Where the Grid meets the Physical World – Research Issues in Grid and Pervasive Computing

David De Roure, Tony Hey and Anne E. Trefethen

Abstract— **The digital world of Grid technologies meets the physical world through a variety of sensors, instruments and interfaces. The e-Science programme has increasingly become aware of the need to focus on this digital-physical interface of the grid, and the ubiquitous computing community is beginning to look towards the Grid for aspects of processing, data handling and access (for example, in sensor networks). There are also opportunities to apply infrastructure and middleware techniques across these distributed computing domains. In this paper we survey the activities which bring together Grid and pervasive computing, drawing examples from the e-Science programme which engage with both areas of research and highlight how the two have a symbiotic relationship. From this we identify areas for further research.**

***Index Terms*—Grid, Pervasive Computing, Ubiquitous Computing, Semantic Grid**

I. INTRODUCTION

THE digital world of the Grid meets the physical world through a variety of sensors, instruments and interfaces. These two significant trends in computing technology – more devices around us, more integration and power behind the scenes – need to be considered together and in this paper we suggest that they have a symbiotic relationship. Grid applications, as exemplified by many projects in the UK e-Science programme, have increasingly become aware of the need to focus on this digital-physical interface of the grid, and the pervasive and ubiquitous computing community is looking towards the Grid for aspects of processing, data handling, integration and access. Meanwhile there is also interest in applying middleware techniques across these distributed computing domains.

Both trends stand to benefit from a third significant movement in computing – the move towards machine-processable explicit knowledge as exemplified by the Semantic Web. This enables the automation and interoperability, which is increasingly necessary in these open, distributed systems and is essential to realise their full ambitions. This is illustrated through the adoption of Semantic Web technologies within the practice of Grid computing, a field of endeavour known as Semantic Grid [1]. These technologies are also being adopted within pervasive computing, for example in representing context information and device ontologies. Significantly, they also facilitate a capability for automated inference – an aspect of 'intelligent' behaviour.

In this paper we will describe the problem space defined by these important technology trends and identify some of the research challenges that need to be met, focusing in turn on grid applications, pervasive computing and the Semantic Web. In doing this we explain how e-Science has been motivated by the data deluge such as that produced by sensor networks, and how e-Science applications have also demonstrated the role of pervasive computing as a means of interaction. We then present case studies and close with the research issues that have arisen through this survey.

II. THE E-SCIENCE VISION

In 2001 the UK initiated a £250M, 5-year e-Science program to develop the tools, technologies and infrastructure to support multidisciplinary and collaborative science, allowing scientists to do 'faster, better or different' research [2]. Such an e-Science infrastructure is in fact very close to the vision that J.C.R.Licklider took with him to DARPA when he initiated the research project that led to the ARPANET [3] and the present day Internet – but the killer applications have so far been email and the Web rather than the distributed computing vision. In the early 1960s, Licklider only envisaged connecting a small number of scarce, expensive computers at relatively few sites. However, the relentless improvements in silicon technology over the past forty years summarized in Moore's Law – Gordon Moore's prediction that the number of transistors on a chip would double about every 18 months and that the price-performance would be halved – has led to an explosion in the

---

Manuscript received June 3, 2005. This work was supported in part by the EPSRC under grants GR/R67729, GR/R85877 and GR/R85143.

D. De Roure is with the University of Southampton, UK (phone: +44 23 8059 2418; fax: +44 23 8059 2865; e-mail: dder@ecs.soton.ac.uk).

A. J. G. Hey is with UK Engineering and Physical Sciences Research Council. (e-mail: Tony.Hey@epsrc.ac.uk).

A. E. Trefethan is also with the UK Engineering and Physical Sciences Research Council (e-mail: Anne.Trefethen@epsrc.ac.uk).

number of supercomputers, mainframes, workstations, personal computers and other devices that are connected to the Internet. It has also led to the availability of the smaller and cheaper devices which drive forward pervasive computing, through which we see increasing numbers of handheld and embedded devices in our everyday environment.

One of the key drivers underpinning the e-Science movement is the imminent deluge of data from the new generations of scientific experiments and surveys. In order to exploit and explore the Petabytes of scientific data that arise from these high throughput experiments, supercomputer simulations, sensor networks and satellite surveys, scientists need assistance from specialised search engines and data mining tools. To create such tools, the data needs to be annotated with relevant metadata giving information as to provenance, content, conditions and so on and, in many instances, the sheer volume of data will dictate that this process is automated. In the next few years scientists will create vast distributed digital repositories of scientific data that will require management services similar to those of more conventional digital libraries as well as other data-specific services. The ability to search, access, move, manipulate and mine such data will be a central requirement for this new generation of collaborative science applications.

For example, in the particle physics and astronomy communities the next decade we will see new experimental facilities coming online that will generate data sets ranging in size from 100s of Terabytes to 10s of Petabytes per year. In the field of engineering, consider the problem of health monitoring of industrial equipment. The UK e-Science programme has funded the DAME project [4] – a consortium analysing sensor data generated by Rolls Royce aero-engines. It is estimated that there are around 100,000 Rolls Royce engines currently in service. Each trans-Atlantic flight made by each engine, for example, generates about a Gigabyte of data per engine – from pressure, temperature and vibration sensors. The goal of the project is to transmit a small subset of this primary data for analysis and comparison with engine data stored in three data centres around the world. By identifying the early onset of problems, Rolls Royce will be able to lengthen the period between scheduled maintenance periods and increasing profitability. The engine sensors will generate many Petabytes of data per year and decisions need to be taken in real-time as to how much data to analyse, how much to transmit for further analysis and how much to archive. Similar (or larger) data volumes will be generated by other high-throughput sensor experiments in fields such as environmental and earth observation, and of course human health-care monitoring.

With this imminent deluge of scientific data, the issue of how scientists can manage these vast datasets becomes of paramount importance. Up to now, scientists have generally been able to manually manage the process of examining the experimental data to identify potentially interesting features and discover significant relationships between them. In the future, when we consider the massive amounts of data being created by simulations, experiments and sensors, it is clear that in many fields they will no longer have this luxury. The discovery process – from data to information to knowledge – needs to be automated as far as possible. At the lowest level, this requires automation of data management with the storage and organisation of digital entities. At the next level, we require automatic annotation of scientific data with metadata describing both interesting features of the data and of the storage and organization of the resulting information. Finally, we need tools to enable scientists to progress beyond the generation of mere structured information towards the automated knowledge management of our scientific data.

3. PERVASIVE COMPUTING

Enabled by technology advances and driven by customer demand for portable devices which are smaller, lighter and which run for longer on batteries, pervasive computing devices are now becoming increasingly prevalent in our everyday life – for example, the mobile phones, PDAs (portable digital assistants), digital cameras, global positioning systems and other electronic accessories that we carry on our person and place in our everyday environment. These portable devices communicate using a range of wireless technologies including Bluetooth, ZigBee, wireless Ethernet and the Global System for Mobile Communications (GSM). Internet protocols are evolving, for example through IPv6, to accommodate increasing numbers of devices, mobility and automatic configuration. Noting Moore's Law, the clear technology trend is towards massive deployment of low-cost and low-power devices, and increasing power in handheld devices.

These pervasive devices form the intersection between the physical world and the digital world. Like the e-Science examples above, they provide a deluge of data that can demand sophisticated processing – effectively we are immersing ourselves in a vast, highly heterogeneous sensor network. Significantly, they also provide the means of interaction with the digital world of the Grid. We discuss these aspects below, and introduce the need for 'intelligent' capabilites.

A. Sensor Networks

Pervasive computing systems have the potential to deliver environmental and social benefits, through improved monitoring, providing access to information to create decision making capabilities and alert mechanisms. There are many application areas, particularly involving wireless sensor networks, including environmental monitoring in such applications as pollution control, disaster recovery, in medical sensors and patient care, and in business applications such as building the smart supply chain [5]. For example, the NASA/Jet Propulsion Laboratory conceived the 'Sensor Web' to take advantage of inexpensive mass consumer-market chips to create platforms that act together as a single instrument and which may be embedded into an environment to monitor and control it [6]. As the technologies continue to improve we will see smaller, lower power and lower cost devices. The latest research has created novel methods for harvesting energy from the environment to provide self-powered microsystems, giving a glimpse of the degree of pervasiveness the future may hold [7].

Prior to sensor networks, measurements would typically be taken manually, infrequently and at a small number of locations. Using pervasive devices, we can achieve a considerably higher spatial and temporal density of measurements. Not only does this increase data volume but the computational task implied by this may be quite considerably higher than before – for example, where the data feeds into models and simulations performing numerical calculation, or where there is a demand for real-time data processing. With increasing numbers of deployments, the data integration task also becomes substantial.

Some applications effectively make use of everyday devices as a sensor network. For example, the more information that can be obtained about the current context of a person, the better able the system is to adapt to their requirements and provide appropriate services. This includes determining location and movement. Other applications focus on using multiple information sources in order to track the movements of individuals, raising a set of privacy issues and a significant technical challenge in drawing reliable inferences from multiple information sources of variable quality.

B. *Interaction*

Pervasive devices enable users to interact with the Grid. Within e-Science projects, we see examples of the use of specialised devices to collect data as well as general purpose devices – such as mobile phones – being used to monitor and control experiments. Context is important for information delivery – providing the right information on the right device at the right time and with the right level of intrusiveness. In fact this rich notion of context becomes the query: using context knowledge, devices can act proactively and autonomously in order to provide the user with the information they require [8]. Context is also crucial to information acquisition where it provides important metadata to aid the subsequent interpretation of the data and to record its provenance.

From the perspective of pervasive applications, the availability of Grid services makes it possible to achieve computational and data integration tasks that would not be possible on the devices themselves. For example, a digital camera can produce a volume of data that would benefit from remote grid processing for tasks such as data compression or feature extraction – even more so with a digital video camera. Many interactive pervasive applications demand real-time processing and therefore require significant computational power, to be delivered on demand – one example of this is mobile mixed reality. Hence we envisage an important role for the Grid in support of new applications, especially as the devices generate (and store) larger quantities of data.

C. *Intelligence*

The data processing tasks in support of pervasive applications require sophisticated data mining, data integration and the use of services such as feature extraction, pattern matching, language translation and gesture classification. Sometimes they involve working with multiple sources of information, such as in situation assessment. It is these aspects, coupled with the context-awareness described above, that take us from a device-centric world of pervasive computing into a world of "ambient intelligence" where the emphasis is on the users, intelligent user interfaces and distributed intelligence [9]. Many of these tasks stand to benefit from Grid processing and can be delivered as Grid services into the pervasive application.

The applications also require sophistication in the behaviour of the infrastructure – dynamic discovery and composition of services and devices in order to meet context-specific requirements. This contains many challenges since it is necessary for it to occur dynamically and there may be many applications competing for the same resources, which may not always be available. This is, for example, the territory of intelligent agents, which conduct negotiations in order to form coalitions to meet current requirements. As grid services become increasingly available, closely related issues will arise there – in both cases we need flexible, secure virtual organisations

III. THE SEMANTIC WEB

The applications discussed in the previous sections demonstrate that there is a very considerable degree of automatic processing, interoperation and integration that is demanded by both Grid computing and pervasive computing. The key to achieving this is to provide the necessary information in a standard, machine-processable form. For this we turn to another significant trend in contemporary computing: the Semantic Web [10].

The Semantic Web is described in the World Wide Web Consortium (W3C) Semantic Web Activity Statement as an initiative "…to create a universal medium for the exchange of data. It is envisaged to smoothly interconnect personal information management, enterprise application integration, and the global sharing of commercial, scientific and cultural data. Facilities to put machine-understandable data on the Web are quickly becoming a high priority for many organizations, individuals and communities. The Web can reach its full potential only if it becomes a place where data can be shared and processed by automated tools as well as by people. For the Web to scale, tomorrow's programs must be able to share and process data even when these programs have been designed totally independently.”

The Semantic Web technologies are based on the W3C’s *Resource Description Framework* (RDF) which provides a standard way to represent metadata – which could for example be data about documents, objects, devices, grid resources or people. The shared vocabularies that are used – and which are the key to interoperability – are called ontologies and can be represented in the W3C’s *Web Ontology Language* (OWL). Currently rule languages are being developed to operate alongside OWL. Tools for RDF are readily available, including for example the ‘RDF triplestores’ that provide a means of working with large volumes of RDF metadata and conducting queries upon it.

RDF can be used to describe the various entities in our pervasive and Grid systems, for example resources, services, devices, context, and user profiles. The creation of ontologies for these entities are topics of current research efforts, as are the techniques of ‘semantic matching’ which help discover appropriate entities; for example, standards for Semantic Markup for Web Services are now under consideration [11]. As well as describing entities, RDF also enables us to model the relationships between things in the physical world. As metadata is recorded in different times and places about a specific object with a unique identifier, it is effectively linked together by that common identifier. Hence the metadata accumulates, forming a rich, machine-processable body of knowledge.

A. Semantic Web and the Grid

In 2001, De Roure, Jennings and Shadbolt introduced the notion of the Semantic Grid which advocated ‘the application of Semantic Web technologies both on and in the Grid’ [12-13]. From the requirements derived from the diverse set of UK e-Science applications they identified a need for maximum reuse of software, services, information and knowledge. Although the basic Grid middleware was originally conceived for hiding the heterogeneity of distributed computing, the authors contended that users now required ‘interoperability across time as well as space’ to cope with both anticipated and unanticipated reuse of services, information and knowledge.

Best practice in Semantic Grid is emerging through a series of e-Science projects which have applied Semantic Web technologies to Grid applications in a variety of ways. For example, the ^{my}Grid e-Science project builds on semantic web technologies and is researching high-level middleware to support personalised in silico experiments in biology [14]. As well as applying Semantic Web technologies to the bioinformatics application [15], the project has also researched service description and discovery at the middleware level [16].

B. Semantic Web and Pervasive Computing

Semantic Web technologies are also being adopted within pervasive computing, for example in representing context information [17] and device ontologies [18], allowing device capabilities and service functionality to be explicitly represented. This allows devices to advertise their services and permits a vision where a device might extend its functionality by contracting the use of a service from another device. The ‘task computing’ activity at Fujitsu Laboratories of America and University of Maryland illustrates the combination of Semantic Web and pervasive computing [19], encouraging device manufacturers to incorporate Semantic Web technologies into their devices in order to give end-users easier and more flexible use of the features of the devices. The Task Computing environment includes Web Services as well as RDF, OWL and Universal Plug and Play (UPnP).

Semantic Web technologies are being used today, but they are typically applied in the context of information on the Web, which evolves slowly. In the context of pervasive computing, vast amounts of metadata may be generated quickly and in a distributed fashion. This poses engineering challenges for the RDF triplestores and in handling distributed (and possibly inconsistent) knowledge. It is also interesting to note that pervasive computing can assist with automatic metadata capture and therefore in building the Semantic Web – the value of the Semantic Web relies on good metadata being available. Hence Semantic Web technologies can help build the pervasive computing infrastructure and pervasive computing can help build the Semantic Web.

Fig. 1 suggests a way of presenting the problem space created by these three major trends – Grid, Pervasive and Semantic – and positions some of these activities within it.

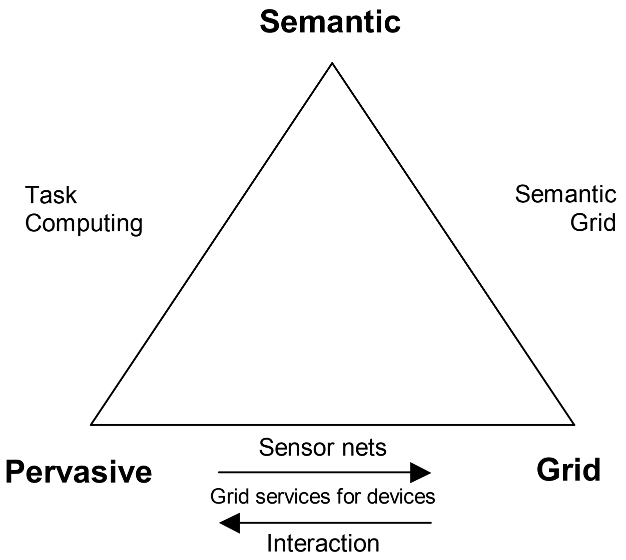

Fig. 1. Sensor networks and pervasive applications demand grid computation; Grid applications demand devices for interaction. Semantic Web technologies are being applied to both, to help realise the integrated vision.

IV. THE SYMBIOSIS OF GRID AND PERVASIVE COMPUTING

The digital world of the Grid meets the physical world through a variety of sensors, instruments and interfaces. These two significant trends in computing technology – more devices around us, more integration and power behind the scenes – have a symbiotic relationship. Grid applications, as exemplified by the projects described above, have increasingly become aware of the need to focus on this digital-physical interface of the grid, and the ubiquitous computing community is looking towards the Grid for aspects of processing, data handling, integration and access.

Here we suggest seven aspects of the intersection between the grid and the physical world where we believe there is value to be had in considering the Grid and pervasive aspects together. In the next section we will then consider two case studies of projects working at this intersection.

A. Devices need the Grid for computation

As sensors and sensor arrays evolve, our pervasive environments are acquiring data with considerably higher temporal and/or spatial resolution than before. This data deluge demands considerably greater computational power to process and analyse it, especially where there are also real-time processing requirements. Currently, many pervasive deployments are relatively small scale, due to small numbers of devices or small numbers of users, but they will demand more Grid processing as numbers inevitably scale up – and, in particular, as the data is combined with other sources which are also scaling up as this gives rise to a very significant increase in complexity. Grid services also provide computational augmentation of devices to support pervasive applications – for example gesture recognition, language translation and the computational tasks demanded by augmented reality.

B. Devices need the Grid for integration

Pervasive/ubiquitous computing research tends to focus on individual devices and deployments of these, with an interoperable infrastructure within a deployment – there is no established common distributed systems infrastructure standard to handle the interworking of multiple diverse sets of devices. Storz argues that, through Grid technologies to link sets of devices together, significant potential exists for building ubiquitous computing applications on a hitherto unprecedented scale [20]. This wider-scale integration is necessary to realise the full benefits of the pervasive computing (and ambient intelligence) vision.

C. The Grid needs devices to interface with the physical world

There are many e-Science applications that use devices for data acquisition, interaction or notification. There are some familiar ‘Grid-enabled’ devices in a laboratory – the pieces of scientific equipment, instruments and ‘grid appliances’ connected directly to the Grid. At one end of the spectrum we have instruments such as telescopes and X-Ray diffractometers, but our interest here is in the multitude of pervasive devices in the surroundings of the Grid user. Traditionally the human user interface to the Grid has been through the graphical user interfaces of applications and portals. In the context of pervasive computing, the interface becomes the devices in the user’s environment. These devices may be used manually or they may be part of the environment, and they may be collecting data or providing notification and information display/visualisation. This shift from the browser on the desktop PC to the mobile device is the subject of W3C’s recently launched Mobile Web Initiative.

D. Virtualisation

Grid computing and pervasive computing are each about large numbers of distributed processing elements. While some aspects are contrasting (synchronous vs asynchronous intercommunication patterns, for example), at an appropriate layer of abstraction (moving up the triangle in Fig. 1) they both involve similar computer science challenges in distributed systems. Specifically, these include service description, discovery and composition, issues of availability and mobility of resources, negotiation of quality of service, autonomic behaviour, and of course charging, security, authentication and trust. Both need ease of dynamic assembly of components, and both rely on interoperability to achieve their goals. The peer-to-peer paradigm is also relevant across the picture.

Both grid and pervasive middleware infrastructures stand to benefit from Semantic Web technologies. Again this is about semantic interoperability, but in the middleware: we need service description, discovery and composition, and indeed research areas such as Semantic Web Services are being applied both to Grid and to Pervasive computing. Hence the Semantic approach sits above the large-scale distributed systems of Pervasive and Grid computing.

E. The information systems perspective

Much of pervasive computing and grid computing is actually about information. As parts of these distributed systems come together, interoperability of the information is essential. Kindberg notes “Too often, we only investigate interoperability mechanisms within environments instead of looking at truly spontaneous interoperation between environments. We suggest exploring content-oriented programming—data-oriented programming using content that is standard across boundaries—as a promising way forward.” [21]

It is not just the content but also the metadata that is essential to achieving interoperability. The techniques of the Semantic Web are quite appropriate here, especially as they lend themselves to automated processing. In particular, some inference capability at this level is a step towards intelligent or ‘smart’ environments and ambient intelligence.

F. *Grid computation on networks of devices*

Though the area is less well developed at this time, some researchers are interested in running grid computations over networks of devices. One case for this is in intelligent sensor networks where it would be advantageous to shift some of the back-end computation out into the field in order to perform some processing on the fly. Another is the future vision of large numbers of devices being available – either through pervasive computing or, one day, through fabrication of large numbers of processing elements that are embedded in everyday materials, as explored in the Amorphous Computing project at MIT [22].

G. *Self-organisation*

Self-configuration, self-management, self-optimisation and self-healing are all desirable properties in both Grid and pervasive systems – sometimes described as 'autonomic' behaviour. There is evidence of this sort of behaviour in *ad hoc* networks, while Grids are typically more highly managed and controlled. The vision is of adding new devices or grid nodes in an arbitrary manner and having the system adapt to them, and similarly adapting to change and failure.

Semantic Web technologies offer a means for the machine-processable knowledge, which is necessary for this level of automation, and we can envisage a self-organising Semantic Grid [23]. Agent-based computing has also been proposed as an appropriate technology in both Grid and pervasive middleware. In addition to a service-oriented approach, agents bring an important notion of autonomous behaviour, and this could be crucial to the necessary degree of automation in these large-scale systems. Agents are also able to respond to their dynamic circumstances using techniques of service negotiation, a capability needed in both the Grid and pervasive contexts.

V. CASE STUDIES

A. *CombeChem*

The CombeChem e-Science project aims to enhance the correlation and prediction of chemical structures and properties, using technologies for automation, semantics and Grid computing [24-25]. A key driver for the project is the fact that large volumes of new chemical data are being created by new high throughput technologies, such as combinatorial chemistry in which large numbers of new chemical compounds are synthesized simultaneously. The need for assistance in organising, annotating and searching this data is becoming acute. The multidisciplinary CombeChem team has therefore developed a prototype test-bed that integrates chemical structure-property data resources with a Grid-based computing environment.

One of the key concepts of the CombeChem project is 'Publication@Source' by which there is a complete end-to-end connection between the results obtained at the laboratory bench and the final published analyses [26] – an illustration of the information perspective. This starts with pervasive computing in the smart laboratory and Grid-enabled instrumentation. By studying chemists within the laboratory, handheld technology has been introduced to facilitate the information capture at this earliest stage (the SmartTea project [27-28]). Using tablet PCs, the system has been successfully trialled in a synthetic organic chemistry laboratory and linked to a flexible back-end storage system. A key finding was that users needed to feel in control and this necessitated a high degree of flexibility in the lab book user interface. The computer scientists on the team investigated the representation and storage of human-scale experiment metadata and introduced an ontology to describe the record of an experiment and a novel storage system for the data from the electronic lab book. In the same way that the interfaces needed to be flexible to cope with whatever chemists wished to record, the back end solutions also needed to be similarly flexible to store any metadata that might be created. Additionally, pervasive computing devices are used to capture laboratory conditions, and chemists are notified in real time about the progress of their experiment using pervasive devices

This data then feeds into the scientific data processing. All usage of the data through the chain of processing is effectively an annotation upon it, and the provenance is explicit. The creation of original data is accompanied by information about the experimental conditions in which it is created. There then follows a chain of processing such as aggregation of experimental data, selection of a particular data subset, statistical analysis, or modelling and simulation. The handling of this information may include explicit annotation of a diagram or editing of a digital image.

Crucially, this digital record is enriched and interlinked by a variety of annotations be they data from sensors, records of use, or explicit interaction. The annotations need to be machine processable, and useful for both their anticipated purpose and interoperable to facilitate subsequent unanticipated reuse. This is achieved by deployment of Semantic Web technologies; RDF is used through the system. At the time of writing there are 80 million RDF triples in the CombeChem triplestore. The target is 200 million, making this a substantial Semantic Web deployment – this is described as the 'Semantic Datagrid'.

Another e-Science project, CoAKTinG [29], has created tools to enhance the collaboration between e-Scientists, and has used CombeChem as a case study. This effectively extends the digital record by including the meetings between the e-Scientists, fully interlinked using semantic annotation. This is a further step towards the complete digital record. Originally conceived to provide real-time 'presence' information about other people in a distributed meeting, the tools have been extended to integrate with devices so that device status is also communicated. Hence the tools support the broader smart environment and can be seen as part of the pervasive computing infrastructure.

B. *Medical Devices*

Devices for health monitoring are an important pervasive computing application. Grid and Pervasive Computing come together in the *Grid Based Medical Devices for Everyday Health* project, in which patients who have left hospital are monitored using wearable computing technology. Since the patient is mobile, position and motion information is gathered (using devices such as accelerometers) to provide the necessary contextual information in which to interpret the physiological signals. The signal processing occurs on the Grid and medics are alerted – by pervasive computing – when the patients experience episodes that need attention.

The interesting infrastructure research question is to what extent the Grid services paradigm can be deployed in the direction of the devices. The project has been conducted using Globus Toolkit 3. The devices and sensors typically have limited computational power and storage and only in some cases may be capable of hosting Grid Services, generally interfacing via a Grid service proxy instead. Intermittent network availability had been found to be a critical problem, and the project has developed a framework for supporting both mobile grid clients and services in an intermittently connected network environment – it has produced a standard Web Services interface and associated software toolkit which provides a generic mechanism for exposing mobile/remote sensing devices as Grid services.

The additional contextual information provided by the wearables – GPS and accelerometer data – is essential to interpretation of the physiological signals. Modelling of context, reasoning about it and managing it, calls once again upon Semantic Web technologies. The project also features an information portal which can itself be accessed by pervasive devices, and can be used to access physiological data but also for monitoring and diagnosis of the deployed system.

The infrastructure of this project is shared with another, involving environmental monitoring in the Antarctic. A number of related projects involving mobile sensing illustrate the increasing volume of data that can be collected on an everyday basis and the science that this enables [30].

VI. RESEARCH CHALLENGES

There are challenges in realising this symbiosis. For example, an infrastructural challenge exists in allowing devices to have a place on the Grid as they often lack the power required to support the web services styles required by many infrastructures. Furthermore the communication patterns differ from a traditional file-compute grid: devices may interact using asynchronous event notification, or they may provide continuous data. Bridging between the grid and these devices raises fundamental issues. Several challenges are articulated in [32], which identifies the need for infrastructure advances in semantic grid, trusted ubiquitous systems, rapid customised assembly of services and autonomic computing systems. Here we consider the issues that arise from our seven points of symbiosis:

1) Devices need the Grid for computation. Grid applications are already being used to handle the data deluge from sensor networks, but the use of grid services on-the-fly by pervasive applications is not yet established, nor the negotiation mechanisms to select appropriate grid services when many are available.

2) Devices need the Grid for integration. We see examples of this, as in the Grid Based Medical Devices project, but on a project-specific basis. A framework is needed to bring devices into the Grid in a, flexible, open and interoperable way. This challenge is attracting attention in the Global Grid Forum through the Appliance Aggregation Architecture Research Group [32].

3) The Grid needs devices to interface with the physical world. The e-Science programme has been application-led and hence usability is a key issue. While many projects interface through a graphical user interface or portal, there are examples of pervasive approaches, as in the CombeChem project. These ideas will mature with further case studies and emerging methodologies. The Semantic Grid aspect is important here, and suggests a new research area in what we could call ‘semantic interactive systems’.

4) Virtualisation. Pervasive computing and the Grid are both distributed systems that, at the appropriate level of abstraction, pose similar challenges in terms of resource description, discovery and composition, in a world where multiple applications compete for resources that may only be intermittently available. The description mechanisms have yet to be agreed, and the techniques for service discovery and composition are not very sophisticated at this time. Semantic Web Services proposals exist but are still some time away from standardisation.

5) The information systems perspective. The principles of information systems design for Grid and pervasive applications are emerging as new systems are designed and deployed, but this perspective does not generally attract much attention in the pervasive or Grid communities. Semantic Web technologies should be part of this.

6) Grid computation on networks of devices. This has attracted very little work so far. Given the inevitable deployment of larger numbers of devices, and the increasing computational power of some deployed devices, the opportunity to create local grids is set to increase.

7) Self-Organisation. Autonomic computing is beginning to attract attention and there are some examples of systems that have some self-configuration, self-optimisation and self-healing properties. These are not the norm, and realising the autonomic vision perhaps requires a paradigm shift in approaches to system design, which will not happen overnight.

Addressing these challenges often requires collaboration across disciplines. The need to build bridges between Grid and AI is discussed in [33] and between Grid and Ubiquitous

Computing in [34]. Many of these efforts have adopted a service-oriented approach based on Web Services, as depicted in Fig. 2.

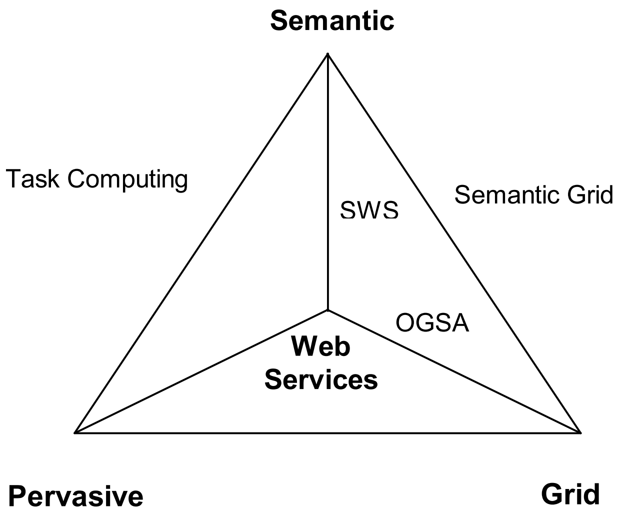

Fig. 2. A service-oriented architecture is being explored through activities in the Open Grid Services Architecture (OGSA) and in Semantic Web Services (SWS).

VII. CONCLUSIONS

In this paper we have discussed the relationship of developments in Grid technologies and pervasive computing and the mutually beneficial bond that exists between the two areas. With advances in semantic capabilities we are able to realise a vision of an intelligent connected world with pervasive computing systems providing personalised access to content, applications and services. The possible applications are numerous and the case studies provided although illustrative, by no means span the spectrum. They do however give an insight into what is likely to become the “normal” mode of operation. Clearly there are many opportunities for social and environmental benefits from these technologies as well as new methodologies for scientific research.

As we have noted, however, there are still many challenges to be met and amongst those not discussed above are security and trust, where there is a substantial amount of effort in both pervasive computing and Grid technologies. Together with the challenges in our seven areas of symbiosis there are clearly many interesting research problems to be solved. It is clear that there is still quite a way to go down the road to coupling Pervasive Computing with Grid systems, but the first steps have been taken and the fundamental building blocks are in place.

REFERENCES

[1] De Roure, D. Jennings, N.R. Shadbolt, N.R. The Semantic Grid: Past, Present, and Future, Proceedings of the IEEE, Volume 93, Issue 3, March 2005, Pages 669-681.

[2] Tony Hey and Anne E. Trefethen, 'The UK e-Science Core Programme and the Grid', Future Generation Computer Systems 18 (2002) 1017-1031

[3] L.Roberts, quoted in ‘Nerds: A Brief History of the Internet’, by S.Segaller, New York:TV Books (1998).

[4] The DAME Project: www.cs.york.ac.uk/dame

[5] Mark Gaynor, Steve Moulton, Matt Welsh, Ed LaCombe, Austin Rowan, and John Wynne, Integrating Wireless Sensor Networks with the Grid. IEEE Internet Computing, July 2004.

[6] Kevin A. Delin, Shannon P. Jackson, David W. Johnson, Scott C. Burleigh, Richard R. Woodrow, J. Michael McAuley, James M. Dohm, Felipe Ip, Ty P.A. Ferré, Dale F. Rucker, and Victor R. Baker, “ Environmental Studies with the Sensor Web: Principles and Practice”, Sensors, Vol. 5, 2005, pp. 103-117.

[7] Glynne-Jones, P. and White, N. M. (2001) Self-powered systems: A review of energy sources. Sensor Review 21(2):pp. 91-97.

[8] NGG Expert Group, Next Generation Grids 2 – Requirements and Options for European Grids Research 2005-2010 and Beyond, IST Programme Grid Technologies Unit, July 2004.

[9] IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans: Volume 35, Issue 1. Special Issue on Ambient Intelligence, IEEE Press, January 2005.

[10] T.Berners-Lee, J.Hendler and O.Lassila, ‘The Semantic Web’, Scientific American, Vol. 284, No. 5, pp 34-43, May 2001.

[11] OWL-S: Semantic Markup for Web Services, W3C Member Submission November 2004, see http://www.w3.org/Submission/2004/07/

[12] D.De Roure, N.R.Jennings and N.R.Shadbolt, ‘Research Agenda for the Semantic Grid: A Future e-Science Infrastructure’, National e-Science Centre, Edinburgh, U.K., UKeS-2002-02, Dec. 2001.

[13] D.De Roure, N.R.Jennings and N.R.Shadbolt, ‘The Semantic Grid: Past, Present and Future’, Proc. IEEE Vol.93, No.3, March 2005, pp669-681

[14] The myGrid Project: www.mygrid.org

[15] C.Wroe, C.Goble, M.Greenwood, P.Lord, S.Miles, J.Papay, T.Payne and L.Moreau, ‘Automating Experiments Using Semantic Data on a Bioinformatics Grid’, IEEE Intelligent Systems, Vol. 19, No. 1, January/February 2004, pp 48 – 55

[16] S.Miles, J.Papay, V.Dialani, M.Luck, K.Decker, T.Payne and L.Moreau, ‘Personalised Grid Service Discovery’, IEEE Proc.-Software, Vol. 150, No. 4, pp252-256,2003.

[17] Thomas Strang, Claudia Linnhoff-Popien: A Context Modeling Survey. First International Workshop on Advanced Context Modelling, Reasoning and Management as part of UbiComp 2004 - The Sixth International Conference on Ubiquitous Computing, Nottingham/England. Jadwiga Indulska and David De Roure (eds), September 2004.

[18] Ryusuke Masuoka, Yannis Labrou, Bijan Parsia and Evren Sirin, "Ontology-Enabled Pervasive Computing Applications," IEEE Intelligent Systems, vol. 18, no. 5, Sep./Oct. 2003, pp. 68-72.

[19] Ryusuke Masuoka, Bijan Parsia and Yannis Labrou, "Task Computing - The Semantic Web meets Pervasive Computing -." In D. Fensel et al. (Eds.), "The Semantic Web - ISWC 2003," the Second International Semantic Web Conference (ISWC 2003), Sanibel Island, FL, USA October 2003 Proceedings, LNCS 2870, pp. 866-881, 2003.

[20] O. Storz, A. Friday and N. Davies, Towards `Ubiquitous' Ubiquitous Computing: an alliance with `the Grid'. First Workshop on System Support for Ubiquitous Computing Workshop (Ubisys 2003) in association with Fifth International Conference on Ubiquitous Computing. Seattle, US, 2003.

[21] T. Kindberg and A. Fox: System Software for Ubiquitous Computing. IEEE Pervasive Computing, vol. 1(1): pp. 70–81, 2002.

[22] Abelson, H. et al, Amorphous Computing, Communications of the ACM, Volume 43, Issue 5 (May 2000), Pages: 74 - 82

[23] D. De Roure, "On Self-Organization and the Semantic Grid," in IEEE Intelligent Systems, vol. 18, 2003, pp. 77-79.

[24] J.G.Frey, M.Bradley, J.W.Essex, M.B.Hursthouse, S.M.Lewis, M.M.Luck, L.Moreau, D.De Roure, M.Surridge and A.Welsh, ‘Combinatorial Chemistry and the Grid’, published in ‘Grid Computing: Making the Global Infrastructure a Reality’, edited by F.Berman, G.Fox and T.Hey, Wiley (2003)

[25] G.Hughes, H.Mills, D.De Roure, J.G.Frey, L.Moreau, m.c.schraefel, G.Smith and E.Zaluska, ‘The semantic smart laboratory: A system for supporting the chemical e-Scientist’, Org. Biomol. Chem. Vol. 2, No. 22, pp3284-3293, 2004.

[26] J.G.Frey, D.De Roure and L.A.Carr, ‘Publication at source: Scientific communication from a publication web to a data grid’, Euroweb 2002 Conf., ‘The Web and the Grid: From e-Science to e-Business’, 2002.

[27] m.c.schraefel, G.Hughes, H.Mills, G.Smith, T.Payne and J.Frey, ‘Breaking the Book: Translating the Chemistry Lab Book to a Pervasive Computing Environment’, published in Proceedings of the Conference on Human Factors (CHI), 2004.

[28] m.c.schraefel, G.Hughes, H.Mills, G.Smith and J.G.Frey, ‘Making Tea: Iterative Design through Analogy’, published in the Proceedings of the Conference on Designing Interactive Systems, 2004.

[29] Bachler, M.et al. Collaborative Tools in the Semantic Grid. In Proceedings Semantic Grid Applications workshop, GGF11 - The Eleventh Global Grid Forum, June 6-10, 2004, Honolulu, Hawaii, USA

[30] Equator e-Science projects, see http://www.equator.ac.uk.

[31] M. Atkinson, J. Crowcroft, C. Goble, J. Gurd, T. Rodden, N. Shadbolt, M. Sloman, I. Sommerville, and T. Storey, Computer Science Challenges to Emerge from e-Science, 2002.

[32] Karan Bhatia, Nobuhiro Koba, Dimitris Lioupis, Ashley D. Lloyd, Sergio Mendiola, Dejan Milojicic, Krishna Sankar & Ian Taylor Appliance Aggregation Architecture Terminology, Survey, and Scenarios. GGF Appliance Aggregation Architecture Group (APPAGG), March 2003.

[33] Goble, C.A. and De Roure, D.: The Semantic Grid: Myth Busting and Bridge Building. ECAI 2004: 1129-1135