Design, Implementation and Adaptation of Sensor Networks through Multi-dimensional Co-design (DIAS-MC)

1 Investigators' Track Records

This ambitious project demands a multi-disciplinary team of researchers with the necessary skills. We have formed a consortium along these lines, constituted by eight partners that bring together the necessary expertise and potential resources to match the scientific/technological challenges of the project. Most importantly, the project partners cover the research areas needed in the project. The organizations and individuals involved are:

Organization	Individual	Expertise
U of Glasgow	Prof J Sventek (P)	network measurement and monitoring
	Dr A Miller	formal methods
U of Kent	Prof I Marshall (P)	embedded operating systems and networks, DTI Envisense programme
	Mr. P Lee	electronic systems design
U of Manchester	Dr A A A Fernandes (P)	distributed data management and query processing
	Prof N Paton	distributed data management and query processing
	Dr S Boult	hydrogeochemistry and deployment of environmental sensor systems
U of St Andrews	Prof R Morrison (P)	generative programming and adaptive systems
	Prof A Dearle	generative programming and adaptive systems
	Dr G Kirby	generative programming and adaptive systems
	Dr D Balasubramaniam	generative programming and adaptive systems
U of Strathclyde	Prof J Dunlop (P)	radio communications
Orange	Dr S Hope	wireless GPRS/3G service provider in the UK
Intelisys	Dr N Boyd	SME deploying and operating water catchment environmental sensors
Xilinx	Dr. C Carruthers	field programmable logic solutions for use in embedded systems
STW	Mr R Borrill	manage upland catcchments for drinking water collection

Joseph Sventek is Professor of Communication Systems in the Department of Computing Science at the University of Glasgow and leader of the Embedded, Networked, and Distributed Systems research group. Prior to joining Glasgow, he had a distinguished career in distributed systems research and managing research teams at LBNL (1979-1986), HP Labs (1987-1999), and Agilent Labs (1999-2002). His research interests include programmable networks, embedded systems, closed-loop network management, and distributed system architectures. He is the PI for two EPSRC grants: AMUSE: Autonomic Management of Ubiquitous Systems for e-Health (GR/S68040/01) and Performance Measurement and Management for Two-Level Optimization of Networks and Peer-to-Peer Applications (GR/S68989/01).

Dharini Balasubramaniam is a Lecturer at the University of St Andrews, with research interests in type theory, generative programming compliant programming languages and software architecture. Dr. Balasubramaniam and Prof. Morrison are partners in the ArchWare EU project (IST-2001-32360) on formal descriptions of reflective software architectures using the π-calculus.

Stephen Boult has received funding from NERC, BBSRC and industrial partners to carry out investigations in hydrogeochemistry through a combination of laboratory experiment and field monitoring. Continuous monitoring of various aquatic environments has been central to this research. In 1996, Dr Boult founded Intelisys Ltd, as a research-led collaborative venture between electronic engineers and environmental scientists. Intelisys designed the stand-alone equipment installed since 2002 in the upland catchment that provides an ideal testbed for distributed sensor networks in this project. Dr Boult is thus able to contribute experience of environment-specific data requirements and sensor design.

Alan Dearle is a Professor of Computer Science at the University of St Andrews. His special interests are in distributed computing, reflective and autonomic systems and component technology. He is the PI on the EPSRC CINGAL and RAFTA projects on thin servers and distributed deployment..

John Dunlop is Professor of Electronic Systems Engineering and Head of the Institute for Communications and Signal Processing, University of Strathclyde. He has been involved in research programmes in communication systems and electronic systems engineering for more than two decades, including RACE Mobile Communications Project (R1043) and RACE Advanced Time Division Multiple Access ATDMA (R2084). He is a full academic member of the UK Virtual Centre of Excellence in Mobile and Personal Communications (and was a director of the Centre from 1999 to 2001). He has held several UK Engineering and Physical Sciences Research Council (EPSRC) awards on local area communications, underwater communications and mobile communications and holds contracts from the Mobile VCE core 3 covering work in the areas of Personal Distributed Environment and Inter-working of Networks.

Alvaro A. A. Fernandes is a Lecturer in Information Systems. He is a co-investigator in an EPSRC grant on adaptive query processing for the grid, in the myGrid e-Science pilot project, and on a DTI grant on high-level grid middleware for data management. His principal research interests are in query optimization and processing. With Prof Norman Paton, he has recently been actively involved in research leading to the development of open-source adaptive, distributed query processing systems in service-based grids. With Prof. Paton, he brings expertise on building well-founded query-processing engines for a range of deployment platforms, underlying data models, and cost parameters.

Graham Kirby is a Senior Lecturer in Computer Science at the University of St Andrews with research interests in programming languages, mobile and distributed computing, system evolution and automated system configuration. He currently holds EPSRC grant GR/S44501/01 on Secure Location-Independent Autonomic Storage Architectures.

Peter Lee joined the University of Kent in 1991 having worked for over 10 years as a VLSI designer at Philips, ITT and AMS (Austria Mikro Systeme) where he was involved in the design of numerous digital and mixed signal ASICs. At Kent he has worked on a number of EPSRC research projects, including: Intelligent Optical Position Sensitive Detectors using CMOS Technologies (GR/J/75883) and A Flexible Design Environment for Modular Image Recognition Development (GR/J/70253). He is the director of the Electronics Systems Design Centre (ESDC), established at the University of Kent to provide support for local SMEs through the DTI sponsored Microelectronics in Business (MiB) and Electronics Design (ED) Programmes. ESDC has been particularly successful in encouraging companies to use new electronic technologies and system design tools for product development.

Ian W Marshall is Professor of Distributed Systems in the Computing Laboratory at the University of Kent and a Visiting Professor in the Electrical Engineering Department at University College London. He is also Technical Director of the DTI funded Envisense research centre and leader of the SECOAS project, which has deployed an intelligent sensor network at the Scroby Sands windfarm site. From 2001-2003, he was a Royal Society Industry Fellow at University College London where he led the initial research on self-organising sensor networks. Previously he worked for BT in intelligent management of communications infrastructures including; active networks & services, optical networks, broadband networks, Internetworks and distributed enterprise systems. He also led the Eurescom funded project CASPIAN, the FP5 project ANDROID, and the BT funded Alpine and MMN projects. All of these projects focused on automated adaptation and management issues.

Alice Miller is a recently appointed lecturer and an ex Daphne Jackson research fellow. Her expertise is in model checking concurrent systems (for feature interaction analysis of telecommunications systems and investigation of the IEEE FireWire protocol, for example). She has developed important mathematical techniques based on abstraction, induction and symmetry reduction to extend the applicability of model checking techniques to large or infinite systems.

Ron Morrison is Professor of Software Engineering at the University of St Andrews, where he leads the Software Architecture Group. He was a director of the Alvey/SERC-funded PISA project, the ESPRIT BRAs $FIDE_1$, 3070, and $FIDE_2$, 6309, co-ordinator for the EC/NSF collaboration grant, ECUS 006:9839, and the Pastel ESPRIT Working Group EP22552. He currently leads the UK Network of Excellence on Distributed Information Management, DIMNET GR/R85907. His special research interests are programming language design, persistent object systems, reflective computing, evaluation of complex systems, distributed garbage collection, and compliant systems architectures.

Norman W. Paton is a Professor of Computer Science. His current activities include developing distributed and service-based query processing technologies for the grid in the OGSA-DAI and ^{my}Grid projects, and designing and implementating spatio-temporal data models and query languages in the Tripod project. He is a principal investigator in grants that address adaptive query processing for the grid and high-level grid middleware for data management. He is co-chair of the Database Access and Integration Services Working Group of the Global Grid Forum. His expertise in the development of query- processing engines will be crucial for application data management aspects of the project.

2 Motivation

Sensor networks are becoming increasingly prevalent in large-scale system monitoring; example applications include city pollution (environmental), communication network quality of service (industrial), and automotive (transportation). The data produced by the deployed sensors is monitored for behavior of the sensed system that is consistent with normal operation; the data can also be used to drive closed-loop control of the sensed system (autonomic management). The sensors can produce large volumes of data; in many contexts, it is essential to save the data for subsequent retrieval.

Creating an effective sensing system requires skills from many independent disciplines. With increasingly flexible hardware structures in embedded systems, there is particular interest in the "optimal" partitioning of functionality across the hardware/software boundary; this approach to optimizing hw/sw concerns in tandem is termed "co-design". Integrated design of such systems must bridge the semantic gaps between independent disciplines while dealing with the complexities inherent in typical, sensed environments – e.g. high expectation of failure, complex statistics, poorly-understood mechanisms.

We will focus on sensor networks that consist of small, battery-powered, wireless sensors that may be tethered or mobile [1]. For such systems, we need to minimize power consumption while delivering full functionality. By contrast, most of the experience in the distributed systems community is with mains-powered, tethered systems; such systems are usually engineered to take full advantage of indirection and layering to achieve the functionality and quality attributes required by a particular distributed application, at the cost of power, computational, and/or communications resources.

In addition, maintainability issues arise for sensors deployed in hostile (to humans) environments (underwater, wide geographic distribution over a watershed, etc.). Such environments may demand that the sensors possess autonomic capabilities while awaiting maintenance, be accepting of control commands from orchestrating components, be supportive to peer sensors in forwarding communications, etc.

One of the primary challenges in the design of effective sensing systems is the sheer number of dimensions (also termed "areas of concern") that must be jointly considered and seamlessly integrated. The major dimensions of sensor systems are: application, monitoring and control, communication network, and operating system. For wireless sensor networks, the network dimension is further subdivided into communication protocols and radio capabilities.

So far, the practice has been to address each of these dimensions orthogonally (independently), with the result that any potential synergies available through co-design are not achieved. Designers are then forced to employ one-off, *ad hoc*, design changes to customize for scarce resource consumption (e.g. battery power) and/or critical deployment requirements (e.g. maintainability).

Co-design is a methodology that integrates orthogonal design methods; the goal of co-design is that the integrated design will better meet system requirements than the sum of the orthogonal designs. We intend to use co-design to integrate the design of the hardware system and ***all*** dimensions of the software system, across the hw/sw boundary. Hw/sw co-design is a hot topic in the embedded systems community, especially for sensor nodes [2]. Such co-design is usually targeted at systems consisting of a general-purpose processor (CPU), a dynamically reconfigurable datapath (e.g. an FPGA), and a memory hierarchy. Many research design tools have been built to address co-design of such systems [3].

Aspect-oriented system design (AOSD) is one approach to software co-design [4]. In AOSD, one specifies the required features in different areas of concern; a tool then weaves these different threads into a solution to the overall problem. While the output of AOSD can be monolithic code, recent research [5] has focused on maintaining a level of modularity in the resulting solution, thus enabling incremental changes to the solution, potentially at runtime. AOSD is typically applied to middleware; there are efforts to apply it to more time-critical software, e.g. operating systems [6].

Our experience in building distributed and sensor systems leads us to believe that the time is ripe for attempting to integrate orthogonal design methods with more traditional hw/sw co-design methods in order to achieve synergistic gains. The end result should be the ability to design, implement, deploy, and adapt sensor systems that are optimized to minimize a much more complex and realistic combined cost function.

3 Approach

We propose to enable a designer to specify the requirements for a particular sensor system in each area of concern orthogonally, and to then use co-design tools to produce integrated designs and implementations that achieve substantial synergy and better reflect the underlying physical environment of the sensed system. This work will extend the notion of hw/sw co-design into multiple dimensions through the application of existing techniques, such as AOSD methods, i.e. the final system design is generated by weaving together the individual components with appropriate software/firmware "glue" to achieve the requisite integration and reap synergistic benefits. This generation phase is optimized with respect to a combined cost function (e.g. aggregate data early but minimize power consumption during operation) for the target system. Iterative specification/integration will most likely be required to maximize synergy.

In order to achieve these goals, we will: 1) formalize each dimension, and define one or more languages to be used for specification of requirements along each dimension; 2) develop mechanisms to formally validate specifications in each such language; 3) build a tool that weaves such orthogonally-produced specifications into one or more designs for the overall system; 4) develop mechanisms to formally validate overall system designs; 5) develop mechanisms to evaluate these designs in terms of the chosen, global, cost function in order to determine the optimal (or near-optimal) design; 6) construct an implementation platform upon which optimal designs can be constructed and deployed; and 7) build two prototype systems using this platform to validate the methodology.

Anticipated outcomes for the project are: specification languages for each dimension, tools for producing integrated designs from orthogonal designs, formal validation mechanisms for each dimension and for integrated system specifications, tools for selecting optimal designs given a global cost function, mechanisms for managing a running system based upon the global cost function, and two prototype systems designed, built, deployed, and managed using the methodology and tools defined above.

4 Research Objectives

The overall objective of the project is to employ generative programming techniques to construct sensor systems that are optimal with respect to a chosen, global cost function, formally validated with respect to required system properties, ***and*** adaptive to changing conditions in the field. This is to be achieved through formal specification of the projection of the system onto each orthogonal area of concern and the specification of each area's contribution to the global cost. Generative programming techniques are used to weave these orthogonal specifications into an overall system specification/design, leveraging any synergies between areas of concern without violating the required system properties. The weaving technique used will maintain some level of modularity in the resulting design such that the system is able to adapt to environmental variations after deployment. Achievement of the overall objective will require investigations in a number of research areas; these are described below. Note that the research described in each area contributes to the overall system specification, validation, and adaptation objective.

4.1 Radio Communications

Recent wireless communications research has focused on the inter-working of heterogeneous networks to achieve ubiquitous service provision [18]. Such a concept is relatively new to sensor networks, and this area of research will involve understanding the exploitation of heterogeneous wireless networks to minimize the cost function(s) associated with wireless transmission of large amounts of data from both pseudo static and mobile sensor nodes. Of particular interest will be the optimization of uplink alternatives, which will allow both synchronized and unsynchronized data transmissions from a large-scale deployment of sensor devices with limited energy. The overall objective is to enable self configuration of sensor networks using the heterogeneous wireless capability, where available, and to consider the dynamic aspects of creation, dispersal and interaction of clusters of devices, as specific sensor networks change in constitution as a function of time. Critical issues that will be addressed include security, data integrity and energy conservation.

The heterogeneous network advantages will be demonstrated by choosing existing standard interfaces such as 802.11 and 802.15 with appropriate modifications to the transmission stack to demonstrate a dynamically reconfigurable network based upon two heterogeneous air interface examples.

4.2 Networking and Operating Systems

The research community has significant experience with layered networking structures. Layering enables isolation of areas of networking concern, and the ability to focus on interactions of these concerns at the layer boundaries. Such a layered approach is inappropriate to the specialized communication environments used in sensor systems. The design may be based upon layers, but the deployed system should flatten the communication structure in such a way as to optimize the global cost function. We will seek a way of formalizing the networking requirements that meets this optimization goal ***and*** enables co-design between this area of concern and the others, especially the radio communication area.

Other groups are active in the formal specification and analysis of network protocols [14]. Kent has an existing collaboration with the group at the University of Western Australia, and expects to be able to leverage the formal work of that group into the language formalization activities associated with this area of concern.

Individual sensors are resource-constrained embedded systems. The operating system used in such sensors must support parallel activities consistent with the sensor's resource constraints, both structural (cycles, memory) and operational (energy conservation). The system must also enable the sensor to adapt to changing conditions in the environment. Ultimately, each sensor is responsible for making decisions about how best to achieve it goals, subject to policy constraints defined when deployed, and current environmental conditions. The appropriate structures needed to satisfy this responsibility are an important research goal of the project.

The Kent group has trialed [17] a range of prototypes of model-free, adaptive control algorithms that adjust the schedules of measurement, aggregation, logging and communication, of both the data and metadata associated with networked sensors. In this project, we will focus particularly on the metadata aspects (which are essential to ensuring the correct interpretation and validity of the data collected) and on dynamic upgrades of the node software. Both of these capabilities have been identified as crucial [20], but neither has been well supported in any prototype.

4.3 Network Management and Monitoring

Sensor networks are large, distributed systems; as such, management functions are required to configure, reconfigure, and adapt the system during its deployed lifetime. These functions are typically built into the system to enable autonomic operation; it must be possible for management responsibility to transfer to an appropriately authorized system when such control is asserted in the proximity of the sensor network.

Typically, these management functions are constructed over the networking and operating system areas of concern, and designed independently of the structure of those components. As such, potential synergies with these layers are not exploited. Glasgow has experience in co-design of management and monitoring functionality with application and network functions for wireline sensor networks [19]. This experience demonstrated that significant improvements in the manageability of a distributed application can be achieved by piggy-backing configuration information with the data streams produced by sensors and by exploiting multicast capabilities in the communication network; this eliminates any management-specific traffic from the sensors to the management controller.

Deployed sensor systems are, by necessity, autonomic. The components must take actions to maintain the system functionality in the face of changes to the environment. Occasionally, the operator of a sensor system visits the site and wishes to temporarily take control of the system to diagnose the system or to reconfigure the system. In such a situation, the sensor system makes itself subservient to the operator's system (a laptop or a PDA) for the duration of the activity. Glasgow is exploring such federations among autonomous systems in the AMUSE EPSRC-funded project.

The goals of the work in this area are: 1) adapt the wireline experiences discussed above to the wireless sensor context; 2) study the potential to exploit the broadcast nature of radio to optimize any management-specific traffic; and 3) investigate the extent to which the federation of autonomic management can be applied to sensor networks.

4.4 Application Data Management

The primary role of a networked sensor is to measure an aspect of the environment and to transmit the corresponding (possibly aggregated) data upstream for analysis by more powerful components of the sensor network. Therefore, the principal application aspect to be considered in multi-dimensional co-design is the management of these data streams.

There is a general trend in database research to capitalize on ever more functional distributed processing platforms for deployment; of particular interest in this context are query processors on data streams [15]. The surge in interest in sensor network data management [16], particularly in query processors (SNQPs) can be seen as another manifestation of this general trend. These software artifacts aim to provide a holistic solution; their overall approach can be characterized as one of reconstructing the classical query-processing stack but with a sensor network as opposed to a general-purpose server as the target execution environment. Current SNQPs are limited to the kinds of sensor network applications that query languages can completely express. If an application's needs are not entirely expressible through queries, then the SNQP approach is unnecessarily constraining, insofar as it leads to isolated decisions being made about how to use radio, processor and network resources, and these may well compromise the global cost of the deployed system. Note that applications in which there is a reverse data flow (i.e., sensors are also sinks to data flows, as needed in network management and monitoring) are not completely expressible in current SNQPs. In this sense, SNQPs can be said to be exclusively concerned with synergy in data retrieval and data aggregation tasks for a given set of constraints to the detriment of orthogonality over a range of data management operations and different kinds of constraints. The holistic SNQP approach is less than ideal in the context of this proposal. We need to make available to the generative programming environment a range of options as to data management capabilities, so that it can make decisions as to how best to deploy them and minimize the global cost.

4.5 Formalization

The specification of the individual areas of concern will need to be tailored initially to the specific concerns of each area. There are a number of modeling languages available, e.g. CSP [8], Petri-nets [9] and queuing models [10]. Higher level languages, e.g. Promela [11], are specifically designed to allow verification via bespoke model checkers. In this project, the type of models created will depend on the characteristics of the individual dimensions. For example, as CSP models can be easily obtained from VHDL hardware designs, CSP would be an obvious choice for specifying the hardware description. However, in order to evaluate the performance (e.g., in terms of message flow) through the sensor network, a language that supports quantitative analysis (e.g., queuing models, stochastic Petri-nets or the probabilistic modeling language Prism [12]) would be more suitable. Some work has already been done on modeling different aspects of sensor networks (e.g., lifetime analysis using the hybrid model checker HyTech [13], and the modeling of performance variation for sensor net diffusion protocols [14]).

The orthogonal dimensions will be primarily designed by appropriate experts in the project team. However, it is crucial to be able to share these design choices across the project team. In particular, it is essential that one partner own a formally specified blueprint of both the individual dimension functionalities, and that of the system as a whole. These formal specifications will not only lead to a better understanding of the individual components and the entire system, but will result in clearer documentation. In addition, as well as traditional testing, we will formally verify the models (e.g., via model checking [7]). This will help us to ensure correctness of the designed system with respect to a global cost function via a design-test-verify-redesign cycle.

4.6 Generative Programming and Adaptation

To support the formal specification of system components, we will take a software architecture approach that describes systems in terms of components and their interactions. To guide self-adaptation within a unified framework, architectures need to specify the behavior of their components as well as constraints on the structure and cardinality of their components and interactions. We bring a software architecture-based approach for self-adaptation where policies may be encoded within languages using mechanisms for supporting constraints, feedback and change.

Various frameworks, languages and methodologies have been proposed for the construction of self-adaptive systems. Containment Units are used to build adaptive systems that deal with anticipated change in [21]. ArchStudio [22] is a tool suite developed to support self-adaptation in the C2 architecture style. Georgiadis et al [23] use specific component managers to identify external architecture changes by listening to events, and then react in order to preserve architecture constraints. Architecture styles, augmented with adaptation operators and repair strategies, are used as the basis for self-repair by Garlan et al [24]. IBM's autonomic computing initiative [25] aims for systems that are self-configuring, self-optimizing, self-protecting, and self-healing. The ArchWare Architecture Description Language (ADL) is a strongly-typed, executable architecture description language designed and implemented as part of the EU-funded ArchWare project [26]. It extends an expression language with typed higher-order polyadic π-calculus and constructs to support composition, decomposition, dynamic evolution and recomposition of systems. The ADL is used to specify the structure and execution of dynamic, self-adapting architectures.

Within this project constraints pertaining to an application will be used as the conditions that must hold true at all times during its execution. We have experience in the mechanisms for feedback and change within a computationally complete architecture description language based on typed higher-order polyadic π-calculus. Feedback is supported through software probes, gauges and an event distribution network, and change is supported through an adaptation engine using the concepts of decomposition, reification, reflection, recomposition and hyper-code.

Both the constraint language and the software ADL will be used to specify each of the components of the desired system. The constraints will be used to determine a deployment cost function against which the generator tool can be used to produce the necessary implementations. Once deployed, adaption may take place within components as long as it does not violate the constraints or, at a global level, where the constraints themselves have changed.

5 Work Packages

The work packages address: generic requirements for sensor networks, the formalization of specification languages for specific requirements in each area of concern, the generative mechanisms for weaving these specifications into an optimal co-designed executable that meets the specifically-stated functional and quality requirements, the formal verification mechanisms to ensure that the desired properties hold for the resulting artefact, an implementation platform for prototyping sensor systems so designed, and the construction of two prototype systems using these techniques. The lead individual for a particular work package is indicated in parentheses; other collaborators participate in each WP. Note that all of the efforts in WP2 and WP3 contribute to WP4.

5.1 WP1 – Architecture (Kent RA)

This activity will: use requirements from existing wireless sensor networks to determine a generic set of requirements for such systems; define the distributed architecture of sensor systems and the node architecture of individual sensors from these generic requirements; determine basic modeling concepts for all components that make up the system and, in particular, those aspects of the system that are adaptable versus those that are assumed fixed.

Deliverables: use case scenarios, generic requirements, system and node architecture documents, and basic modeling concepts,

5.2 WP2 – Formalizations and Metrics (Glasgow RA)

This activity will: develop a blueprint of formalizations of individual dimension functionalities and the system as a whole; design and document a suite of appropriate cost functions (individual areas of concern will be formalized primarily by individual experts in each respective field - see sub-workpackages below); and ensure that reasoning and optimization can address the system as a whole, and transitively, each individual area of concern.

Deliverables: formal descriptions of each dimension, integration of these descriptions to yield system description, and specification of per-dimension cost functions.

Each sub-workpackage below will produce: a) a formalization of the constraints or costs associated with that particular dimension, expressed in a language expected by the generator tool (WP3); b) design appropriate components that can be specified using the ADL expected by the generator (WP3); and 3) implement these components for use in the integration platform (WP4). These outputs are described in each item below as constrain/specify/implement.

5.2.1 Application Data Management (Manchester RA)

This activity will: identify and characterize distinct (radio-, processor-, data- and network-related) decision-making processes embedded into existing SNQP query stacks; add reverse flow characteristics to query language capabilities to allow life-cycle management, reconfiguration and adaptation actions; redesign and enhance the functionality provided by current SNQP stacks into a collection of software components for application data management; participate in global decisions regarding requirements and architecture; design and implement data management components; formalize components and cost behaviours using constraint and software architecture description languages (see sections 4.6 and 4.5); implement components and make them available to generator tool.

Deliverables: formal descriptions of the application data management dimension as constraints and architectural descriptions of the data management software components in the languages expected by WP3, a suite of data management software components whose costs are described by the constraints above and whose behavior conforms to the architectural descriptions mentioned.

5.2.2 Network Management and Monitoring (Glasgow PhD student)

This activity will: leverage wireline co-design theory and structure, with appropriate modifications, into the wireless sensor network environment, relating to the work in 5.2.1 and 5.2.3; constrain/specify/implement; and leverage self-managed cell federation techniques developed in the AMUSE project.

Deliverables: formal descriptions of the network management and monitoring dimension as constraints and architectural descriptions of the components expected by WP3, a suite of system management software components whose costs are described by the constraints above and whose behavior conforms to the architectural descriptions; mechanisms in place to federate control of a deployed sensor system with a juxtaposed management console.

5.2.3 Sensor Operating System and Networking (Kent RA)

This activity will: focus on adaptive control of measurement process and metadata associated with networked sensors; explore adaptive concepts in sensor OS and networking stacks; and constrain/specify/implement.

Deliverables: formal descriptions of the OS and networking dimension as constraints and architectural descriptions of the componenents expected by WP3, a suite of OS and networking software components whose costs are described by the constraints above and whose behaviour conforms to the architectural descriptions.

5.2.4 Radio Communications (Strathclyde RA and PhD student)

This will focus on security, data integrity and energy conservation of the radio system and will develop and assess the performance of self configuration processes with multiple heterogeneous radio interfaces. This will involve a) cross layer optimisation techniques, including the physical layer, for device management b) development of a transmission integrity strategy for dynamically changing networks with multiple radio interfaces and c) integration of data security and energy conservation constraints. [PhD to concentrate on a) linking with WP4]

Deliverables: Formal descriptions of the communications stack and networking dimension as constraints and architectural descriptions of the components expected by WP3, a suite of OS and networking software components whose costs are described by the constraints above and whose behaviour conforms to the architectural descriptions.

5.3 WP3 – Generator and Adaptive Constraint Solution (St Andrews RA & PhD student)

This activity will: use the formal descriptions of each of the systems components, both software and hardware, to generate efficient, integrated solutions against a particular cost function; determine an evolving system model against which to evaluate observations via: a) development of specifications of components (layers); b) development of heuristics for balancing conflicting and evolving requirements of the layers; and c) deployment and change of the system by observing feedback from probes to modify initial model assumptions; and formal validation of the resulting design during development and (potentially) of adaptations at run-time will be performed by Glasgow. (Glasgow RA)

Deliverables: the generator system, specification of software probes and gauges, the adaptation system, the formal verification techniques and technologies.

5.4 WP4 – Prototypes (Manchester PhD, Kent RA and PhD student, St Andrews RA)

This activity will: integrate components produced by other work packages into the implementation platform and design and construct two different prototypes using the system: i) environmental sensors applied to hydrology of a watershed and ii) an active badge system (in conjunction with an externally-funded PhD student at St Andrews). Kent will produce the sensor systems to be deployed, and Strathclyde will produce the communications circuitry.

Deliverables: integrated implementation platform, design documents for prototypes, implemented and deployed prototypes.

6 Resource Requirements and Justification

University	RA	PhD	Sec'y	SysProg	Tech	Desktop	Laptop	Area of PhD focus
Glasgow	1	1	10%	10%	-	2	1	The management activities in WP2, § 5.2.2.
Kent	1	1	-	10%	30%	2	1	Design and implementation of reconfigurable sensor nodes based upon this methodology..
Manchester	1	1	-	-	10%	2	1	Devise model of catchment processes, specify sensor network, interpret results of deployed system.
St Andrews	1	1	-	10%	-	2	1	The adaptation system in WP3, § 5.3.
Strathclyde	1	1	-	-	20%	2	1	Incorporation of cross layer optimisation procedures into the testbed, contributing to the radio aspects of WP4.

The RAs are responsible for leading the effort at their respective universities, as documented in section 4 above. All RAs are requested at point 8 on the RA1A scale, as the research challenges require a reasonable degree of maturity: 1) each area involves re-factoring and co-design; 2) the ability to aggressively track cutting edge literature in their specific research areas is crucial; and 3) identification and formalization of their specific dimension will require complete understanding of the dimension. The PhD students are justified in the table above. The secretarial support at the lead organization is requested to facilitate the interactions required of such a large collaboration. System programmer support has been requested by each university in which a significant amount of software development will take place. Technical support is requested as follows: at Kent, for design, layout, fabrication and initial testing of the prototype sensor systems proposed by the other partners in the earlier workpages; at Strathclyde, for construction of communications hardware for the sensor systems; and at Manchester, for analysis of water and sediment samples taken weekly from the various sites within the catchment, as well as to accompany the student on field visits, when necessary, particularly during network installation. Additionally, Kent is requesting funds for the construction of the environmental sensor systems, estimated at £1000/system; approximately 20 systems will be needed. A desktop system is requested for each RA and PhD student, and a laptop per university.

Each university is requesting sufficient travel funds for all investigators, RAs, and PhD students to attend the semi-annual project meetings, 6x£250 to cover the PI travel to management meetings between the project meetings, and sufficient funds for each investigator, RA, and PhD student to travel to one international conference during the life of the project. The total budget is approximately £1.25M over 3 years.

7 Relevance to Beneficiaries

Academics will benefit from the extended notion of co-design that the project focuses on; besides being able to take more than 2 dimensions into account when performing co-design, the formal validation of the resulting system should encourage others to further the work that we will have started.

Many different parts of the commercial, sensor system, value chain will benefit from this research: 1) chip vendors, especially those focusing on field programmable logic, will be able to make their chips more attractive for use in sensor systems by providing customizability that can be exploited by the design and development tools; 2) wireless service providers can benefit from the multi-air-interface support facilitated by the methodology and tools; 3) SME's that design, build, and deploy sensor-based systems will be able to sell systems that are optimal with regards to the cost function that is important to them and their customers; and 4) customers that need sensor systems can be more assured that the systems delivered are optimized to meet their requirements.

8 Collaboration

The project consists of a collaboration of computer science, electrical engineering, and environmental science departments at five UK universities. Each of these departments is contributing both unique and essential expertise for the success of the DIAS-MC proposed research.

We have also industrial partnership from several companies at different points in the sensor system value chain:

- ***Intelisys, Ltd*** is an SME that will provide consultation regarding the requirements of hydrochemical monitoring. They have been operating an existing sensor system which did not have the benefit of DIAS design methodology: by instrumenting the same catchment area, we will be able to compare a number of characteristics of the system generated using our methodology against the existing system.
- ***Orange PCSL*** is a mobile communications service provider; they will provide us with GPRS/3G data communications capability to support our prototyping efforts in the field.
- ***Severn Trent Water Plc*** manage upland catchments for drinking water collection and will provide further calibration data.
- ***Xilinx Development Corp.*** is the worldwide market share leader for field-programmable logic devices; they will provide us with development systems and FPGAs for incorporation into our prototype sensors.

In addition to the contributions in kind, each of our partners have agreed to the participation of a senior technical individual in our discussions, minimally at our 6-monthly progress meetings, but potentially through regular involvement via our mailing lists and phone conferences. This arrangement bodes well for requirements transfer into the project, technology transfer from the project, and continued interaction to maximize the commercial relevance of the research.

9 Exploitation and Dissemination

Industrial, commercial and academic contacts will be used to disseminate the results of the proposed research. The participating research groups are all well-established and well-connected with leading academic and commercial researchers at both national and international level. The ideas associated with the research, lessons from its implementation, and the results of investigations, will be published in technical reports, workshops, conferences, books, journals and on the Internet. Free access to any systems will be given to academic institutions. A publicly accessible web site will be maintained. We intend to make any software artefacts available via open source distribution.

Through Kent, we will have strong interaction with the DTI EnviSense program. Potential commercial exploitation may be achieved through our partners, Intelisys, France Telecom, and Xilinx; I. Marshall's link with the ITI TechMedia sensors project may also provide potential exploitation opportunities. Any commercial exploitation will be undertaken in agreement with appropriate enterprise organizations at the universities involved.

10 Project Management

A project of this size and scope requires careful attention to management; to that end, Glasgow will assume a leadership role in the management of the project. Besides scheduling and leading the required meetings (discussed below), Glasgow will be responsible for preparing and submitting any required progress updates to the EPSRC, as well as representing the project at EPSRC workshops/meetings.

The PIs will meet on a quarterly basis to discuss progress to date, issues that need to be proactively addressed, and any change of project plans required. The entire project team will meet every 6 months to discuss recent results, to brainstorm based upon those results, and to plan activities over the next 12-month period. All such meetings will be minuted, and the minutes will be available to project members on the web site.

For the day-to-day management of the project, the RA's, PhD students, and involved investigators will conduct regular phone conferences (anticipated to be approximately every 2 weeks) to discuss detailed issues affecting the project. Additionally, a mailing list will be maintained and archived for non-real-time discussion of issues at this level.

References

[1] J Kahn, R Katz and K Pister, "Mobile Networking for Smart Dust", Proc of ACM/IEEE Intl. Conf. on Mobile Computing and Networking, August 1999.
[2] Y Li, T Callahan, E Darnell, R Harr, U Kurkure and J Stockwood, "Hardware-Software Co-Design of Embedded Reconfigurable Architectures", Proc. of the Design Automation Conference, Los Angeles, California, 2000.
[3] CodeSign, http://www.tik.ee.ethz.ch/~codesign/.
[4] G Kiczales, J Lamping, A Mendhekar, C Maeda, C Videira Lopes, J-M Loingtier and J Irwin, "Aspect-Oriented Programming", Proc. of the European Conf. on Object-Oriented Programming, June 1997.
[5] D Orleans, "Incremental Programming with Extensible Decisions", Proc. of the 1st International Conference on Aspect-Oriented Software Development, April 2002.
[6] Y Coady, "Crosscutting the Great Divide: Exploring an Aspect-Oriented Approach to OS Code", http://www.cs.ubc.ca/~ycoady/papers/prop.pdf.
[7] E M Clarke, O Grumberg, and D Peled, Model Checking, The MIT Press, Cambridge, MA, 1999.
[8] C A R Hoare. Communicating Sequential Processes. Prentice-Hall, 1985.
[9] J L Peterson. Petri net theory and the modeling of systems. Prentice-Hall, 1981.
[10] S K Bose. An introduction to queueing systems. Kluwer, London, 2002.
[11] G J Holzmann. The SPIN model checker: primer and reference manual. Addison Wesley, Boston, 2003.
[12] M Kwiatkowska, G Norman, and D Parker, "Probabilistic symbolic model checking with PRISM", Proc. of the 8th International Conference on Tools and Algorithms for Construction and Analysis of Systems, Grenoble, France, April 2002.
[13] S Coleri, M Ergen, and T John Koo, "Lifetime analysis of a sensor network with hybrid automata modelling", Proc. of the 1st ACM International Workshop on Wireless Sensor Networks and Applications, Atlanta, Georgia, USA, September 2002.
[14] S Nair and R Cardell-Oliver, "Formal specification and analysis of performance variation for sensor network diffusion protocols", Proc. of the 7th ACM/IEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile systems, Venice, Italy, 2004.
[15] A Arasu, B Babcock, S Babu et al. STREAM: The Stanford Stream Data Manager. IEEE Data Eng. Bull. 26(1): 19-26 (2003)
[16] J Gehrke, S Madden. Query Processing in Sensor Networks. Pervasive Computing, Jan-Mar 2004, pp. 46-55.
[17] L Sacks, M Britton, I Wokoma, A Marbini, T Adebutu, I Marshall, C Roadknight, J Tateson, D Robinson, and A G Velazquez, "The Development of a Robust, Autonomous Sensor Network Platform for Environmental Monitoring", IoP Sensors and their Applications (S&A XII), University of Limerick, Ireland, 2003.
[18] H-Y Hsieh, H-H Kim, and R Sivakumar, "An end-to-end approach for transparent mobility across heterogeneous wireless networks", Mobile Networks and Applications, Vol 9, No 4, pp 363-378, August 2004.
[19] G Pollock et al, "The Asymptotic Configuration of Application Components in a Distributed System", http://www.dcs.gla.ac.uk/~joe/auxiliary/papers/Personal/AsymptoticConfig.pdf.
[20] C S Raghavendra et al., Wireless Sensor Networks, Kluwer Academic Publishers, ISBN 1-4020-7883-8, May 2004.
[21] P Oreizy, N Medvidovic, and R Taylor, "Architecture-Based Runtime Software Evolution", Proc. of the International Conference on Software Engineering, Kyoto, Japan, 1998.
[22] E Dashofy, A van der Hoek, and R Taylor, "Towards architecture-based self-healing systems", Proceedings of the 1st ACM SIGSOFT Workshop on Self-Healing Systems, Charleston, SC, USA, 2002.
[23] I Georgiadis, J Magee, and J Kramer, "Self-Organising Software Architectures for Distributed Systems", Proc.1st ACM SIGSOFT Workshop on Self-Healing Systems, Charleston, SC, USA, 2002.
[24] D Garlan, S-W Cheng, and B Schmerl, "Increasing System Dependability through Architecture-based Self-repair", *Architecting Dependable Systems*, LNCS 2677, Springer-Verlag, 2003.
[25] IBM Autonomic Computing, http://www.research.ibm.com/autonomic/.
[26] F Oquendo, BC Warboys, R Morrison, R Dindeleux, F Gallo, and C Occhipinti, "ArchWare: Architecting Evolvable Software", Proc. 1st European Workshop on Software Architecture, St Andrews, UK, 2004.

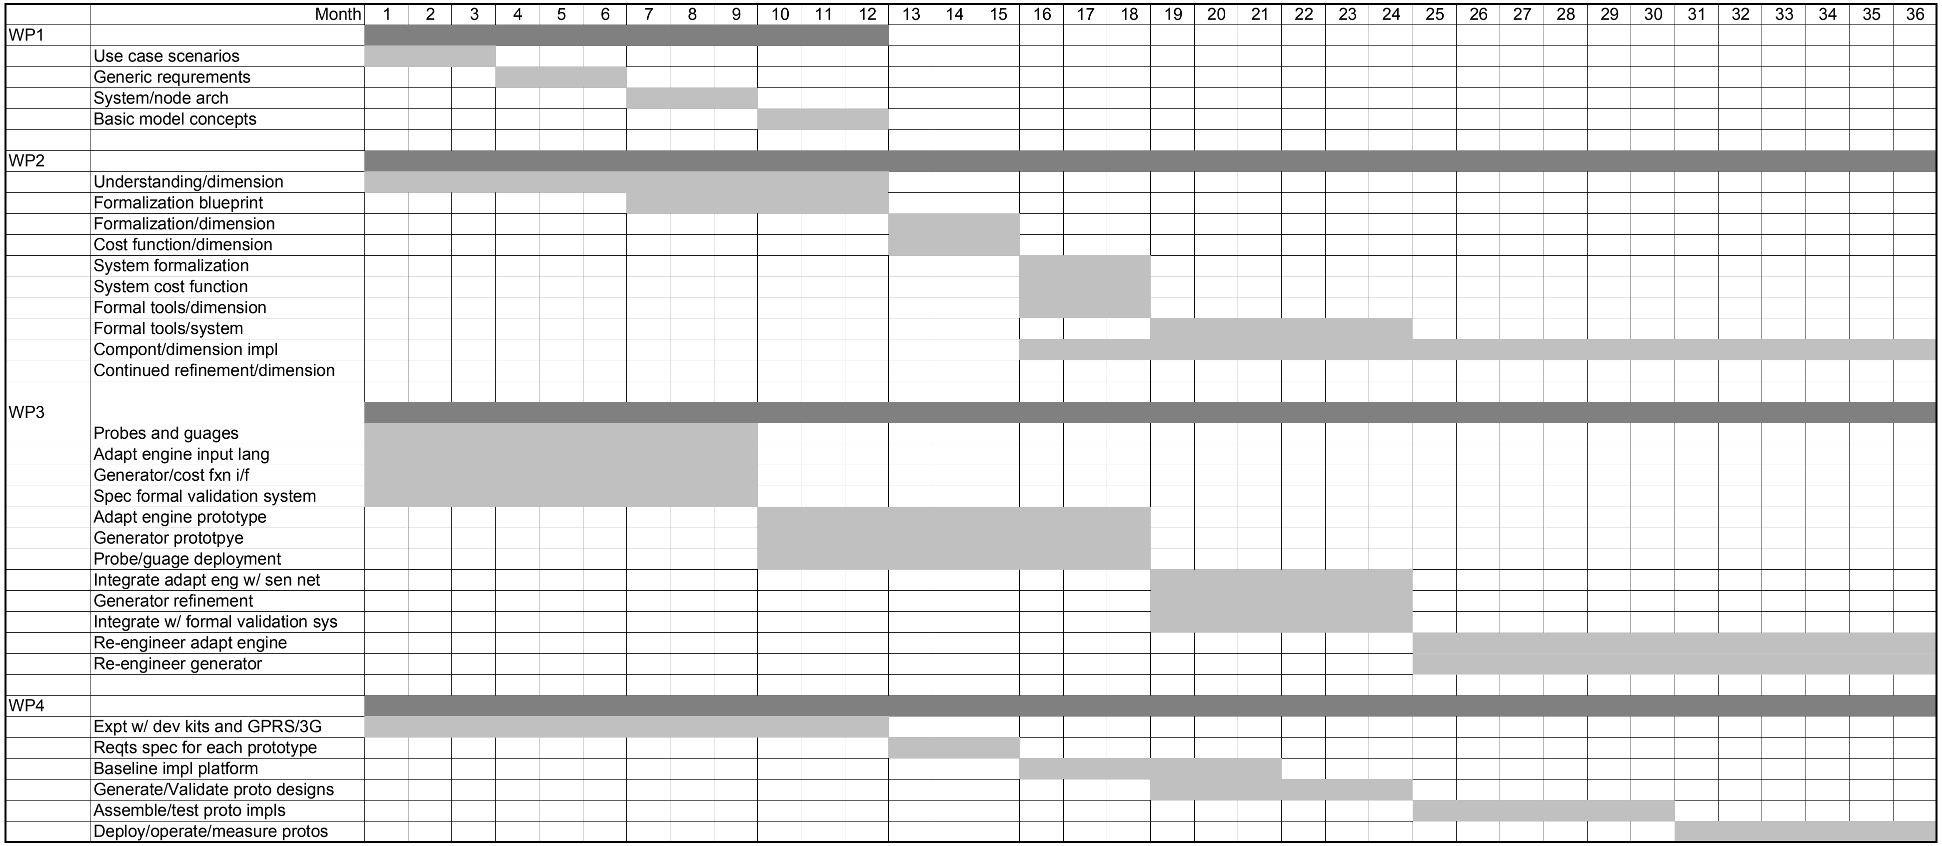
Month
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36
WP1
Use case scenarios
Generic requrements
System/node arch
Basic model concepts
WP2
Understanding/dimension
Formalization blueprint
Formalization/dimension
Cost function/dimension
System formalization
System cost function
Formal tools/dimension
Formal tools/system
Compont/dimension impl
Continued refinement/dimension
WP3
Probes and guages
Adapt engine input lang
Generator/cost fxn i/f
Spec formal validation system
Adapt engine prototype
Generator prototpye
Probe/guage deployment
Integrate adapt eng w/ sen net
Generator refinement
Integrate w/ formal validation sys
Re-engineer adapt engine
Re-engineer generator
WP4
Expt w/ dev kits and GPRS/3G
Reqts spec for each prototype
Baseline impl platform
Generate/Validate proto designs
Assemble/test proto impls
Deploy/operate/measure protos