UbiVal: Fundamental Approaches to Validation of Ubiquitous Computing Applications and Infrastructures

Prof David S, Rosenblum,
Dr Cecilia Mascolo
Dept. of Computer Science
University College London

Prof Marta Z. Kwiatkowska,
Dr Dan Ghica, Dr Mark Ryan
School of Computer Science
University of Birmingham

Dr Naranker Dulay,
Dr Emil Lupu
Dept. of Computing
Imperial College London

Part I: Investigators' Previous Research Track Record

1. University College London

The UCL group have extensive experience in software systems testing and mobile computing, including realistic mobility models for simulation. Recent work extends this to context-awareness and modelling of ubiquitous computing systems.

Prof David S. Rosenblum (the Project Director) is Professor of Software Systems at University College London (UCL) and Director of London Software Systems (LSS), a research institute established jointly by the Software Systems Engineering group at UCL and the Distributed Software Engineering group at Imperial College London. He has worked for many years in distributed systems, software testing and runtime checking of software systems, including work on assertion checking, regression test selection, and testing of component-based software systems, as well as instrumentation approaches to support these techniques. In past work he extensively studied the problem of selective regression testing, in which program analyses are used to identify a minimal set of necessary test cases to run after changes are made to create a new version of a software system. His work on this problem represented a departure from the many competing approaches that had been developed, in that it employed coarse-grained rather than fine-grained analysis of the relationship between a test suite and system under test [8, 37, 5]. More recently he developed an approach to testing component-based systems that uses specialised component metadata providing limited information about a component (such as flow graphs) that a component developer packages with a component and makes available through the component's interface [33, 32]. He is currently investigating approaches for lightweight probabilistic modelling and analysis of publish/subscribe systems.

Rosenblum is currently chair of the steering committee of the International Conference on Software Engineering (ICSE). In 2002 he received the retrospective ICSE Most Influential Paper Award for his ICSE 1992 paper [36], and he recently gave a keynote presentation at the 2005 International Conference on Distributed Objects and Applications. He is a Chartered Fellow of the British Computer Society, a Fellow of the IEE, a Fellow of the IEEE (effective 2006), and Vice Chair of the ACM Special Interest Group in Software Engineering (ACM SIGSOFT). He has been an associate editor of the *IEEE Transactions on Software Engineering* and is currently an associate editor of the *ACM Transactions on Software Engineering and Methodology*.

Dr Cecilia Mascolo is an EPSRC Advanced Research Fellow and UCL Senior Lecturer. Her research has focused on enhancing adaptation via middleware services that provide a flow of context information from lower layers to the application layer and vice versa [7]. In the past she also worked on design and model checking techniques for mobile systems [9]. She has recently produced some initial studies on the use of social models in simulation and has devised routing algorithms that exploit these models [30]. In the EU Project RUNES (Reconfigurable Ubiquitous Network Embedded Systems) she is researching middleware solutions for ubiquitous systems. A large part of the project focuses on devising verification techniques for the middleware. She is part of the ESF MINEMA network of excellence on Mobile Computing Middleware.

In the EU Project PLASTIC (Providing Lightweight and Adaptable Service Technology for pervasive Information and Communication), both UCL investigators are studying a wide range of service provisioning issues for B3G (Beyond 3G) distributed computing platforms.

2. University of Birmingham

The Birmingham group's strengths are in software verification via model checking, security analysis, and probabilistic model checking.

Prof Marta Kwiatkowska is Professor in the School of Computer Science. She leads a group of ten postdoctoral research staff and students who work on probabilistic modelling and verification, and applications of model checking technology to industrially-relevant applications. Probability is frequently used to achieve self-organisation, adaptation and decentralisation in distributed systems, and therefore probabilistic modelling is particularly important in ubiquitous computing. *Probabilistic model checking* is an automatic method that enables the construction of probabilistic models from high-level descriptions and their analysis against probabilistic temporal logic properties.

Kwiatkowska led the development of PRISM (**PR**obabi**lI**stic **S**ymbolic **M**odel Checker) [21, 35]. Her contribution has been wide-ranging, and includes languages and models for probabilistic systems, logics and semantic foundations for randomised distributed systems, algorithms and compact data structures for probabilistic model checking, symbolic model checking of probabilistic real-time systems, and applications to protocol analysis [3, 25, 38]. PRISM is extensively used worldwide for research and teaching purposes; there have been 2600 downloads since 2001 and 100 publications using and/or contributing to PRISM, half of those external to Birmingham. PRISM has been used to model and analyse static configurations of numerous real-world Internet wireless, security, anonymity, non-repudiation, contract signing, negotiation and quantum cryptographic protocols, such as Bluetooth device discovery [13], IPv4 Zeroconf [24] and Crowds anonymity [39], as well as power management schemes, nanotechnology designs, controller dependability and biological processes [31, 22]. Recent improvements to PRISM include a simulator and sampling-based approximate probabilistic model checking that employs statistical hypothesis testing to establish whether a property holds with a given confidence level [42]. Current projects are addressing novel routing schemes for ad hoc network protocols inspired by swarm-intelligence [26] and extension of probabilistic model checking with mobility and abstraction. PRISM is supported by EPSRC, DTI, QinetiQ, Microsoft Research Cambridge and Formal Systems Ltd.

Kwiatkowska was invited speaker at the Logic in Computer Science Symposium (LICS 2003) [20] and is highly active in organisation and programme committees of leading conferences in verification (CAV, TACAS), semantics and concurrency (LICS, CONCUR, FOSSACS), and quantitative verification (QEST, FORMATS). In 2007 she will co-chair the Programme Committee of QEST 2007 (Quantitative Evaluation of SysTems). She serves on the Editorial Boards of the *Transactions on Computational Systems Biology* and the *Journal on Logical Methods in Computer Science (LMCS)*.

Dr Mark Ryan is a Reader in the School of Computer Science at Birmingham. He works on applications of logic in computer science, particularly in verification of security systems via model checking. Ryan has recently worked in analysing security protocols against their expected properties. Prior to that he developed techniques to analyse access control systems. Ryan's book *Logic in Computer Science: Modelling and Reasoning about Systems* [17], co-authored with Michael Huth, is the first major book introducing techniques and methods for verification by model checking to non-specialist readers. Now in its second edition (2004), and having sold over 10,000 copies, the book has pioneered the teaching of verification in graduate and undergraduate curricula the world over, and is recommended as a core text for over 100 university courses world-wide.

Ryan is on the programme committee of several conferences related to security and verification of mobile systems, including Computer Security Foundations Workshop (CSFW, 2006), 12th International Symposium on Temporal Representation and Reasoning (TIME, 2005), International Colloquium on Theoretical Aspects of Computing (2005, 2006), Foundations of Computer Security (FCS, 2005), and Model Checking and Artificial Intelligence (2005).

Dr Dan Ghica is a Lecturer in the School of Computer Science at Birmingham. His expertise is in compositional software model checking. In collaboration with McCusker he developed a method, based on game semantics, of modelling and verification for open procedural programs. He obtained his doctoral degree in 2002 from Queen's University, Canada with the dissertation *A Games-Based Foundation for Compositional Model Checking*. Subsequently, he worked as a Research Officer at the Oxford University Computing Laboratory, where he continued to investigate foundational aspects of game semantics and their application to model checking [2] and static analysis [12].

Dr Gethin Norman is the named Research Fellow on this project. He obtained a degree in Mathematics from Oxford University, followed by a PhD in Computer Science at Birmingham, and is currently employed on EPSRC project GR/S46727 concerning probabilistic model checking of mobile ad hoc network protocols, of which the proposed project is a natural continuation. He has contributed in a major way to the foundations of probabilistic model checking, particularly for real-time and quality of service, and PRISM modelling of protocol case studies and dynamic power management [22, 31]. Recently he collaborated with Shmatikov on analysing probabilistic contract signing and fair exchange protocols [35]. He is an experienced researcher whose expertise is unique in the UK and particularly well suited to this project.

3. Imperial College London

Dr Naranker Dulay is a Senior Lecturer in the Department of Computing. He obtained his BSc in Computer Science from the University of Manchester and his PhD from Imperial College. His research interests include middleware and specification languages for distributed, mobile and ubiquitous systems. He has over 50 publications in international conferences and journals and has served on the program committees for Middleware, ICDCS, Persec, Policy, SMPW, PM4W, SEM, DAIS. He was program co-chair for Policy 2002, and he was an invited speaker at Mathematical Methods, Models, and Architectures for Network Security Systems (MMM-ACNS 05). He is currently principal investigator of two projects funded by the EPSRC, and co-investigator on five others.

Dr Emil Lupu is a Senior Lecturer in the Department of Computing. He obtained his Diplôme d'Ingénieur from the ENSIMAG, Grenoble, France and his PhD from Imperial College. His research interests include network and systems management, security and design issues in large distributed systems. He has over 40 publications in international journals and conferences and has served on the organising and program committees of over 20 international conferences, including the IFIP/IEEE symposia and workshops on Network and Systems Management (IM, NOMS, DSOM). He was program co-chair of the Policy Workshop 2001 & 2004 and of the 5th IEEE Enterprise Distributed Object Computing Conference. He is currently principal investigator of three projects funded by the EPSRC, EU or industry and co-investigator on five others.

Both investigators are members of the Distributed Software Engineering research group at Imperial, which combines practical work on tools for the design and implementation of distributed and ubiquitous systems with more formal software engineering approaches. The investigators have an excellent track record of application-driven projects in relevant areas and will lead the research in applying UbiVal's validation tools and techniques to real-world ubiquitous systems via their participation in Cityware, a project funded through the first WINES call that combines pervasive systems with design of urban spaces. Imperial's research in Cityware covers middleware and specification techniques for context-based services, overlay networks, management policies and space-syntax provisioning in pervasive systems. The investigators are also involved in ongoing work with Sussex University on modelling uncertainty of contextual data. Imperial co-ordinate the UK Ubi-Net workshop series. In autonomic computing the investigators have two projects investigating scalable architectures for the management and monitoring of ubiquitous systems: AMUSE for ubiquitous healthcare, and SMMC for autonomous vehicles. In trust, security and privacy, the group currently have two projects: Caregrid, an EPSRC project researching trust negotiation, adaptive security and privacy for e-Health, and Trustcom, an EU project investigating trust and contract management for dynamic virtual organisations.

4. Complementary Strengths of the Investigators

The consortium brings together a number of leading UK researchers who have not worked together previously. The three groups provide highly complementary areas of expertise and skills to the project. The UCL group have expertise in mobile middleware, simulation models and software testing. The Birmingham group are expert in probabilistic modelling and verification of software, and also verification of security properties. The group at Imperial provide expertise in ubiquitous computing, distributed systems and instrumentation techniques. Furthermore, they are investigators on the EPSRC-funded Cityware project awarded during the first year of the WINES initiative and thus will play an important liaison role for UbiVal. All three groups have a common philosophy of applying modern software engineering practices to the design and implementation of adaptive distributed systems. The groups have contributed to the UK Grand Challenge exercise in ubiquitous computing, with Kwiatkowska being a member of the Steering Committee of the Ubiquitous Computing Grand Challenge (UbicompGC) on behalf of the UKCRC (UK Computing Research Committee).

5. Industrial Partners

The activities of UbiVal will be supported by British Telecom, QinetiQ and HP (a partner in Cityware); letters of support from these companies are attached to this Case for Support. The group at BT (Nigel Walker) are contributing to the 21C Next Generation Network, and are keen both to provide input into the project as well apply the techniques being developed at certain network layers. QinetiQ (Colin O'Halloran) are long-standing sponsors of the PRISM tool and collaborators on the FORWARD project. Their work on safety and security assurance for pervasive systems is synergetic with the goals of UbiVal and will provide ample scenarios of high-security computing systems to exercise the integrated suite. HP (Phil Stenton) are major contributors to the Mobile Bristol project investigating the use of pervasive computing to enhance user experience as they interact with their physical environment and with each other in urban spaces. They will offer input as to the type of devices and everyday usage scenarios.

In addition, we anticipate that the Cityware project will provide further potential industrial partners as its work progresses.

6. Some Related Grants and Contracts of the Investigators

- CAREER: Mechanisms for Ensuring the Integrity of Distributed Object Systems, US National Science Foundation CCR-9701973, 1997–2001, studied novel mechanisms for design and validation of distributed component-based software systems.

- Specification and Dynamic Checking of Composition Constraints in Distributed Component-Based Systems, US Air Force Office of Scientific Research F49620-98-1-0061, 1997–2001, studied architecture-based validation of distributed component-based software systems.

- A Future Of Reliable Wireless Ad hoc networks of Roaming Devices (FORWARD), DTI/QinetiQ, 2003–2005, collaboration between QinetiQ, Formal Systems Europe Ltd, and Universities of Oxford and Birmingham to analyse security and QoS of ubiquitous computing protocols, focused on the Bluetooth protocol.

- Probabilistic Model Checking of Mobile Ad Hoc Network Protocols, EPSRC GR/S46727, 2003–2006, addressing mobility in probabilistic systems and the corresponding model checking algorithms.

- Automated Verification of Probabilistic Protocols with PRISM, EPSRC GR/S11107, 2003–2006, investigating abstraction and model reductions in probabilistic model checking.

- Cityware: Urban Design and Pervasive Systems, EPSRC EP/C547586/1, 2005–2009, a WINES project investigating smart urban spaces.

- BiosensorNet: Autonomic Biosensor Networks for Pervasive Healthcare, EPSRC EP/C547705/1, 2005–2009, a WINES project investigating issues for body-area sensor networks for patient care.

- Autonomic Management of Ubiquitous Systems for e-Health (AMUSE), EPSRC GR/S68033/01, 2004–2007, investigating ubiquitous healthcare.

- Self Managed Mobile Cells (SMMC), BAE Systems, DTC/RAO/WPE /N03751/SEAS, 2005–2008, investigating autonomous vehicles.

- Autonomous Trust Domains for Healthcare Applications (CareGrid), EPSRC EP/C537181/1, 2005–2008, providing a grid infrastructure for trust-negotiation, adaptive security and privacy.

- Trust and Contract Management framework (Trustcom), EU IST-2002-2.3.1.9, 2004–2007, investigating trust and contract management for dynamic virtual organisations

- Reconfigurable Ubiquitous Network Embedded Systems (RUNES), EU-IST-004536, 2004–2007, investigating a framework for development and verification of networked embedded systems.

- Providing dependabLe & Adaptive Service Technology for pervasive Information & Communication (PLASTIC), EU-IST-26955, 2006–2008, investigating service composition and orchestration in ubiquitous systems.

- Coordination and Reliability Mechanisms for Adaptive Mobile Middleware (CREAM), EPSRC EP/C544765/1, 2005–2010, investigating primitives for communication and coordination in mobile computing middleware.

- The Divergent Grid: Dealing with Extreme Heterogeneity and Dynamicity in Next Generation Grid Middleware, EPSRC EP/C010345/1, 2005–2008, investigating the extension of grid systems to heterogeneous and ubiquitous environments.

Part II: Description of the Proposed Research

1. Introduction and Motivation

Real exemplars of ubiquitous computing environments have begun to emerge. In the UK there are EQUATOR, SmartIts and Mobile Bristol for technology and interactive experiences with mobile devices; the Envisense centre for environmental monitoring and control; and UbiCare for ubiquitous medicare and patient monitoring via wearable sensors. Similar projects are also being undertaken abroad.

However, the primary emphasis of these demonstration projects has been on engineering, deployment and user experiences in specific scenarios, *not* on the generic lessons and fundamental principles for their *effective engineering*.

Consider the design of a smart urban space comprising a network of cooperating servers, displays, sensors, wearable tags and wireless access points. Visitors and residents interact with the smart space via mobile devices such as PDAs and mobile phones. The smart space provides context-aware applications such as tourist guides and restaurant booking services. Visitors and residents also may be interested in forming mobile ad hoc communities for collaborative activities, such as shopping, attending a concert or rugby match, and gaming. Examples of relevant context include location, time of day, environmental readings (e.g. temperature, humidity), physiological state (e.g. heart rate), user preferences (e.g. spoken language, spending limits) and current role.

The design of such applications is fraught with difficulties unique to the nature of ubiquitous systems, and their development requires sophisticated validation techniques. Consider the periodic, asynchronous delivery of location-based information to mobile users in the urban space. The correct and efficient functioning of this application will be influenced quite strongly by changes in the context of the application and by the mobility of the users. Validation of the application thus requires adequate exploration of the space of context inputs and mobility trajectories to which the application will be subjected. Existing techniques for validation such as model checking, testing and simulation can provide only limited confidence in the correctness and efficiency of the application, because they employ a very simplistic representation of external forces on the application. For instance, while test methods are traditionally built around suites of *initial* test inputs to the system under test, testing a ubiquitous computing application such as the one described above requires *continuous* feeding of context inputs to the system. Greater confidence in the application can be gained only through the use of validation techniques that account for the application's context-awareness. Consequently, methods able to *predict* and systematically *validate* the behaviour of ubiquitous systems in advance of deployment are necessary, yet no such methods exist at present.

At the same time, the risks to business and society posed by flawed ubiquitous systems are apparent from phenomena such as bluesnarfing[1], where it is possible to gain access to private data on a Bluetooth phone without the user's knowledge, or even to overhear a conversation, due to security loopholes. Privacy, security and trust in ubiquitous computing environments demand new validation approaches and mechanisms because of the potential of new forms of malicious attacks in the wireless medium.

As illustrated above we believe that engineers are building ubiquitous computing systems in the absence of sound engineering methods tailored to the unique characteristics of these systems. *The aims of this project are thus to rectify the dearth of sound engineering methods and to define a comprehensive, principled approach to validating ubiquitous systems with an associated suite of software tools.*

2. Background

As argued above, ubiquitous computing applications and environments exhibit characteristics that raise numerous challenges for developers of applications and their underlying middleware infrastructures. Ubiquitous systems offer a high degree of complexity, which is due to the variability of *context*, *mobility* and *heterogeneity* of the devices involved, and the possible *decentralisation* of some of these systems. Still, ubiquitous systems need to offer all those non-functional capabilities such as security, trustworthiness, reliability and responsiveness that are required in traditional systems.

In the development of traditional distributed systems, automated validation of software is carried out through some combination of *testing*, *simulation* and/or *model-checking*. Furthermore, these validation activities are supported by additional automated techniques such as *model extraction* from source code via static analysis, *test case generation*, and *code instrumentation*. Testing invokes the software directly, whereas simulation involves building a model and exploring its behaviours by generating possible trajectories for anticipated scenarios. Instrumentation itself can be performed for a variety of purposes, such as *test coverage monitoring*, *runtime consistency checking* between execution behaviour and model properties, and *fault recovery*. Model checking [17] refers to the activity of exhaustively analysing the behaviour of the model, either built directly or extracted from source code, in order to establish that the requirements hold. The array of existing automated approaches to validation is powerful and broad, yet they were never designed for the unique nature and complexity exhibited by ubiquitous systems.

Testing is a common feature of software development environments but does not yet adequately account for context-awareness, mobility, and low-level languages and platforms such as J2ME. Indeed, the only work in testing ubiquitous systems of which we are aware is recent work on using *metamorphic testing* to generate test case variants for testing context-aware middleware [40]. Previous work in software reliability testing has developed statistical methods for sampling a program's input space in order to develop test suites that adequately 'cover' the input space (for instance, see Podgurski and Yang [34]). However, in order to account for context in testing ubiquitous systems and to ensure adequate coverage of context, fundamentally new statistical methods of *context sampling* are required such that the sampled contexts can be applied in systematic fashion *throughout* the execution of the system under test via instrumentation.

Simulation techniques are the main analysis method applicable to mobile systems, but they suffer similar difficulties with the explosion of contexts. An additional problem is that existing mobility models are too simplistic [6]. One example is the *random waypoint* model that, unfortunately, produces predictions that are inadequate and irreproducible in practice [41]. Some realistic mobility models have begun to emerge [30]. For instance, Jardosh et al. presented a technique for the creation of mobility models that include the presence of obstacles [18]. The specification of obstacles is based on the use of Voronoi graphs in order to derive possible pathways in the simulation space. Maeda et al. also present a realistic pedestrian mobility model [27]. However, these works are limited, very recent and not yet used for validation. The US National Science Foundation (NSF) has funded a repository of execution traces called CRAWDAD to enable further understanding of mobility [10]. However, no matter how many traces are collected, the explo-

[1] http://www.thebunker.net/security/bluetooth.htm

sion of possible combinations of context, mobility and adaptive behaviour calls for more generic tools that, starting from these traces, can offer more flexible mechanisms for the validation of the designed systems.

Probabilistic model checking [38] addresses the analysis of stochastic models, such as randomised protocols and networks. In contrast to simulation, also applicable in this case but yielding approximate results because of partial coverage of executions, probabilistic model checking is exhaustive, and can produce exact numerical quantitative answers. Examples of properties that can be established are the probability that the level of quality of service will remain above minimum over some specified time, and the maximum energy consumption of a power-saving scheme. Verification of *security, privacy and access control* aims to address user-specified policies by employing tools based on calculi for mobile processes (i.e., pi-calculus) [1, 19]. This approach has been successful in discovering and fixing faults [29] but lacks generality. More advanced techniques such as probabilistic model checking promise to be able to capture context and mobility in a stochastic fashion, by taking into account the likelihood of context or action. They have been successfully applied to *quantitatively* analyse finite-state probabilistic ubiquitous system models consisting of a fixed number of agents (e.g., IPv4 Zeroconf dynamic link-local addressing [24] and dynamic power management [31]). Using the PRISM tool, a fault was found in an anonymity protocol [39] and the *worst case time* for Bluetooth device discovery was obtained for the first time for two devices [13]. However, probabilistic model checking is presently limited in that it is not scalable, has not been applied to real code, and cannot be applied to obtain predictions in real time. Some scalability improvement can be obtained with sampling-based statistical Monte Carlo techniques [42], but context and mobility are poorly understood.

The above observations lead us to conclude that the engineering foundations of ubiquitous computing are still in their infancy. Developing ubiquitous systems is a challenging task, and we believe that researchers consequently must revisit much of the current theory and practice of software validation and to formulate new validation approaches. Furthermore, non-functional properties such as security, and quality-of-service attributes such as performance, are of paramount importance in ubiquitous systems and must be first-class concerns in system design and validation.

We therefore also argue, in agreement with the UbicompGC Grand Challenge initiative[2] for which this proposal represents one of the *foothill projects*, that an interaction between researchers in theory and practice of ubiquitous computing is necessary to advance the state-of-the-art.

3. Research Objectives

The project will have the following specific objectives:

1. To develop a *comprehensive suite of validation techniques and their implementation* that are tailored to *mobile*, *adaptive*, *context-aware* ubiquitous systems. In particular,
   (a) To develop *formal models* and associated *model-checking techniques* for ubiquitous systems
   (b) To develop *testing techniques* for ubiquitous systems
   (c) To develop *simulation techniques* for ubiquitous systems
   (d) To investigate ways of leveraging one technique to support another (e.g., using simulation results to guide test case selection)
   (e) To apply these techniques for validation of both functional and non-functional properties (especially *security* and *performance*)
2. To develop the necessary *scientific and engineering foundations* that will support the validation techniques:
   (a) To develop *probabilistic and stochastic representations* that capture the ways ubiquitous systems achieve decentralisation and self-organisation, and the ways their quality-of-service attributes (such as response time and and energy efficiency) vary dynamically
   (b) To develop *realistic mobility models* based on social network theory to account for the human-driven mobility patterns typical of ubiquitous systems
   (c) To develop *techniques for transparent instrumentation* of the real environment as a driver for accurate simulation models, for test coverage analysis, for trace generation, and for model extraction from source code
3. To *evaluate* the validation techniques on *significant cases studies* in a realistic application domain of ubiquitous computing

When we refer to ‘ubiquitous systems’, we mean both the end-user ubiquitous applications as well as the middleware used to support them (especially context-awareness and mobile middleware).

Clearly, the effectiveness of this new engineering method must be demonstrated through empirical evaluation on realistic ubiquitous computing systems. We therefore will partner with the Cityware project, funded under the first WINES call issued in 2004, because it will provide an ideal range of systems for such empirical evaluation, including applications deployed on mobile handheld devices and infrastructure systems deployed in smart urban spaces. To achieve our objective of realistic evaluation, we will instrument the applications and middleware developed in Cityware and then predict and validate their behaviour using the suite of validation approaches we develop. We will also undertake case studies in other application domains in conjunction with our industrial partners.

We note that there are other related objectives that could be mentioned but will be deferred to future projects in order to maintain a reasonable and coherent scope for this project. *Debugging* and *fault isolation* in particular are important issues for any validation technique, but it is difficult to promise development of systematic debugging approaches without knowing exactly how the validation approaches will operate (and interoperate), or what kinds of information the approaches will reveal from case studies on realistic ubiquitous systems.

4. Methodology and Work Programme

This section describes the breakdown of the project into work packages, and Appendix A presents a GANTT chart of the schedule of work packages and deliverables, and some of their key interdependencies (with other dependencies implied by temporal separations). Work package relationships are also depicted in Figure 1. We propose a 3.5-year project that will implement the objectives described in Section 3. We believe that this timescale is necessary to enable development and transfer of results from theory to implementation and the flexibility

[2] `http://www-dse.doc.ic.ac.uk/Projects/UbiNet/GC/`

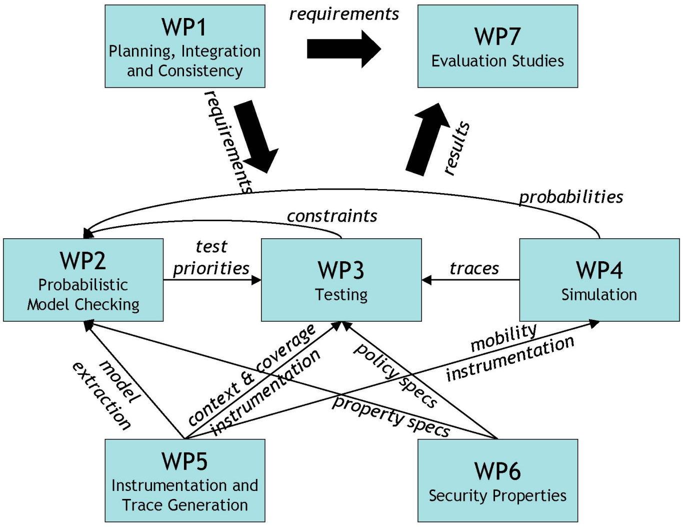

Figure 1: Work Package Dependencies.

needed for experimental validation using a testbed (the Cityware systems) that is yet to be deployed, expected 2008. Shortening the project would carry the risk of being unable to perform and obtain feedback from the experiments. We have built in explicit provision for all students to spend time with Cityware partners in order to familiarise themselves with the demands and needs of ubiquitous systems development to guide their research. Papers will be published on many of the results of the work packages and are thus implicit deliverables for all work packages.

WP1: Planning, Integration and Consistency

Lead Institution: UCL (Rosenblum, Mascolo)

Research Personnel: All PhD students and RAs

Description: In order to provide a suitable focus for the development of our suite of validation techniques, we will study the applications and infrastructures under development by our partners in the Cityware project, which is our primary target for case studies. To avoid the bias that might result from looking at a single application domain, we will also study applications of interest to our industrial partners and to other groups with which we have contact, such as Mobile Bristol and NICTA. In all of these cases we will identify important system features and typical usage scenarios. We will also need to study the context-awareness middleware that will be developed by Imperial for the Cityware project, since some work packages require knowledge of the middleware API. If necessary we may also select a more established middleware for initial research, such as the Context Toolkit [11]. Clearly, any middleware-dependent techniques and tools should be designed in such a way that the specific API dependencies are isolated and minimised.

This work package will also have responsibility for ensuring consistency and interoperability among the tools. This will be aided by definition of scenarios in which use of one validation approach leverages the outputs obtained from another. There are many ways in which the approaches can leverage each other. For instance, simulation traces obtained from the work in WP4 will be used in WP3 to feed context values to test runs in order to make testing context-aware. As another example, results obtained from simulation runs in WP4 will be used to parameterise probabilistic models developed in WP2 in order to make the models more realistic. As a third example, results from probabilistic model checking obtained in WP2 will be used to prioritise test case selection and context trace selection in WP3 in order to help minimise the number of test runs. A more detailed description of the relationships and integration of the work packages is depicted in Figure 1.

Deliverables: (D1.1) a detailed characterisation of the features and properties of ubiquitous systems that will be studied in the project; (D1.2) a selected context-awareness middleware that will be used in the project; (D1.3) initial project Web site; (D1.4) interoperability solutions for the tools developed in the project.

WP2: Probabilistic Model Checking

Lead Institution: Birmingham (Kwiatkowska, Ghica)

Research Personnel: Birmingham RA and PhD Student 1

Description: Probabilistic model checking can provide feedback about system development early in the lifecycle, when abstract behavioural models are constructed and before a working implementation is developed. It can also provide feedback later in the lifecycle through analysis of behavioural models automatically extracted from implementation code or derived from protocol standards. UbiVal will investigate both of these approaches in the context of ubiquitous systems.

The former is intended to provide early feedback on the effects of context-awareness, context adaptation and mobility on system design, and will entail formulating appropriate abstract probabilistic models of ubiquitous systems and middleware infrastructures. These models are intended to represent the system behaviour at a high level, so that its design can be improved based on preliminary analysis; we expect to draw on the the concept of context-awareness program points (WP3) and on (non-probabilistic) process calculi and logics for mobility (e.g., see Milner [28]). The abstract models will be parameterised by context and mobility traces, ultimately to be obtained from WP3 and WP4, and our probabilistic model checking algorithms [38] will be reformulated for this setting by adapting our approximate Monte Carlo methods [42]. We will focus on efficient algorithms and implementation of the quality-of-service and performance properties (offline methods), such as estimating the probability or expected time of delivery of a service.

In another strand, we will consider model checking of ubiquitous software components, for example, protocol standards such as Bluetooth, distributed coordination, anonymity, etc.

Many of these are necessarily probabilistic. Based on our earlier work [13, 23], we will develop techniques for conformance checking of protocol implementations (assumed to be Java/J2ME) in the style of SLAM [4] against the specification, in the presence of environmental constraints (e.g., location, power level, etc) specified as *context-likelihood pairs*. We will adapt static analysis methods, drawing on the (non-probabilistic) type systems for mobile calculi such as Godskesen et al. [14] and work on compositional game-theoretic approaches to model checking [2]. In developing these methods, we will emphasise algorithms for verification of quality-of-service properties for dynamically reconfigurable systems, including the efficiency of power management schemes, focussing mainly on online prediction methods. These are anticipated to involve machine learning based on simulation and instrumentation traces (WP4 and WP5). A particular challenge will result from online model checking for runtime evolution and reconfiguration (such as the on-the-fly service construction planned for Cityware).

We aim to derive theoretical frameworks, algorithms and prototype implementations. The work will be distributed according to expertise of researchers. This WP will be reinforced by an existing PhD student (Matthias Fruth at Birmingham) researching formal verification of system-level power management schemes for ubiquitous devices.

Deliverables: (D2.1) principles for the design of languages and specifications for ubiquitous software; (D2.2) a verification approach for the analysis of quality-of-service properties of ubiquitous systems; (D2.3) techniques and prototype software for model extraction and conformance checking; (D2.4) online analysis methods for reconfigurable systems and their implementation.

WP3: Testing

Lead Institution: UCL (Rosenblum)

Research Personnel: UCL PhD Student 1

Description: Our approach to testing ubiquitous computing systems will be built around standard processes used for system testing of software, with development of a representative suite of black-box test cases, execution of each test case on the system under test, identification of successful and failed test runs, and evaluation of the resulting test coverage to provide an indicator of test effectiveness. A test case is typically an initial vector of input values used to launch a test run, and an appropriate notion of coverage is applied depending on the specific purpose of the black-box test cases (functional testing, reliability testing, quality-of-service testing, etc.).

We will enhance this basic testing process in three ways. First, execution of a test case on the system under test must be enhanced to account for variation in context. We envision two possible approaches for this. In one approach, a representative set of context traces produced by simulation (WP4) will be used to produce a corresponding set of test runs for the test case. In the other approach, a range of context-dependent execution paths will be systematically explored within a single test run, in a manner similar in spirit to the operation of the Java PathFinder [16]. Both of these approaches require appropriate instrumentation of the system under test, and we will develop the requisite instrumentation tools as part of WP5.

Our second enhancement follows naturally from the first. In order to instrument the system under test in a way that allows systematic and controlled feeding of context values to a test run, we will develop a model of *context-awareness program points* (CAPPs), which characterise systematically the points in source code where the system receives and processes context information. Given that a ubiquitous system will typically run on top of a context-awareness middleware, identification of CAPPs during execution must use knowledge of the middleware API. It will be important to identify the CAPPs of both *synchronous* and *asynchronous* context awareness. The former correspond to direct invocations of context querying operations provided by the API; the latter correspond to callbacks implemented by the system under test and registered with the API.

Our third enhancement is to define appropriate test adequacy criteria for ubiquitous computing systems and to evaluate the criteria during testing. The use of context as a separate input driver for test execution naturally suggests that we develop an appropriate notion of *context coverage* to evaluate the extent to which the space of context variation is covered during testing. We will also define criteria for *mobility coverage* as a by-product of developing realistic mobility models for simulation (WP4) and for *security policy coverage* as a by-product of research on security properties (WP6). In addition to improving the suite of test cases, test coverage results can also be used to help guide and constrain state space exploration in model checking (WP2).

Deliverables: (D3.1) a method for testing ubiquitous systems, including a model of CAPPs; (D3.2) criteria for evaluating test adequacy of ubiquitous systems; (D3.3) a prototype test tool that implements the method and adequacy criteria.

WP4: Simulation

Lead Institution: UCL (Mascolo)

Research Personnel: UCL PhD Student 2

Description: Just as probabilistic model checking can provide useful feedback both prior to deployment and after deployment, so can simulation, which also offers a method of validation that is very close to implementation. To obtain more accurate characterisations of ubiquitous system properties than has been possible in the past, we will investigate how realistic mobility models based on social network theory can be devised and applied to the verification of typical ubiquitous computing scenarios. Repositories such as CRAWDAD [10], which contain traces of real movement, will be used to validate our models.

Prior to deployment of the system, these generic but quite realistic social mobility models can be used to drive a simulation of an application or synthetic execution scenario running on top of a simulation substrate. However this type of validation is still quite generic. Therefore, a second aim of this work package is to study how realistic mobility models can be derived from real traces of prototype-stage systems (through the instrumentation work done in WP5 and for the case study defined in WP7). This will allow more extensive validation of behaviour, as the traces generated by the simulation runs will then respect the patterns of mobility of the real traces but will offer sufficient variability for use in other validation activities such as testing (WP3).

Deliverables: (D4.1) mobility models based on social network theory; (D4.2) a prototype simulation tool that uses real execution traces to generate customised, application-dependent mobility models.

WP5: Instrumentation and Trace Generation

Lead Institution: Imperial (Dulay)

Research Personnel: Imperial and UCL RAs

Description: In this work package we will investigate and implement instrumentation methods needed by the other work packages and will explore techniques for automatically tailoring instrumentation tools to context-awareness middleware APIs (such as the use of simple languages for specifying metadata that characterise middleware API features), and aspect-oriented techniques for designing ubiquitous programs.

We will also investigate generic methods for analysing traces of ubiquitous system executions (from simulation runs and test executions) and apply statistical methods such as sampling-based probabilistic model checking or machine learning to derive procedures for calculating coverage estimates.

More specific use of this work is described in other work packages.

Deliverables: (D5.1) instrumentation methods for the tools developed in other work packages; (D5.2) prototype implementation of instrumentation for testing; (D5.3) prototype implementation of model extraction for probabilistic model checking.

WP6: Security Properties

Lead Institution: Birmingham (Ryan)

Research Personnel: Birmingham PhD Student 2

Description: There are distinct security risks associated with ubiquitous computing, arising from new one-off contexts being formed through mobility and ad hoc networking, for example the Bluetooth mobile phone security loopholes mentioned earlier. This work package will address this significant challenge by developing techniques for validation of security requirements.

First, drawing on typical Cityware scenarios we will develop languages for expressing access control policies that allow system designers to ensure that legitimate users have flexible access to the data and resources they require, while barring other users from illegitimate access. We will extend our existing access control language [15] and its associated model checker with constructs for handling mobility and context awareness, via the use of context-awareness program points (WP3). In particular, the diagnostic traces obtained from model checking will generate input into testing of security policies (WP3).

In another direction, we will focus on new protocol standards necessitated by smart urban spaces, for example device authentication, service creation, access to data and resources, etc. User goals can be subtle and apparently conflicting; for example, a user may wish to identify himself to a remote server in order to gain access to protected data or resources, while preserving anonymity to a local untrusted device mediating the access. We will develop languages for expressing such requirements and techniques for their verification, based on existing frameworks such as ProVerif which has already been successful for protocol analysis in the static case [29, 1]. We will implement those techniques and evaluate them on Cityware protocols (WP1), as well as other high-security protocols provided by our collaborators. In addition, we will design and test new protocols for smart urban spaces (WP7).

Deliverables: (D6.1) access control verification algorithms for context-aware mobile systems; (D6.2) prototype security protocol validation techniques for mobile systems based on ProVerif; (D6.3) exemplar protocols for security and privacy.

WP7: Evaluation Studies

Lead Institution: Imperial (Dulay, Lupu)

Research Personnel: Imperial RA

Description: A crucial part of the project will be to apply Ubival's techniques and tools to realistic ubiquitous applications, particularly those developed in Cityware and by our industrial partners. This will allow evaluation of Ubival's techniques and tools, provide feedback for refinement of the techniques, act as exemplars that others can use, and adapt and compare to their own validation techniques. The work package will develop a methodology to test the effectiveness of using the validation techniques and tools for ubiquitous applications. We also hope to liaise with the DTI's UbiCare centre in order to validate ubiquitous healthcare systems.

We will also evaluate the developed techniques on applications provided by our industrial partners. Success will be measured by applying the techniques to validate selected Cityware software components, predicting their performance and other behavioural characteristics, and evaluating the predictions against testbed experiments.

Deliverables: (D7.1) a range of evaluation studies applying UbiVal tools and techniques to Cityware applications; (D7.2) other such evaluation studies applied to applications identified by our industrial partners.

5. Project Management

All three institutions have extensive experience with managing and working in large, multi-institution collaborative projects. The project with be managed by a management committee comprising Prof Rosenblum from UCL, Prof Kwiatkowska from Birmingham, and Dr Dulay from Imperial. There will be an initial project kickoff meeting followed by semi-annual review meetings (alternating between Birmingham and London) on project progress and direction, to which industrial participants will be invited, with additional ad hoc meetings between the researchers as necessary. The staff will hold regular teleconferences to discuss technical and management matters. There will be annual meetings with the Cityware project, and staff will travel to Bath to participate in testbed experiments. To encourage integration and familiarisation with ubiquitous systems technology, students will travel on extended visits to Bath and partner institutions.

6. Relevance to Beneficiaries

This proposal fully addresses the criteria specified in the WINES II call and thus directly benefits the scientific and technical objectives of the EPSRC:

- It directly addresses the challenge of software development for ubiquitous systems, recognising the central role of context-awareness and mobility and the need for novel approaches to validation.
- It addresses the areas of wireless communications; programming and design tools; context-awareness; and trust, security and privacy.
- It is much larger in size, scope and ambition than traditional responsive-mode EPSRC proposals.
- It will be carried out by a newly formed consortium of three groups providing highly complementary expertise.
- It will involve multi-disciplinary research across many areas, including software engineering, networking, mobile computing, formal methods, systems theory, social network theory, and urban architecture.

- It is highly adventurous, because rather than seeking incremental improvements to one class of existing methods or tools, its ambition is to discover and establish a comprehensive, sound, scientific basis for engineering a new and important class of systems.

- Its results will be applied in a significant application domain, namely pervasive urban spaces.

- It will provide a significant bridge to and cross-fertilisation with existing WINES projects through direct collaboration with the Cityware project (with the investigators from Imperial providing coordination with Cityware), and via additional dissemination plans described in Section 7. It will also provide a bridge with the DTI Next Wave, through FORWARD and Mobile Bristol.

If successful, the research will establish scientific principles that will raise the understanding of ubiquitous systems, improve their quality and robustness, reduce their development cost, and (for embedded software) reduce the cost of recall in case of faults. Additionally, the project will establish a model for additional future cross-project collaborations.

Beyond the needs of the WINES initiative, the results of this project will benefit UK science and industry in their ability to build and validate ubiquitous computing systems. In particular, development environments for ubiquitous software will be in demand by mobile phone and wearable devices manufacturers in the UK, and by companies engineering sensor network systems for home, work and healthcare scenarios. Device manufacturers have identified a pressing need for validation technology for such systems. Our validation approaches and accompanying tools will thus enhance engineers' ability to build and validate robust and efficient ubiquitous systems with greater confidence. Ultimately this will further establish the UK as a leader in technology research and development for ubiquitous systems.

7. Dissemination and Exploitation

The investigators have long and successful track records of publication in top-tier conferences and journals and dissemination of popular software prototypes. Furthermore, our goal of developing engineering foundations for ubiquitous computing means that our project results should be of significant interest to other WINES projects.

We anticipate publishing the results of the project at the premier domestic and international software engineering and ubiquitous computing meetings, such as the International Conference on Ubiquitous Computing (UbiComp), the International Conference on Mobile Communications and Networks (Mobicom), the International Conference on Software Engineering (ICSE), and the Conference on Computer-Aided Verification (CAV). Likely venues for journal publication include the IEEE Transactions on Software Engineering (IEEE TSE), and the IEEE Transactions on Mobile Computing. All provide an excellent forum for presentation and discussion of this work.

We will offer tutorials and software demos for Summer Schools and for the conferences mentioned above; for instance, PRISM will be taught at the FIRST Summer School in Copenhagen in 2006. We also will continue to feature a range of techniques and tools (e.g. J2ME/Bluetooth phone programming, PRISM modelling) in undergraduate and postgraduate teaching at our own institutions, thus enhancing the portfolio of case studies and the degree of relevant experience. We will further disseminate through our industrial partners.

All prototypes, models, statistics and videos created by the project will be made freely available through the Web, for use within the WINES initiative and by the industrial and scientific communities at large. Our partnership with the Cityware project is one avenue of dissemination and exploitation already discussed elsewhere in this proposal. In order to encourage and facilitate the uptake of UbiVal project results within the broader WINES initiative, the investigators will approach other WINES projects for invitations to attend their meetings and to brief them on UbiVal work and results. Any additional collaboration resulting from such exploratory visits would need to be undertaken at the expense of the other projects.

8. International Collaborations

In order to develop additional opportunities for dissemination of results and creation of follow-on research projects, and to increase the visibility of the UK as a leader in technology research and development for ubiquitous computing, we intend to host short- and medium-term visits by some of our international colleagues during the project.

We anticipate that Prof Sebastian Elbaum of the University of Nebraska–Lincoln will visit UCL during his upcoming sabbatical year, and that Prof Rosenblum will visit Prof Elbaum twice during the project. Prof Elbaum is a leading expert in software testing, instrumentation and monitoring and specialises in controlled experiments with testing and monitoring techniques. He and Prof Rosenblum have initiated a collaboration on testing ubiquitous systems, and their initial ideas are reflected in the description of WP3. A 2-page CV for Prof Elbaum is provided as an attachment to this Case for Support.

For WP2, and WP7, we have continuing research collaborations with a number of institutions, such as Bonn (Baier), Verona (Segala), Aachen (Katoen), Urbana-Champagne (Sanders), Texas (Shmatikov), CMU (Younes) and Rice University (Vardi), and we are planning brief mutual visits with them. We have established an ARC-funded exchange programme with NICTA (Professor Carroll Morgan and Dr Annabelle McIver), who are embarking on analysing sensor network protocols with PRISM, and we are requesting Travel and subsistence for a visit to allow for the collaboration to continue beyond the current level of funding. For WP4 we have collaborations with Aachen University (Mahonen) and Oslo University (Plagemann), who are interested in similar realistic mobility models.

9. Justification of Resources

The project has been budgeted using Full Economic Costing in full compliance with EPSRC guidelines, including budgeting with un-inflated personnel costs and with separate budgeting of VAT and base costs.

The scope and ambition of the project require a corresponding investment in human resources. We request one post-doc RA for each of the three project sites and two PhD students each for Birmingham and UCL, with the work allocated to these researchers as described in Section 4. The PhD studentships are for the full duration of the project (3.5 years) at institutional rates, to allow for in-depth familiarisation of the Cityware technology, to alleviate the possible risk of being unable to perform testbed experiments in view of their deployment plans, and to allow time for writing up theses. We are fortunate to be in a position to name an experienced RA (Dr Gethin Norman) to assist with the work on WP2, and we request a salary commensurate with his substantial experience. Dr Norman will focus on using and extending the probabilistic model checking techniques to ubiquitous computing systems, and already has relevant experience via the EPSRC project GR/S46727 concerning probabilistic model checking of mobile ad hoc network protocols on which he is currently employed. A proposal to further develop the core probabilistic model checking technology and the PRISM tool, which is distinct yet synergetic with this research, will be submitted shortly in responsive mode. Salaries at UCL and Imperial are budgeted to include the standard London Allowance, and both sites request funds for recruitment of their RAs.

The investigators will spend their time allocations managing work packages and deliverables, monitoring the research work to ensure its consistency with project goals, contributing directly to the research work, doing background reading and investigation, defining tasks and milestones for the project, participating in UbiVal meetings and teleconferences and meetings of other WINES projects, and collaborating with other domestic and international research colleagues in support of the project goals. The percentage allocations of investigator time to the project reflect both estimates of actual time that will be spent each week working on the project as well as expectations about the relative amount of effort each investigator will put into the project. As project director, Prof Rosenblum assumes a significant project management burden in ensuring administrative and technical cohesion of the project across the three sites and thus will contribute 15% of his time to the project. Prof Kwiatkowska will contribute 13% of her time to coordinate the significant Birmingham involvement in the project and to serve as a member of the project management committee. Drs Ryan and Ghica, and Drs Dulay and Lupu, will contribute 8% and 5% respectively, for supervision of the research work at their sites. Dr Mascolo will contribute 5% of her time but incurs zero overhead because she will be working under an EPSRC Advanced Research Fellowship throughout the project.

We request travel funds for semi-annual project meetings, for travel to UK/EU and overseas conferences, and for travel to other WINES project sites. Each investigator, PhD student and RA will travel by rail once each year from London to Birmingham or vice versa for a two-day project meeting with overnight hotel stay. Each site requests funds for three person-trips per year to UK/EU conferences and three person-trips per year to overseas conferences. And each site requests funds for three overnight person-trips per year to Bath for Cityware project meetings, and two overnight person-trips per year to other WINES project sites. To facilitate his collaboration with Prof Sebastian Elbuam described in Section 8, Prof Rosenblum requests funds for two visits to UNL during the project, and round-trip airfare and 6 months of bench fees for Prof Elbaum's sabbatical visit to UCL. Birmingham request similar funds for the collaborative visits described in Section 8.

The project will require equipment both for the everyday work of the personnel, for simulation experiments, and for demonstration and evaluation activities with ubiquitous computing systems. We request one high-end laptop machine for each investigator, PhD student and RA. UCL request funds for a server to conduct simulation studies. Birmingham will use an existing PC for benchmarking and a SRIF-2 purchased cluster for compute-intensive simulations and verification. Each site requests funds for purchase of PDAs, smart mobile phones, sensors, and wearable tags, plus monthly charges for GPRS/-MMS/SMS/UMTS wireless communication between devices. To mitigate the risks associated with this significant investment in computing equipment, each site requests funds for equipment maintenance.

The project will require purchase of software packages, books, manuals, computer printing allocations, archival storage space, and each site requests annual funds for this purpose.

References

[1] M. Abadi and B. Blanchet. Analyzing Security Protocols with Secrecy Types and Logic Programs. *Journal of the ACM*, 52(1):102–146, Jan. 2005.

[2] S. Abramsky, D. R. Ghica, A. S. Murawski, and C.-H. L. Ong. Applying game semantics to compositional software modeling and verification. In *Proc. TACAS'04*, pages 421–435, 2004.

[3] C. Baier, E. Clarke, V. Hartonas-Garmhausen, M. Kwiatkowska, and M. Ryan. Symbolic model checking for probabilistic processes. In *Proc. ICALP'97*, volume 1256 of *LNCS*, pages 430–440. Springer, 1997.

[4] T. Ball, B. Cook, V. Levin, and S. K. Rajamani. SLAM and Static Driver Verifier: Technology transfer of formal methods inside Microsoft. Technical Report MSR-TR-2004-08, Microsoft Research Cambridge, 2004.

[5] J. Bible, G. Rothermel, and D. S. Rosenblum. A comparative study of coarse- and fine-grained safe regression test selection techniques. *ACM Transactions on Software Engineering and Methodology*, 10(2):149–183, Apr. 2001.

[6] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless Communication and Mobile Computing Special Issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, 2(5):483–502, 2002.

[7] L. Capra, W. Emmerich, and C. Mascolo. Carisma: Context-aware reflective middleware system for mobile applications. *IEEE Transactions on Software Engineering*, 29(10):929–945, 2003.

[8] Y.-F. Chen, D. S. Rosenblum, and K.-P. Vo. Testtube: A system for selective regression testing. In *Proc. 16th International Conference on Software Engineering*, pages 211–220, Sorrento, Italy, May 1994.

[9] P. Ciancarini, F. Franzé, and C. Mascolo. Using a coordination language to specify and analyze systems containing mobile components. *ACM Transactions on Software Engineering and Methodology*, 9(2):167–198, Apr. 2000.

[10] CRAWDAD web site. http://crawdad.cs.dartmouth.edu/.

[11] A. K. Dey, G. D. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16:97–166, 2001.

[12] A. Dimovski, D. R. Ghica, and R. Lazic. Data-abstraction refinement: A game semantic approach. In *SAS*, pages 102–117, 2005.

[13] M. Duflot, M. Kwiatkowska, G. Norman, and D. Parker. A formal analysis of Bluetooth device discovery. *International Journal on Software Tools for Technology Transfer (STTT)*, 2006. To appear.

[14] J. C. Godskesen, T. T. Hildebrandt, and V. Sassone. A calculus of mobile resources. In *CONCUR'02*, volume 2421 of *LNCS*, pages 272–287. Springer, 2002.

[15] D. P. Guelev, M. D. Ryan, and P.-Y. Schobbens. Model-checking access control policies. In *Seventh Information Security Conference (ISC04*. LNCS, Springer-Verlag, 2004.

[16] K. Havelund and T. Pressburger. Model checking java programs using java pathfinder. *International Journal on Software Tools for Technology Transfer*, 2(4), Apr. 2000.

[17] M. R. A. Huth and M. Ryan. *Logic in computer science: modelling and reasoning about systems*. Cambridge University Press, New York, NY, USA, 2000.

[18] A. Jardosh, E. M. Belding-Royer, K. C. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. In *Proceedings of MobiCom'03*, San Diego, California, USA, September 2003.

[19] S. Kremer and M. D. Ryan. Analysis of an electronic voting protocol in the applied pi-calculus. In *Proc. 14th European Symposium on Programming (ESOP'05)*, volume 3444 of *LNCS*, pages 186–200, Edinburgh, U.K., 2005. Springer.

[20] M. Kwiatkowska. Model checking for probability and time: From theory to practice. In *Proc. 18th Annual IEEE Symposium on Logic in Computer Science (LICS'03)*, pages 351–360. IEEE Computer Society Press, 2003.

[21] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 2.0: A tool for probabilistic model checking. In *Proc. QEST'04*, pages 322–323. IEEE Computer Society Press, 2004.

[22] M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic model checking in practice: Case studies with PRISM. *ACM SIGMETRICS Performance Evaluation Review*, 32(4):16–21, 2005.

[23] M. Kwiatkowska, G. Norman, and D. Parker. Quantitative analysis with the probabilistic model checker PRISM. In *Proc. QAPL'05*. Electronic Notes in Theoretical Computer Science, 2005. To appear.

[24] M. Kwiatkowska, G. Norman, D. Parker, and J. Sproston. Performance analysis of probabilistic timed automata using digital clocks. In *Proc. Formal Modeling and Analysis of Timed Systems (FORMATS'03)*, volume 2791 of *LNCS*, pages 105–120. Springer-Verlag, 2003.

[25] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282:101–150, 2002.

[26] Z. Liu, M. Kwiatkowska, and C. Constantinou. A biologically inspired QoS routing algorithm for mobile ad hoc networks. In *Proc. 19th International Conference on Advanced Information Networking and Applications (AINA 2005)*, pages 426–431. IEEE CS Press, 2005.

[27] K. Maeda, K. Sato, K. Konishi, A. Yamasaki, A. Uchiyama, H. Yamaguchi, K. Yasumotoy, and T. Higashino. Getting urban pedestrian flow from simple observation: Realistic mobility generation in wireless network simulation. In *Proceedings of International Workshop on Modeling Analysis and Simulation of Wireless and Mobile Systems*, September 2005.

[28] R. Milner. Theories for the global ubiquitous computer. In *Proc. FoSSaCS'04*, volume 2987 of *LNCS*, pages 5–11. Springer, 2004.

[29] A. Mukhamedov and M. D. Ryan. On anonymity with identity escrow. In *Third international Workshop on Formal Aspects in Security and Trust (FAST'05)*, LNCS. Springer, 2005.

[30] M. Musolesi, S. Hailes, and C. Mascolo. An ad hoc mobility model founded on social network theory. In *Proceedings of the 7th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2004)*, October 2004.

[31] G. Norman, D. Parker, M. Kwiatkowska, S. Shukla, and R. Gupta. Using probabilistic model checking for dynamic power management. *Formal Aspects of Computing*, 17(2):160–176, 2005.

[32] A. Orso, H. Do, M. J. Harrold, G. Rothermel, and D. Rosenblum. Using component metadata to regression test component-based software. *Software Testing, Verification & Reliability*, to appear 2006.

[33] A. Orso, M. J. Harrold, D. Rosenblum, G. Rothermel, M. L. Soffa, and H. Do. Using component metacontent to support the regression testing of component-based software. In *Proc. IEEE International Conference on Software Maintenance 2001*, pages 716–725, Florence, Italy, Nov. 2001. IEEE Computer Society Press.

[34] A. Podgurski and C. Yang. Partition testing, stratified sampling, and cluster analysis. In *Proc. 1st ACM SIGSOFT Symposium on Foundations of Software Engineering*, pages 169–181, Redondo Beach, CA, USA, Dec. 1993. ACM.

[35] PRISM web site. www.cs.bham.ac.uk/˜dxp/prism.

[36] D. S. Rosenblum. Towards a method of programming with assertions. In *Proc. 14th International Conference on Software Engineering*, pages 92–104, Melbourne, Victoria, Australia, May 1992. ACM.

[37] D. S. Rosenblum and E. J. Weyuker. Using coverage information to predict the cost-effectiveness of regression testing strategies. *IEEE Transactions on Software Engineering*, 23(3):146–156, Mar. 1997.

[38] J. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems,* P. Panangaden and F. van Breugel (eds.), volume 23 of *CRM Monograph Series*. American Mathematical Society, 2004.

[39] V. Shmatikov. Probabilistic model checking of an anonymity system. *Journal of Computer Security*, 2003.

[40] T. Tse, S. Yau, W. Chan, H. Lu, and T. Chen. Testing context-sensitive middleware-based software applications. In *Proceedings of the 28th Annual International Computer Software and Applications Conference*, pages 458–466. IEEE Computer Society Press, Sept. 2004.

[41] J. Yoon, M. Liu, and B. Noble. Random waypoint considered harmful. In *Proc. INFOCOM'03*. IEEE, 2003.

[42] H. Younes, M. Kwiatkowska, G. Norman, and D. Parker. Numerical vs. statistical probabilistic model checking. *International Journal on Software Tools for Technology Transfer (STTT)*, 2005. To appear.

Appendix A. Project GANTT Chart

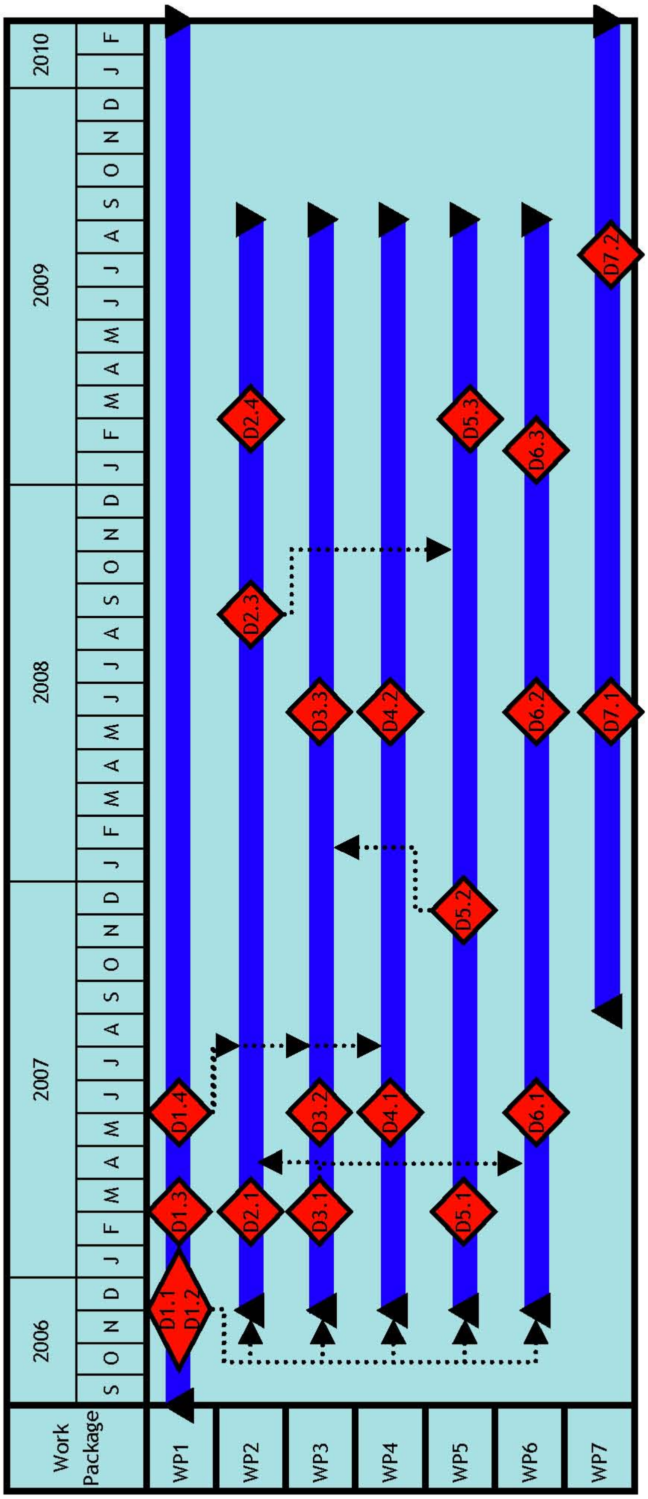
Work Package
WP1
WP2
WP3
WP4
WP5
WP6
WP7
S O N D J F M A M J J A S O N D J F M A M J J A S O N D J F M A M J J A S O N D J F
D1.1
D1.2
D1.3
D1.4
D2.1
D2.3
D2.4
D3.1
D3.2
D3.3
D4.1
D4.2
D5.1
D5.2
D5.3
D6.1
D6.2
D6.3
D7.1
D7.2