Data-Intensive Ubiquitous Computing Needs the Grid

Erica Y. Yang Jie Xu

School of Computing
University of Leeds
Leeds LS2 9JT
{ericay, jxu@comp.leeds.ac.uk}

Abstract

The rapid advances of wireless network technologies together with the undergoing commercial deployment of sensors shed lights on many aspects of the practicability of large scale Ubiquitous computing (UbiCOMP). The implication of this development is that we now have to handle data at an amount and density previously unattainable. Real-time data processing and analysis are two key capacities for realising the full potential of such data. Existing distributed systems, in particular in the commercial domains, cannot currently cope with the data challenges arise in such data intensive ubiquitous computing systems. The Grid, as a technology introduced to handle data intensive applications, is an ideal candidate to harness and utilize the massively available information produced such UbiCOMP applications. We discuss the unique research challenges resulting from such the integration.

1. Introduction

It has been well over a decade since the concept of UbiCOMP was proposed by Mark Weiser in 1991. We have yet to see the flourish of UbiCOMP applications in our every life and work. Several researchers have argued that the lack of *a common system infrastructure* is the cause of the problem and suggest that the Grid, as a ubiquitous, uniform, standardised system platform, holds the promise of the future of deployment of UbiCOMP in a wide scale [SFD03]. The authors in [SFD03] have also explored the need for the synergy as "the requirements for building the Grid significantly overlap with those for supporting ubiquitous computing applications" [SFD03]. This paper arrives at the same conclusion: the combination of the Grid and UbiCOMP is crucial to the success of both but on a different ground. The cause of this synergy is *far more complicated than* a simple combination or just a quest for a common infrastructure as there exists decades of experience of building distributed computing platforms.

On the one hand, the huge amount of data that is being and are going to be produced by emerging data intensive UbiCOMP applications pushes the demand for a coherent integration with the Grid (we shall explore this point in Section 4). On the other hand, emerging UbiCOMP applications can extend the potential of the Grid to commercial domains in a very timely fashion. This is important to the long term sustainability of the Grid as it is strategically important that the Grid is *not* something that is *only useful* in a scientific research domain.

Certain ubiquitous computing technologies are undergoing large scale deployment in commercial domains. This is typified by the latest mandatory use of Radio-Frequency Identification (RFID) technology [WCD05] in the top 100 suppliers of the US retail giant - Wal-Mart and in all goods in the supply chain of the US Department of Defense (DOD) [zdnet]. Coupled with the rapid advances of wireless sensor technologies, the practicability of large scale UbiCOMP has never been so promising.

2. Existing UbiCOMP Research

As described in [Wei93], "Ubiquitous computing has as its goal the enhancing computer use by making many computers available throughout the physical environment, but making them effectively invisible to the user". This vision of Ubiquitous computing (UbiCOMP) has been around for well over a decade. However, many UbiCOMP projects remain in small scale in that they are mostly *isolated* experiments involving objects *inside* a confined space (e.g. a house or a city). To realise Weiser's vision, we have much challenges to tackle.

Most early work of UbiCOMP has been focused on demonstrating its possibility by making small computing devices, enabling wireless networking, or writing software to manipulate these small artefacts. Examples are Xeox's PARCTAB and AT&T's Active Badges. In the late 90's, Wireless networking technologies, such as Wi-Fi, later Bluetooth, open up a new era of UbiCOMP developments. Many recent projects in the area are based on these wireless standards and emphasis on the more practical aspects of UbiCOMP, such as usability, Human-Computer Interactions (HCI), software infrastructure support, and feasibility studies of novel application domains.

Most of the technologies are proprietary which are based on in-house hardware and/or software design. This is probably the main reason for its limited presence in a wide scale. There are other reasons which contribute to the slow adoption of UbiCOMP, for example, the demanding power requirement and cost of enabling wireless infrastructures. As Storz, Friday, and Davies pointed out in [SFD03], among many other reasons, the lack of a commonly available system infrastructure is considered to be one of the key obstacles [SFD03, DG02, KF02]. There are more fundamental reasons why we should put data intensive UbiCOMP and the Grid together. To answer this question, we must understand what the new types of large scale UbiCOMP applications are.

3. Large Scale UbiCOMP Applications

The latest advances of wireless sensor network technologies shed the light on enabling UbiCOMP applications in a massive scale, which typically include supply chain management, inventory control, building security and access control, environmental monitoring, logistics control, factory automation and control, and massive scale health care monitoring. In fact, the possibilities in commercial domains brought by these advances are numerous and the promise of realising them in the near future is encouraging. The excitement of the new advancements can be summarised in the following three aspects [GF05]:

1. the massive deployment of low-cost and low-power sensors makes it possible to continuously collect reliable context aware information (e.g. identification, location, temperature, time);
2. the standardised wireless networking technologies (e.g. ZigBee - http://www.zigbee.org) and open source embedded systems platforms (e.g. TinyOS – www.tinyos.net) means that it is easier to integrate sensors and processes data collected from geographically distributed sensor networks; and
3. the business incentive for processing and aggregating object (e.g. goods, vehicle) monitoring information, possibly in real-time or on a regular basis, has never been higher.

However, the advances in sensor network technologies only answer part of the problems posed by this exciting new territory of UbiCOMP. To enable truly meaningful applications, it is mandatory that the pervasive provision of system level infrastructure to effectively and efficiently harness these massively available “islands of information”. That is because such information is often collected from highly dynamic and flexible application settings in different deployment environments, ranging from big businesses with well established IT infrastructure to small businesses with extremely limited IT resources.

Large scale UbiCOMP applications involve a massive amount of computers that inhabit in geographically distributed spaces and are capable of gathering information about the physical world through sensors. All such applications share a common feature: they all produce a massive amount of data that demand a significant level of computation and data storage resources which is often too demanding for any single organisation to offer. We shall begin with a description of exiting e-Science applications and then present some emerging applications in this area.

A. Astrophysics

Some existing e-Science/Grid projects can be viewed as data intensive UbiCOMP applications as they use sensors (e.g. telescope and environmental sensors) to detect and capture information about the physical world. Several virtual observatories, including USA’s NVO (http://www.us-vo.org), Europe AVO (http://www.euro-vo.org), and the UK AstroGrid (http://www.astrogrid.org), all are examples of using astronomy data collected from lots of geographically distributed telescopes from all over the world to construct a virtual observatory.

B. Environmental Science

NASA’s Sensor Web project [DJJ+05] developed a pervasive environmental monitoring system over large spatial areas with the use of low-power and low-cost wireless networked sensors to provide continuous and monitoring presence. So far, this project has explored a wide range of application areas, ranging from wild life monitoring, to Antarctica condition monitoring. Environmental monitoring projects, like these, can be applied to a much wider applications in urban city environment, such as traffic monitoring.

C. Wide Scale Commercial Applications

Most of the projects described above are largely research oriented by aiming to serve the scientific research community. The more exciting development of data intensive UbiCOMP are in commercial domains, such as predictive maintenance, supply chain management, industrial automation, and logistics management. Due to the space restriction, we only select two typical applications for our discussion.

Predicative Maintenance

The Distributed Aircraft Maintenance Environment (DAME) project[1] relies on sensors attached to jet engines to collect diagnostics data (e.g. vibration and performance) of flight engines while the plane is in flight. Rolls-Royce hopes to exploit Grid technologies to carry out fleet-wide predictive maintenance for around 100,000 engines that are currently in service [HT03].

The Heterogeneous Sensor Networks project being carried in Intel research aims to enable predictive maintenance in industrial environments. For example, BP is carrying out a trail experiment with this technology to continuously monitor engine vibration on BP's oil tankers [bp].

Location Tracking

RFID is simply a tagging technology with radio communication capability. Although this technology has been around for several decades and its potential has been well recognised, it has not really taken off until recently. The technology has been severely restricted by the cost of deployment, tag designs, and network connectivity. The recent massive deployment of RFID technology in the supply chains of the major US retailer – Wal-mart and in that of the Department of Defense (DoD) of the US government means that the low-cost massive manufacturing of RFID capable objects will quickly become a reality [zdnet]. Coupled with the pervasiveness of networking technologies and low-cost sensors, it is expected that RFID technologies will prevail in numerous commercial applications [FG02].

4. The Quest for the Grid

This section discusses the fundamental motivation for utilising the Grid for data intensive UbiCOMP applications, by taking a close look at their data challenges brought by the applications described in the previous section.

A. Existing Data Intensive UbiCOMP Applications

It is well known that the data volume generated in Astrophysics and environmental science projects is in the order of Terabytes per year [HT03]. In DAME, the daily engine vibration and performance data can easily mounted to Gigabytes per engine per day [dame]. Such amount of data is too overwhelming for any single organisation to handle, i.e. capture, transmit, filter, classify, analyse, store, and act upon. That is precisely the motivation of introducing the Grid to harness the vast amount of data generated in these applications.

B. Emerging Data Intensive UbiCOMP Applications

What is much less known is the data challenges introduced by the RFID technologies. In a more general term, it should be sensor technologies, as some types of RFID tags are bundled with sensors. For example, Alien Technology's ALR2850 kit provides RFID tags with continuous real-time temperature readings. It has been argued that "RFID is the stepping-stone to sensor networks" [Mul05]. For the simplicity of presentation, the following discussion will be focussed on the data challenges brought by RFID technology. We expect the amount and density of data generated by sensor networks will be at a similar level.

According to the current de-facto RFID standard – EPCGlobal [epc], each RFID tag has a unique identifier of at least 12-byte long. Let us just take Wal-mart's recent RFID deployment as an example to illustrate the amount of data we are talking about. In the long run, Wal-mart aims to use RFID to trace all cases and pallets as it is expected that this technology can produce huge saving (in the order of billions per year) for the retailer if RFID is "fully deployed throughout its supply chain and in stores" [Rob03]. Based on the estimation method developed in [BM05], we can do a similar but much conservative calculation as follows. It is estimated that Wal-Mart receives from eight to ten billion cases *a year* [Mur03]. Let's say 100 million cases per day are actually in the supply chain. Even if only one tenth these cases *each day* are RFID enabled, that will be 10 millions (i.e. $10 \times 10^9 / 10^3 = 10 \times 10^6$) cases in Wal-

[1] http://www.cs.york.ac.uk/dame/

mart's supply chain on a daily basis. The total amount of RFID data of these cases would be 0.12 gigabytes (12bytes x 10 x 10^6 = 0.12 x 10^9 bytes). If all these cases are read once every 5 minutes for 10 hours a day, the amount of data would be around 14 gigabytes (0.12 gigabytes x 12 times x 10 hours) per day. This is in fact a rather *conservative* estimation which is at a *case-level* tagging rather than at an item-level tagging. In [BM05], the daily estimation of RFID data collected by one major retailer amounts to 15 *terabytes* per day on an *item-level* tagging.

C. The Quest for the Grid

This amount of data is no less than the amount of data handling by the existing data Grid. The Grid is specifically designed and developed to cope with such huge amount of data by embracing the power of lots computers to offer unmatchable processing and data storage capabilities. Therefore, we believe that the Grid is an ideal for carry out the task for such data intensive UbiCOMP applications. The importance of introducing the Grid to such applications is further reinforced by the fact that little enterprise computer systems are designed to handle this level of data volume in real-time.

Data intensive UbiCOMP applications can also leverage the Grid's existing software infrastructure and standards to enable data and resource sharing among multiple organisations involved in large scale UbiCOMP applications. It needs an enormous amount of processing power to deal with the mounds data in a large distributed environment where many organisations (suppliers, warehouses, manufacturers etc.) are concurrently involved to manage different portion of the entire RFID data in a supply chain. To fully realise the benefit of RFID, it is recognised that data has to be shared among the organisations involved at varying levels depending on the status of the goods [BH05]. The widely Grid standards, relatively mature Grid technologies, much Grid experience we have gained over the past few years, and most importantly, the readily available large scale resource-sharing and problem-solving platform all certainly make the Grid an ideal candidate to cope with the above challenges.

5. Setting the Research Agenda

The undergoing transformation of UbiCOMP applications bring several unique research challenges that call for new solutions. This section focuses on the challenges which we believe will produce significant impacts on the existing Grid technologies if the Grid is used to enable large scale data intensive UbiCOMP applications. These challenges go way beyond the conventional challenges (e.g. HCI, power management, routing, and data aggregation) widely recognised in existing UbiCOMP research.

A. Data Engineering

Continuously monitoring and tracking not only produce a vast amount of data but also collect data at a dense rate previously impossible. The pervasiveness of data intensive UbiCOMP applications in a wide scale means that data streams will be processed in a ubiquitous way at different locations and be aggregated from highly distributed data sources at a varied rate. The density of these data streams varies significantly between applications. It depends on the amount of sensors deployed in an area, and the mobility, reliability of the deployed sensors. Small data streams may aggregate into a surge of data streams which may well be possible to disperse into small streams again. For example, the continuous transportation of RFID tagged goods in a supply chain means that the mobility of the goods and the environment determines the varied densities and the origins of the readings. Therefore, such applications must be able to deal with a sudden surge of a large among of data steams in a very short time scale. In other times, there may be little data coming in. This requires systems to dynamically adapt to the ever-changing situations. All these characteristics of data intensive UbiCOMP applications demand a re-examination of the scalability and flexibility of existing Grid protocols, in particular resource allocations and scheduling.

B. Formation of VOs

Virtual Organisation is a fundamental concept in the Grid. To apply Grid technologies for data intensive UbiCOMP, a more flexible collaboration model should be supported to accommodate a wide range of business models. For example, the business relationship among suppliers and supermarkets highly depends on the current market situation.

C. Charging and Auditing

This is an issue that demands significant attentions from the Grid community. Currently, most Grid Toolkits (e.g. Globus) provide little support for charging and auditing [Sie05]. To support commercial large scale data intensive UbiCOMP with the Grid, this must be properly addressed.

6. Summary

As the rapid prevalence of RFID like ubiquitous computing technologies, UbiCOMP is at its turning point. If such technologies are proven to be successful, the next move will be naturally extended to large scale deployment of sensor network technologies. This will bring us one step further to realise Weiser's UbiCOMP vision. Without a proper infrastructure support, UbiCOMP will remain to be a vision rather than something that is ready to be exploited in a large scale. It is now the crucial time for unifying the Grid community and the UbiCOMP community so that we can make UbiCOMP a reality in our daily life and work.

However, the introduction of the Grid to enable large scale data intensive UbiCOMP applications is not simply because of the lack of a common system infrastructure. Fundamentally, it is because of the data challenges brought by the prevalence of such applications. The data mounds generated in such applications indicate that it will demand overwhelming resources on existing IT infrastructure. Without the Grid, handling such data in real-time will be demanding for any *single* organisation. Hence, a flexible and dynamic system infrastructure, such as the Grid, can really remedy the problem to a large extend.

We also discuss the research challenges that pertain to the emerging large scale Ubiquitous computing applications. We have been focused on the technical aspects of such applications. The pervasiveness of large scale UbiCOMP will clearly bring the privacy issue into the spotlight of our research agenda, which clearly deserves further attentions.

References

[BM05] Manish Bhuptani and Shahram Moradpour, "RFID Field Guide – Deploying Radio Frequency Identification Systems", Prentice Hall (PTR), 2005.

[DG02] N. Davies and H. Gellersen, "Beyond Prototypes: Challenges in Deploying Ubiquitous Systems", IEEE Pervasive Computing, 1(1), 2002.

[DJJ+05] 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, http://sensorwebs.jpl.nasa.gov/resources/Delin_et_al_Sensor_Web_Hydrology_2005.pdf.

[FG02] Andrew Fano and Anatole Gershman, "The future of business services in the age of ubiquitous computing", Communications of the ACM, 45(12), Dec. 2002.

[GF05] Anatole Gershman and Andrew Fano, "Examples of Commercial Applications of Ubiquitous Computing", Communications of the ACM, 48(3), Mar. 2005.

[HT03] Tony Hey and Anne Trefethen, "The Data Deluge: An e-Science Perspective", in Wiley's "Grid Computing – Making the Global Infrastructure a Reality", Jan. 2003.

[KF02] T. Kindberg and A. Fox, "System Software for ubiquitous Computing", IEEE Pervasive Computing, 1(1), 2002.

[Mur03] Jean V. Murphy, "Get Ready! Wal-Mart Mandate Puts RFID, Smart Tags on Fast track", SupplyChainBrain.com, Sept. 2003, http://www.glscs.com/archives/09.03.rfid.htm?adcode=10.

[Mul05] EuroTag, "Intel and Oracle Form Alliance Around RFID and Sensors", http://www.eurotag.org/news/archives/Apr_2005/Intel_and_Oracle.html.

[Rob03] Mark Roberti, "Analysis: RFID – WalMart's Network Effect", CIO Insight, Sept. 2003, http://www.cioinsight.com/print_article/0,3668,a=61672,00.asp.

[SFD03] Oliver Storz, Adrian Friday and Nigel Davies, "Towards 'Ubiquitous' Ubiquitous Computing: an alliance with the Grid", System Support for Ubiquitous Computing Workshop at UbiCOMP'03, Seattle, USA, October 2003.

[Sie05] Frank Siebenlist, the presentation of "The Globus Security Architecture", U.K. e-Science Core Programme Town Meeting – "Security for e-Science, approaches and interoperability", Monday 11th April 2005, London.

[WCD05] The Webopedia Computer Dictionary, a definition of RFID (Radio Frequency Identification) technology, http://www.webopedia.com/TERM/R/RFID.html.

[Wei93] Mark Weiser, "Hot Topics: Ubiquitous Computing", IEEE Computer, October 1993.

[zdnet] Matt Hines, "Barcode Replacement comes in from cold – Special Report on Radio Frequency Identification (RFID)", ZdNET insight special report series, UK, http://insight.zdnet.co.uk/specials/rfid/0,39026568,39154051,00.htm.

