE Transactions on Vision and Computer Graphics and ACM CHI, ACM UIST, Eurovis, Neurocomputing, IJHCS and Information Visualisation (UK).

Malcolm Hall is a PhD Student writing up his thesis in Computing Science. Hall is a specialist in software and networking infrastructure for ubicmp. His initial PhD work was on infrastructure and application design in systems such as *Treasure* [4] and *Feeding Yoshi* [5]. Most recently, he has drawn inspiration from the everyday

social practices of using and adapting software in developing Domino [6] which, like Replayer, is one of the main foci of the proposed research. He has also developed highly successful mobile applications, such as Navizon, WiFiFoFum and Dynamo, which have over a million downloads. He has served as a reviewer for ACM CHI, Ubicomp and IEEE Pervasive.

Partners

Prof. Jon Crowcroft (U. Cambridge) is Marconi Professor of Communications Systems, and a fellow of the ACM, the British Computer Society, the IE[ET], the Royal Academy of Engineering and the IEEE. Prior to that he was professor of networked systems at UCL. He was a member of the IAB, general chair for the ACM SIGCOMM 95-99, on the editorial team for COMNET, and on the program committee for ACM SIGCOMM and IEEE Infocomm. He has published 5 books, the latest being *The Linux TCP/IP Implementation* (Wiley, 2001). He is the PI for the CMI funded Communications Research Network, a £3M government/industry/academia multidisciplinary collaboration to automate the successful exploitation of disruptive communications technologies. He is a leading member of the EU Hagggle project (www.hagggleproject.org), creating middleware for fixed and mobile ad hoc networks. In *Contextual Software*, he aims to extend his middleware work out of closed settings into applications used in everyday life, focusing on trust architectures, and experiments that support understanding and design responses to people's mobility, social structure and technology use. As the attached letter of support shows, Cambridge is not asking for funding but has offered Glasgow formal collaboration as a partner that can feed design requirements into the ongoing Hagggle EU project. Crowcroft will also actively encourage self-funded 2007- and 2008-intake Cambridge PhD students to work on topics bridging between Hagggle and Glasgow, such as trust architectures and 'cheat-proofing', and using tools such as the HagggleSim emulator on our system log data so as to quantitatively assess low-level networking design choices.

Dr. Barry Brown starts as Associate Professor at U.C. San Diego in Sept. 07, having left the Glasgow ubicomp group in June 07. He combines the social and computing sciences, pioneering the serious study of leisure and enjoyment, and examining existing leisure practice, new technologies for leisure, and trials of systems in real settings. He has organised 6 international workshops, published 2 edited books, served on the committees of 4 conferences (including ACM CHI). He is a Visiting Scholar for *Habitable Cars* (ESRC RES-000230758), with 2 months per year spent in Scotland over the next 2 years—which he will extend to study this project's practices of evaluation and design. He was a visiting scholar at the University of Saskatchewan, the Interactive Institute in Stockholm and the Digital World Research Centre at U. Surrey. He was a keynote speaker at the Socio-Technical Systems conference (2004, Edinburgh) and Virtual Conference (2005, Stockholm).

Selected Publications

1. B. Brown, I. MacColl, M. Chalmers, A. Galani, C. Randell, A. Steed (2003) *Lessons from the lighthouse: Collaboration in a shared mixed reality system*, Proc. ACM CHI, Fort Lauderdale, pp. 577-585.
2. B. Brown, M. Chalmers, M. Bell, I. MacColl, M. Hall, P. Rudman (2005), *Sharing the square: collaborative leisure in the city streets*, Proc. ECSCW, Helsinki, pp. 427-447.
3. I. Anderson, J. Maitland, S. Sherwood, L. Barkhuus, M. Chalmers, M. Hall, B. Brown, H. Muller (2007), *Shakra: Tracking and sharing daily activity levels with unaugmented mobile phones*, To appear in ACM Mobile Networks and Applications.
4. L. Barkhuus, M. Chalmers, P. Tennent, M. Hall, M Bell, S. Sherwood, B. Brown (2005) *Picking pockets on the lawn: The development of tactics and strategies in a mobile game*, Proc. Ubicomp, Tokyo, pp. 358-374.
5. L. Barkhuus, M. Hall, S. Sherwood, P. Tennent, B. Brown, D. Rowland, S. Benford, M. Capra, A. Hampshire (2006), *Interweaving Mobile Games with Everyday Life*, Proc. ACM CHI, pp. 417-426.
6. M. Bell, M. Hall, M. Chalmers, P. Gray, B. Brown (2006), *Domino: Exploring Mobile Collaborative Software Adaptation*, Proc. Pervasive, Dublin, pp. 153-168.
7. Morrison A, Chalmers M, *Improving Hybrid MDS with Pivot-Based Searching*, Proc. IEEE Information Visualization (InfoVis) 2003, Seattle, pp. 85-90.
8. A. Morrison, P. Tennent, M. Chalmers, J. Williamson (2007), *Using Location and Motion Data to Filter System Logs*. Proc. Pervasive, Toronto, pp. 109-126.
9. B. Brown, M. Bell (2004) *Social interaction in 'There'*, Proc. ACM CHI, pp. 350-359.
10. M. Chalmers, I. MacColl and M. Bell et al. (2003) *Seamful Design: Showing the Seams in Wearable Computing*, Proc. IEE Eurowearable, Birmingham, pp. 11-17.
11. M. Chalmers and A. Galani (2004) *Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems*, Proc. ACM DIS, Cambridge MA, pp. 243-252.
12. M. Chalmers (2004), *A Historical View of Context*, J. CSCW 13(3), pp. 223-247.
13. A. Crabtree, S. Benford, C. Greenhalgh, P. Tennent, M. Chalmers, B. Brown (2006) *Supporting Ethnographic Studies of Ubiquitous Computing in the Wild*, Proc. ACM DIS, 60-69.
14. M. Chalmers, *Informatics, Architecture and Language*. In: *Social Navigation in Information Space*, A. Munro, K. Hook & D. Benyon (eds.), Springer, 1999.
15. P. Gray, D. Salber (2001) *Modelling and Using Sensed Context Information in the Design of Interactive Applications*, M. Reed Little and L. Nigay (Eds.): EHCI 2001, LNCS 2254, pp. 317-335.

Part 2: Proposed Research and its Context

This project aims to make user experience, design and evaluation in ubiquitous computing (ubiquitous computing) more integrated and synergistic than before. Ubiquitous computing has emerged as a key area of ICT. It deals with systems that fit with user context and interaction, and takes a holistic view spanning technology, use and users, in which “the unit of design should be social people, in their environment, plus your device” [16]. Robin Milner, one of the authors of the UKCRC Ubiquitous Computing Grand Challenge, reflected on the ubiquitous computing ‘vision’ in [24] and advocated “exploratory projects that aim to define the kinds of experience that lie at the core of the vision. This requires experiments that create specific socio-technical environments and ask humans to enter them. [...] Here we look for synergy between the societal vision on the one hand, and the development of scientific models and engineering principles on the other.” Our proposed project is in accord with this, in that it supports and connects users, evaluators and developers engaged in the process of creating ubiquitous computing systems of sustained contextual utility and value, i.e., **contextual software**. Understanding and improving this iterative socio-technical process is vitally important for ubiquitous computing because this broader process is what creates and sustains’ systems contextual fit. The central models underlying the proposed work are richer usage histories than have been used before, combining system logs of software components’ installation and combination, locations visited and other tracked features, and video and audio recordings made by evaluators and users. Our system engineering principles focus on analysis and sharing of patterns within these histories. Our experiments with socio-technical environments are designed to address two **key problems** in ubiquitous computing: adaptation and evaluation.

Making systems more adaptable or adaptive is crucial to ubiquitous computing’s progress because contexts, needs and uses are often more dynamic, subtle and hard to predict than in other areas of ICT. Actual use of ubiquitous computing systems may differ from designers’ preconceptions when, for example, mobile users are interacting in the uncontrolled environment of city streets. Software based on such preconceptions may become increasingly unhelpful or inappropriate unless it adapts or is adapted with use. In addition, despite major advances in the Equator IRC (www.equator.ac.uk) and elsewhere, the problem of evaluation in ubiquitous computing is becoming critical. Those aiming to create new ubiquitous computing systems or adapt existing ones in a timely fashion need to understand users’ changing contexts and uses, but it is prohibitively difficult for evaluators to be with users all the time in their everyday lives. Observing and recording where they go, what they do, and their interaction with people nearby and—via networked mobile devices, for example—other people in quite different contexts [13].

The proposed research is timely, in that it addresses adaptation and evaluation in ubiquitous computing by building on tools and techniques that have only recently emerged from the Equator IRC. As our track record (given in Part 1) shows, we have the experience and ability to address adaptation and evaluation in ubiquitous computing in this holistic way. Our unique skill set, built up in Equator, spans the range of areas needed, including ubiquitous computing infrastructure [6, 15], application and user experience design [1-5], information visualisation [7,8], evaluation [13] and theory [11,12]. Here we take Equator work in a new direction, in ubiquitous computing systems that can be adapted within minutes or hours rather than the normal days or weeks. We will demonstrate the first example of dynamic adaptation of system structure involving new software tailored to a new context of use: a mobile phone user will be able to integrate and use a component written or modified for him/her, and tailored to the software he/she already has on the phone and to an urban area previously unknown to the evaluator/developer team—all within an hour of the user entering that area.

This is an advance beyond prior work in terms of pace of the loop of design, use and evaluation. Logged usage data will be streamed to evaluators in real-time, who work with developers to make software changes that are quickly distributed among users, whose use of the new software is logged... thus starting the loop again. The dynamics of context make such a tight feedback loop particularly relevant to ubiquitous computing, but we also aim to show whether this loop is productive and practical in terms of new software and new design practices, i.e. understanding the strengths and constraints on the evaluations and the software made within it. While evaluators and designers will develop new analyses and new software components in a more timely way, users are the final arbiters of contextual fit and value. Therefore the project also aims to understand and design for users adapting systems and use: choosing which new components to integrate and which old ones to halt or discard, and finding new ways to use their systems and new contexts their systems might be used in.

In practical terms, what binds these advances together is taking advantage of a drastically under-utilised resource: logs of system execution and use. Such usage histories afford powerful new models of context and functionality that support evaluation, design and use in the ways described above, especially when we exploit the *Haggle* mobile communications infrastructure (www.haggleproject.org) recently developed by Crowcroft et al., and three prototype systems, each piloted near to the end of Equator:

- *Replayer*, a set of visualisation tools specialised for ubiquitous computing evaluation [8], allowing user studies experts and developers to collaboratively analyse system log data and related audio and video recordings;
- *Domino*, a toolkit for distribution of new software components among mobile devices, and analysis of components’ usage histories that drives integration of components into applications [6]; and
- The *Castles* software recommender, tailored to *Domino* and based on patterns abstracted from usage histories. This lists new relevant components, so users can choose which to integrate into their running system [6].

These systems’ potential synergy stems from how they could help evaluators, designers and users interact, so as to create contextual software. This synergy will be achieved in this project through a number of advances: new

Replayer visualisations that handle streams of real-time data, rather than only static files, will let evaluators analyse and steer a user trial as it happens, determine new uses and contexts, and work with developers on design issues to be resolved or exploited; other new visualisations for developers, tightly integrated with everyday programming tools, will show components' distribution, execution and combination in use, supporting the timely creation of new Domino components; using Huggle, we will spread new components efficiently and securely—and comprehensibly to users—via fixed infrastructure and mobile ad hoc networks; new visualisations in Domino-based recommenders will let users better see new components' past and potential use, and help them judge whether and how to adapt their system, their use or their context of use; and Huggle will be used to feed ongoing logged usage data back to evaluators and developers either implicitly, via log data alone, or explicitly, i.e. usage data along with annotations and discussion. Secure and efficient interconnection among devices will support near to real-time communication between evaluators, developers and users, but it will also afford greater collaboration in their practices and roles. Understanding and designing for this social interaction forms another novel contribution of the project. We aim for a broad treatment of people's mobility, collaboration and use of networked mobile devices, based on a view of them as a dynamic whole, e.g. how individuals' mobility and interaction shape and are shaped by the system connectivity that comes with mobility, and how mobility, interaction and the system itself are adapted to suit patterns such as being central in one social group, or 'promiscuously' moving between groups and roles.

In summary, the project aims to improve the process of creating and sustaining contextually fitting software, changing not only applications and tools but also the practices of evaluators, designers and users. More specifically, we will answer a number of **key research questions** about evaluation, design and use in ubicomp, reflecting different aspects of this novel dynamism and holism:

RQ1: How can evaluators' analysis of ongoing context, operation and use help them run a user trial?

RQ2: How can developers' analysis of ongoing context, operation and use help them make new timely software?

RQ3: How can users' analysis of ongoing context, operation and use help them in adapting all three?

RQ4: How can robust, efficient and secure communication via fixed and MANET infrastructure support collaborative adaptation?

RQ5: How can timely and informed social interaction among users, evaluators and developers be used in creating and sustaining systems' contextual utility?

RQ6: What theoretical reflections and formal models clarify this dynamic set of people, software components and contexts, and generalise from it in the form of design guidelines and analytic frameworks?

Answers will be based on demonstrations of new applications, trialled with users in urban settings. We will build on our experience in making visualisations that afford powerful analysis, and creating ubicomp experiences that are not only engaging to users but which also serve as vehicles for new computer science techniques and concepts [1-6]. Our **research outputs** will include research papers that offer analytic findings and methods, design guidelines and formal models, as well our software and an evidential record of the user trials: a publicly accessible resource for others to examine and reuse, consisting of system usage histories, evaluators' notes, multimedia recordings, and interviews and diaries of users.

Background

A central but difficult issue for ubicomp system design is how to model the changing environment or context of each person using a system. Context and use are varied, dynamic and subjective, even though many systems rely solely on discrete geographical locations of users. The well-known paper by Dey et al. proposed that context includes "the location, identity and state of people, groups and computational and physical objects" [17]. This could include almost anything in human activity, but a system designer has to make some choice *a priori* as to what contextual features to model in the formal and objectifying medium of software. He or she has to define how to combine and relate such models, how to feed into them from sensors and other inputs, which system functionality to trigger in response to a change in a model, and which infrastructure to build the user experience on—all in a way that will be useful, interesting or engaging to users in their contexts. Gray and Salber, reflecting on [17], advanced system engineering for context-based systems by taking account of the inherent uncertainty of sensor data [13]. Chalmers described how history is a part of context, with consequences for ubicomp theory as well as system design [12], and Dourish explained how, in their activity, people continually negotiate and adapt the contextual features that are significant or relevant [18]. People reflect on how they are modelled, i.e. the people using a system may take account of its model of their identities, locations and states, accommodating it or even appropriating it to their own ends rather as one controls and adjusts one's self-presentation to others in human-human communication. This reinforces the view that ubicomp design should be treated as an iterative process that spans users, evaluators and developers, and that systems should be adaptive or adaptable.

Through repeated phases of design, implementation and evaluation, developers can change and improve their systems, and their understanding of use and context. However, evaluation in ubicomp can be complex in itself, and it can be difficult to communicate useful information back to the system developer. As Carter and Mankoff [19] put it, "Ubicomp systems [are] more difficult to evaluate than desktop applications. This difficulty is due to

issues like scale and a tendency to apply Ubicomp in ongoing, daily life settings unlike task and work oriented desktop systems.” Observational techniques founded in ethnography may be well suited in principle but in practice are often hampered because keeping up with the activity is difficult, small devices such as mobile phones and PDAs can easily be occluded from view, and people’s use may be intimately related to and influenced by the activity of others far away [13]. Several video cameras may be used to record activities in several locations set within some larger activity, but this brings the practical problem of synchronisation, and how to overview this material and combine it with other relevant data, such as system logs gathered from the mobile devices. Furthermore, network connectivity may be intermittent or costly enough to hamper attempts to keep in continuous contact with users and their devices, e.g. to stream log data back to evaluators or developers.

Some researchers have explored ‘experience sampling’ methods, that bring up a questionnaire on-screen when the mobile device detects that it is in a context of interest [20]. Carter and Mankoff recently developed Momento [21], which supports experience sampling, diary studies, capture of photos and sounds, and messaging from evaluators to participants. It uses SMS and MMS to send data between a participant’s mobile device and an evaluator’s desktop client. Combining quantitative log data and more qualitative data to give a more holistic view of ubicomp systems in use is the aim of our *Replayer* system [8], an integrated set of tools for display and interaction with different parts of the record of system operation, use and context. Log data visualisations combine many usage histories so as to support analysis of users acting in greater numbers, and at larger geographic and time scales, than we can directly observe. Complementary Replayer components support synchronised replays of audio and video recorded during user trials. Separate to these, we have demonstrated a rudimentary plug-in for the Visual Studio IDE (or ‘add-in’ in Visual Studio terminology) that adds logging methods visibly into the source code. While not robust or complete enough for real-world deployment, it shows that a project, class, methods or variables could potentially be instrumented with a single click, so that code to create logs in a Replayer-compatible format might be created more easily.

Ubicomp software is even harder to evaluate and analyse when it supports the adaptation and customisation needed to sustain contextual fit in long-term use. Such support is an especially important goal in ubicomp, as computers’ use may extend beyond work activities, focused on pre-planned tasks, into leisure and domestic life. Recently, Humble et al. supported end-user adaptation in ubicomp via a distributed shared state model and a simple set of categories of transformation between physical effects and digital effects [22]. This was expressed using a ‘jigsaw piece’ editing metaphor, in which manual assembly of components is based on display of what a component could be connected to. *Speakeasy* [23] relied on a fixed set of interaction patterns and ‘task templates’ so as to allow a mobile component to be written with less prior knowledge of the particular components it will communicate with when deployed. In effect, many of components’ combinations and uses are fixed in advance by the form of the pattern or template, and it was noted in [23] that a more dynamic approach would be preferable, i.e. one in which software structure is tailored to the user’s context and history.

Using context and history in this way is the focus of our *Domino* toolkit [6]. *Domino* manages and supports adaptation of the set of software components making up a ubicomp application. It uses mobile ad hoc networks and controlled epidemic algorithms to spread components and usage histories between users. It tracks and logs the current system ‘context’ in terms of the set of components currently running in the application. This set is used to filter usage histories in the course of making recommendations of new components to install and run. Therefore, instead of relying solely on predefined templates or patterns of use, like *Speakeasy*, *Domino* also takes advantage of emergent patterns of use in recommending and integrating components. This is done in a way that hides from the user much of the technical details of discovery and integration of new software, but reveals enough to let him or her maintain control over the system. When a user accepts a recommendation, *Domino* checks whether other currently running components satisfy the component’s dependencies—required interfaces declared *a priori* by the developer—in an effort to ensure that its execution is technically feasible before trying to install it. If not, it suspends installation until it finds such components, which then can be recommended, accepted by the user, and installed. Dependencies specify objective constraints on component combination, like the connections in Humble’s jigsaw editor and *Speakeasy*’s templates, but *Domino* goes beyond other work in this area by also taking advantage of the evolving patterns of use that represent users’ subjective preferences about component combination. In pilot runs of the *Castles* game, built using *Domino*, we created Replayer-compatible system logs, and found Replayer’s visualisations invaluable for post-trial analysis: they made understanding the dynamics and patterns of system structure and use tractable. Increasing the coupling and synergy of *Domino* and Replayer is one of the goals of the proposed work, as the following section discusses.

Programme and Methodology

We take a holistic and exploratory approach, in accord with Equator work and the UKCRC Ubicomp Grand Challenge. We propose research in a number of areas, each outlined separately below as work packages, and laid out temporally in the Work Plan appendix. Our **project structure and management** builds on our Equator experience by working as a tightly integrated group in ways that blur the boundaries between these areas, e.g. members primarily active in theory and user studies will take part in system design discussions for both applications and infrastructure, and theoretical reflections and study findings will be discussed with developers

and evaluators. We will alternate between periods in which the group works closely together, collectively focusing on user trials of applications and associated tools, and ‘loosely coupled’ periods in which tools and infrastructure are developed, background user studies are made, and conceptual frameworks are created and refined. We propose to develop and evaluate applications in two areas: mobile multiplayer games and cultural tourism. By comparing our findings from these related but distinct areas, we should be able to draw useful generalisations and conclusions with regard to system design, evaluation and use.

The RAs and students will work together in one large office so as to share and maintain awareness of each others’ work, even when loosely coupled. More formal group meetings and student supervision meetings will happen weekly. We also use instant messenger and a wiki for networked interaction, and use Subversion as our internal code repository and Sourceforge for the public version of our code. In order to link well with our Cambridge collaborators, we intend to hold regular joint workshops for all project members, in particular intense multi-day ‘codefests’ in which all project members, from coders to evaluators to PIs, gather to design, implement and test systems, to observe systems in use, and to be presented with ongoing findings from larger-scale user trials. Lastly, after one year we will run a project review workshop, with Dr. Ken Wood (Deputy Managing Director, Microsoft Research Cambridge), Prof. Tom Rodden (U. Nottingham and Equator IRC Director) and Prof. Michel Beaudouin-Lafon (U. Paris–Sud and ACM Fellow) acting as external reviewers (letters of support attached). We will interact with these reviewers more informally throughout the project, but in this review workshop they will formally comment of our progress and plans, and we will make their report public. Project PhD students will, with the support outlined in Dr. Wood’s letter, apply for internships at MSR so as to deploy our tools in his lab.

WP1: Visualisation tools for evaluators and developers

We propose new visualisations for analysis of both ‘live’ usage data streaming in from all the mobile devices used in a trial (see WP3), and accumulations of that data (i.e. usage histories). These tools will also support real-time collaboration between evaluators and coders working on that data. Apart from being coupled with evaluators’ tools, developers’ tools will be tightly coupled with the everyday coding tools inside the the Visual Studio IDE (VS), offering detail and overviews of the use in system operation of components, classes, methods and variables. For example, one might choose a particular area of the city in which a tourism trial is taking place, and analyse real-time log and video data from the trial participants there, comparing it with data from prior test runs. WP6’s studies of these tools (and those of WP2) will answer research questions (RQs) 1, 2 & 5. Morrison will be the RA primarily involved in this WP.

WP2: Extending Domino’s Model of Context

At present, Domino’s model of context is univariate: the modelled dependencies and associations are only between software components. Building on *George Square* [2], which used patterns of association in usage histories to combine locations and URLs, we will enrich Domino’s information on context, operation and use. New VS tools will allow the programmer to specify and analyse dependencies and patterns of association between components, locations and other logged contextual features. The programmer will then use the infrastructure of WP3 to disseminate new components among user trial participants, along with any specified usage histories needed to ‘bootstrap’ the recommendation process. For example, a programmer might specify that a new component offering an audio tour of a museum can only be started up when the user’s device is in that museum. In a trial, he could then analyse the sets of components running when the new one was installed and used, where in the museum this occurred, and video recordings of user discussion at those times. WP6’s study of the products of WP2 (and WP1) will answer RQ1, 2 & 5. The RA primarily involved in this WP is Hall.

WP3: New Infrastructure for Domino

The distribution and integration of Domino components is, at present, relatively inefficient and insecure. By porting Domino to run on top of the .NET implementation of Huggle, we can use secure opportunistic forwarding via both fixed and mobile ad hoc networks, using multiple forms of connectivity, e.g. WiFi, SMS, GPRS and Bluetooth. Domino will keep its facilities for recommendation and adaptation, but will use Huggle for device discovery, transmission of usage histories and components, resource management and security. We will thus support collection of usage data from mobile devices, and distribution of new components and usage data from the IDE to those devices—including new Domino components that offer users choice over Huggle’s facilities for communications, resource management and security. Hall will be the RA for this WP, and the primary link to the Cambridge Huggle group, while a PhD student (RS1) will focus on usable security and privacy built on Huggle’s facilities and techniques such as audit trails, virtual machine sandboxing and OS virtualisation. Quantitative testing in this WP and user trials in WP4 & WP5 will demonstrate an answer to RQ4.

WP4: Application Design 1: Mobile Multiplayer Games

We will build and trial games that use Domino components for position tracking, revealing and recording different mapped information, and for different forms of Huggle connectivity to control the spread of data and software among teammates and opponents. Web-based interfaces to central servers will offer overview and

summary information to game players, but the bulk of the use will be on smartphones. We intend to advance mobile games with by-products, i.e. games that serve as an effective means to collect data that can be used in other applications, building on the ESP Game [25] and MSM [26], games with by-products of databases of text labels for images. Game rules, goals and interactions will exploit features of WP3's new infrastructure. As an indicative example, players might compete to 'own' a Glasgow location by using Domino software to find features there, georeferencing, photographing and labelling it, and then selectively spreading this data to others for validation. User studies in will use quantitative and qualitative methods—and the new tools of WP1—to show how users' analysis of ongoing context, operation and use help them in adapting all three, i.e. to answer RQ3. Barry Brown (UCSD), acting as external observer, will help us document how timely and informed social interaction among users and the project members working on this WP—Hall, Morrison and RA3, assisted by Chalmers, Gray, and PhD students RS1 and RS2—is used in creating and sustaining systems' contextual utility, i.e. providing parts of our answers to RQs 1, 2 and 5.

WP5: Application Design 2: Cultural Tourism

The products of WP4 will be a starting point for software for cultural tourists, to enhance and enrich their visits to the cities of Glasgow and Edinburgh. We will focus on the numerous arts and science festivals run throughout the year, in particular the Edinburgh Festival, the world's largest arts festival. In effect, we will extend our prior work on city visiting and cultural tourism [2] in scale, mobility and dynamics. Like WP4, we will demonstrate processes of design and user-driven adaptation in experiments involving users with smartphones and web-based overview tools. Sets of components will be deployed and evaluated over a series of relatively long visits (e.g. a week, like those we of the recruited tourists in [2]) so that each trial participant has time to understand, adapt and appropriate our system. Building on our earlier collaboration with Scottish tourist offices, we will give visitors smartphones running a simple map-based guide for the city centres, and each visitor will have access to a growing set of additional software components that highlight different urban areas, events and topics, and support the selective sharing of photographs, locations, comments and components. User studies will use quantitative and qualitative methods—and the new tools of WP1—to show how users' analysis of ongoing context, operation and use help them in adapting all three, i.e. to answer RQ3. Barry Brown (UCSD), acting as external observer, will help us document how timely and informed social interaction among users and the project members working on this WP—Hall, Morrison and RA3, assisted by Chalmers, Gray, RS1 and RS2—is used in creating and sustaining systems' contextual utility, i.e. providing parts of our answers to RQs 1, 2 and 5.

WP6: Studies of Use, Design and Evaluation by Others

This WP studies socio-technical environments other than our own, documenting the tools and processes of design, use and evaluation in each. This will let us compare our use of our own tools/infrastructure with how others use either our tools/infrastructure or their own. We will run 'hackfest' tutorials and workshops, lasting a few days and using our tools/infrastructure to develop and test mobile applications. These will involve Huggle members, Glasgow colleagues from *Designing for the Augmented Stadium* (EP/E04848X, 10/07-3/10) which also involves MANETs and has Chalmers as PI, and other interested researchers. We will also carry out studies of other groups designing and evaluating mobile systems, locally and among our former Equator colleagues. We will concentrate on how they handle adaptation and change in their systems and trials, spanning the evaluators' and developers' workplaces, and the places in which trial participants use their systems. Focus on evaluators, developers and users in turn will contribute to our answers to RQ1, 2 and 3 respectively, and broader treatment of the social interaction among all three roles, creating and sustaining systems' contextual utility, will be part of our answer to RQ5. RA3 would be the primary RA engaged in this area of work, aided by RS2.

WP7: Theoretical framework and abstract models

This WP aims to codify and generalise from the shifting set of people, software, uses and contexts in WP1–6, in ways that may be useful for the design of new user experiences, applications and tools, i.e. to answer RQ6. We will create more abstract formal representations of Domino's context model, and of the systemic aspects of the recommendation and adaptation process, in ways that complement and generalise over the observations and analyses of WP6. The basic historical representations and recommendation algorithms used in Domino draw directly from structuralist linguistics and Hillier's 'space syntax' theory of urban form [14], and we will again combine structures and concepts from those fields with computer science theory in developing useful abstractions of the process of recommendation, adaptation and use. We aim to express the *type* of a component not only in the traditional terms of static names, interfaces and dependencies, but also in terms of dynamically evolving patterns of association. We will feed back into IDE, infrastructure and interaction design features of WP1–3, e.g. widening Domino's resource management to include techniques such as predictive caching, and feed design concepts to explore in WP4 and WP5. Primarily, Chalmers and Gray will be involved in this area.

WP8: Public web resource of code and data

This WP will make a publicly accessible web resource from the tools, applications and data arising from the project, forming an evidential record of our system trials and supporting others' research. We will make source

code available for others' use, via Sourceforge. These tools and applications will be complemented by anonymised data sets from the user trials: records of our work and material for others' analysis, comparison and simulation. Our new visualisation tools will present the quantitative data on behaviour and performance alongside qualitative data. In order to ease access to the software and data sets, we will create demonstration systems and tutorial materials of a level suitable for undergraduate and masters students, but primarily aimed at workshops we run locally and at U. Cambridge. Morrison and Hall will be primarily responsible for WP8.

Relevance to beneficiaries

The computing industry and research community, including those working in mobile computing, have long wrestled with the problem of customisation and adaptation of systems. The techniques for design and evaluation explored here have the potential to be used in many areas of systems design involving assemblies of software components, and in development processes involving distributed systems and IDEs. Similarly, the tools for analysis and visualisation may have significant benefits for design processes related to mobile computing, but would also be useful in more general design contexts. Standalone Replayer tools will be applicable to a still wider range of application areas beyond computing, supporting other researchers' analysis of the project's data sets in understanding people's mobility, social interaction and use of mobile devices. Our collaboration with Prof. Crowcroft's group at U. Cambridge should offer the Huggle project the benefit of application trials, experimentally grounded data sets and design requirements for future middleware.

Given the deployment on smartphones, we see the potential for phone and handset companies to have new dynamic types of application that could boost their market. We also see the ubicomp research community as beneficiaries, as developing and debugging new mobile systems is a complex yet common task—which could be eased by the proposed data collection, distribution and analysis tools. The application area of games is of growing cultural, commercial and academic importance, and the new forms of gaming proposed in this work would help us maintain the UK status as leaders in this burgeoning field, and feed into the emergent market in this area. The novel area of mobile 'games with by-products' may have many positive spinoffs, with the energy of game players helping work in areas as varied as cartography, urban planning and tourism. The tourist industry is of huge economic importance to the UK in general and to Glasgow and Edinburgh in particular, and our new forms of adaptive application would support types of visit experience that are rich in information and yet sociable and woven into the fabric of the place visited, thus enhancing visit quality and the economic health of related tourist industries.

Dissemination and exploitation

We will vigorously publish our work in top level conferences and journals, maintaining our record of keynotes, papers and demonstrations. We will run a conference workshop on iterative design processes at Ubicomp or Pervasive, and run a UK tutorial/workshop on the use of our tools and processes. We will place robust documented versions of our code on Sourceforge as well as our own web site, offering easy access to a broad community of developers and users. We will use the web to publish suitably anonymised assemblies of usage history data, system log data and audio/video data collected from field trials, along with our tools to visualise and analyse them, thus supporting others' research and offering an *evidential record* of our system trials. We will also make available demonstrator systems and teaching materials for undergraduate/masters courses.

References

16. M. Weiser (1994) *Creating the Invisible Interface*, Invited talk, ACM UIST, p1.
17. A. Dey, G. Abowd, D. Salber (2001) *A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications*, Human Computer Interaction, pp. 97-167.
18. P. Dourish (2004) *What We Talk About When We Talk About Context*, Personal and Ubiquitous Computing, 8(1), pp. 19-30.
19. S. Carter, J. Mankoff (2005) *Prototypes in the Wild: Lessons Learned from Evaluating Three Ubicomp Systems*, IEEE Pervasive, 4(4), pp. 51-57.
20. J. Froehlich, M. Chen, I. Smith, F. Potter (2006) *Voting With Your Feet: An Investigative Study of the Relationship Between Place Visit Behavior and Preference*, Proc. Ubicomp, pp. 333-350.
21. S. Carter, J. Mankoff, J. Heer (2007) *Momento: Support for Situated Ubicomp Experimentation*. Proc. ACM CHI, San Jose, pp. 125-134.
22. J. Humble, A. Crabtree, T. Hemmings, K.-P. Åkesson, B. Koleva, T. Rodden, P. Hansson (2003) *'Playing with your bits': user-composition of ubiquitous domestic environments*, Proc. Ubicomp, pp. 256-263.
23. W.K. Edwards, M. Newman, J. Sedivy, T. Smith, S. Izadi (2002) *Challenge: Recombinant Computing and the Speakeasy Approach*, Proc. ACM MobiCom. Atlanta, pp. 279-286.
24. R. Milner (2006) *Ubiquitous Computing: Shall We Understand It?* The Computer Journal 49 pp. 383-389.
25. L. von Ahn, L. Dabbish (2004) *Labeling images with a computer game*, Proc. ACM CHI, pp. 319-326.
26. V. Tuulos, J. Scheible, H. Nyholm (2007) *Combining the Web, Mobile Phones and Public Displays in Large-Scale: Manhattan Story Mashup*, Proc. Pervasive 2007,