1. Introduction

Current trends in mobile computing devices and wireless networking technologies are helping to broaden the domain of computing. Increasingly, computing is occurring within settings far removed from the desktop, such as in the car or the shopping centre, where computing applications compete with other tasks for user attention. The spread of computing to these new environments, supported by a variety of interconnected mobile and embedded devices, is commonly referred to as pervasive or ubiquitous computing.

Mobile computing devices have become increasingly invaluable to users across a variety of fields undergoing continual improvements in computing power, storage and connectivity. Mobile devices have become ubiquitous as advances in technology have meant these devices are no longer just electronic replacements for paper diaries. Current models can perform a wide range of functions, from document processing, communications, global positioning systems (GPS) and entertainment hubs. Embedded devices are also increasingly being used to enable novel applications such as automated inventory management, which involves the augmentation of goods with chips that can be interrogated by remote readers and programmed with information such as shipping and expiry dates. Similarly, computational capabilities are being added to a variety of consumer products, leading to appliances such as Internet-enabled refrigerators that support browsing recipes and online grocery stores, or systems for smart homes that can be remotely controlled through mobile device interfaces.

The mobility of computing devices, applications and people leads to highly dynamic computing environments. Unlike desktop applications, which rely on a carefully configured and largely static set of resources, pervasive computing applications are subject to changes in available resources such as network connectivity and input/output devices. Moreover, they are frequently required to cooperate spontaneously and opportunistically with previously unknown software services in order to accomplish tasks on behalf of users. Consequently, pervasive computing software must be highly adaptive and flexible. Similarly, pervasive computing applications must accommodate changes in user requirements that arise as the user's activities and goals evolve. As an example, an application may need to modify its style of output following a transition from an office context to a meeting room, in order to be less intrusive.

When humans interact with other people in the surrounding environment, they make use of implicit situational information. People naturally assume and interpret the context of the situation they are in and respond fittingly. Computers are not as good as humans in interpreting situational information from their environment and in using it in interactions. Computers cannot easily take advantage of such information in a transparent way, and if they can, they usually require that it is explicitly provided. In a mobile computing environment, it will be ideal if a user can obtain services and information according to the current location and activity. In the future, the applications of mobile computing will continue to multiply, with the result that humans will inhabit environments in which vast numbers of invisible computing devices will provide seamless support for a broad spectrum of human activities. Consequently, the success of this technology in these environments is dependent on their ability to adapt to the available resources and limitations.

Context-aware computing is a compile computing paradigm which proposes that applications can discover and use to take advantage of contextual information. Context-aware applications exploit information about the context of use, such as the location, tasks and preferences of the user, in order to adapt their behaviour in response to changing operating environments and user requirements. Context-aware pervasive systems are becoming increasingly important and research has begun to investigate context-aware systems more generally, independently of specific applications, including context middleware, toolkits and ontologies to provide vocabularies to describe context. Research in pervasive computing and, in particular, context awareness has not yet addressed in a generic way the problem of adapting to changing situations, including events that are not recognisable by the particular application (Henricksen and Indulska, 2006). In the adaptation process, many applications make use of human-user interaction to guide the system, which can often be intrusive or complex to be handled by a user. Many current applications deal with only a finite set of foreseeable contexts, each employing specific methods in order to handle and adapt to well-known predefined context environments.

A problem that arises is the inability to respond to contextual information that cannot be classified to any known context. Many context-aware applications require all discovered information to exactly match a previously known type of context; otherwise, the application will not react responsively. The ability to learn and recall contexts based on the contextual information discovered have not been very well addressed in previous context-aware applications and research. Additionally, a mechanism to reason about the contextual information being discovered from past application interactions will be beneficial to induce contexts for future experiences.

1.1 Research Aim and Objectives

The aim of this research is to develop a component middleware technology for mobile computing devices in the discovery and capture of contextual information, using the situated reasoning concept of constructive memory. This research will investigate the ability to use constructive memory and past experiences to induce possible contexts, given a small amount of training and application interaction. This research defines context or context environment as *place* and does not identify people, tasks or activities. A *place* often describes a unit of context which users may wish to be notified about. Labelling a geometric region assigns it semantics and can help prune and improve place predictions (refer to Section 1.5 for an extended definition of *place*).

Depending on what contextual information is discovered from the current place, and what has previously been stored, this system will interpret this information to determine an appropriate context. This interaction will then be stored as an experience and will form a basis from which the system can learn and recall contexts to determine future contexts, based on discovering contextual information of previous experiences. To be able to reason about contextual information and construct the memory of past experiences in uncertain environments of context will be important in advancing context-aware applications, as not all contextual information will be required. The basic premise is to perceive, construct from memory and reason about a context environment.

To achieve the stated aim, the following objectives are established:

- analyse the research disciplines of context awareness and memory, identifying the approaches adopted and the challenges posed;
- highlight the fundamental areas of context awareness and memory in which solutions fall short of requirements;
- propose a framework for conceptual components and modified architecture for context awareness, incorporating a constructive memory model;
- design, implement and deploy a middleware application of this modified architecture model for mobile computing;
- conduct a variety of experiments to investigate the behaviour of the middleware application; and
- evaluate the overall approach of this model and application from the experiments conducted.

1.2 Motivation

This research has had a number of motivating factors. This initial motivation for this research has come from mobile technologies offering experiences which can effectively engage users and are often markedly different from those afforded by desktop computers. These devices are used dynamically, in many different settings, giving access to a broad range of uses and situational experiences. The personal nature of these technologies means they are well suited to engaging users in individual experiences and to giving an increased ownership and responsibility over their use. People continuously learn through their experiences and utilise knowledge whenever they want to carry out some tasks. Simply, human interaction in a particular context depends on the knowledge gained through their past experiences and interactions. This research is motivated by the idea that mobile devices enable the discovery of contextual information and they are well placed to enable memory construction and learning for future contexts, based on the recollection of past experiences and interactions. The focus of contextual information in mobile computing concerns the ability of a mobile device to detect, interpret and respond to aspects of the user's physical environment. This is very similar to the processes of constructive memory in perceiving, constructing and reasoning about an experience. The effect of these processes, in context awareness and constructive memory, promises to add value to existing uses, as well as creating new types of applications. Mobile devices that intelligently act autonomously on knowledge in these environments will be useful in the future.

The research is also motivated by the proliferation of stand-alone and embedded computing devices in work and home contexts. Combined with a variety of networking technologies, they increase the importance of context awareness in distributed applications. There is an interesting issue in the ways knowledge is characterised and used. The view of context awareness is that some knowledge is required, such as information about the user or the user's environment. Nevertheless, this is not adequate. The lack of ability to reason about context and learn from experiences and interactions motivates this research. Within a constructive memory model, the very act of constructing a memory of an associated experience affects the memory through its grounding in the interaction with the context. Any knowledge about the current context is used as cues in the memory construction process. This is what distinguishes it from a standard information retrieval system whereby the knowledge evolves and is allowed to operate within a context that is not pre-programmed into the system.

A final motivation comes from the observation that computers have become the vehicle for an increasing range of everyday human activities. Pervasive computing aims to provide computation ubiquitously and its vision is to provide a user with the transparency of any where, any place, any time computing. According to Wesier (1999), the most profound technologies are those that disappear and weave themselves into the fabric of everyday life until they are indistinguishable from it. Pervasive computing's highest ideal is to make a computer so embedded, so fitting, so natural, that it is used without even our thinking about it. This means a shift to human-centred computing, where technology is no longer a barrier, but works for us, adapting to our needs and preferences and remaining in the background until required. This implies a change in the relationship with technology to a much more natural way of interacting and using the power of networked computing systems, which will be connected not just to the Internet or other computers, but, to places, people, devices and objects in the world around us. An increase in the ratio of computing applications to humans is steadily reducing the amount of attention that users can expend upon each application. The design of applications that possess the required degrees of autonomy and invisibility without removing the users' sense of control presents a significant motivation for this research. This trend coincides with the shift of computing away from the desktop and into environments in which human-computer interaction must take a back seat to other activities. These factors have led to the proposal of invisible computing (Norman, 1998); a novel design approach will see context-aware applications provide users with unobtrusive support for the tasks at hand.

1.3 Research Claims

This thesis contends that the traditional methods in training context-aware mapping are ill-equipped to meet the challenges inherent in developing applications for future intelligent pervasive computing environments. The problem exists in training the system in all possible context sensor data and context-type environment combinations, which is a long and arduous process. The research contribution of the thesis lies in developing a modified architecture for context-aware systems, using a constructive memory model as a way to learn and recall context. Using a constructive memory model, previous experiences can be induced to construct potential contexts, given a small amount of learning and interaction. The learning process is able to map the many variations of context sensor data currently discovered with the user with a predicted type of context seen previously. It only requires a small amount of contextual information to predict a context, something common context-aware systems lack, as they require all information before a type of context is assigned. Moreover, such learning algorithms are typically processing-power and memory-draining intensive. However, this learning algorithm is designed to be lightweight and implemented on resource-constrained mobile devices for developing context-aware applications. The application prototype implementation of the framework incorporating constructive memory is presented as the basis for the claims and performance studies to illustrate these claims.

1.4 Organisational Structure of Thesis

A review of the related work and the theoretical foundations underpinning this research are introduced in Chapter 2. Chapter 3 identifies issues relating to the literature reviewed and based on these issues, describes the conceptual framework. Additionally, the proposed modified architecture used in this research is introduced. The implementation of this framework and architecture for the middleware application are addressed in Chapter 4. Chapter 5 describes experiments that are conducted based on the prototype implementation and illustrates some of the claims of the research. Chapter 6 reviews and summarises the research, and concludes with the contributions made. Moreover, it presents future work and directions for this research.

1.5 Terminology

Throughout this thesis a number of key terms are used to describe the research. Some of the terminology used is reiterated in the thesis to help frame that section of the research. Listed below are some of the definitions of those terms to help understand what the research is describing and discussing:

- context – defines *place* in an environment. A *place* describes a region with a common use, such as a bank, a home or a park. In addition, a *place* often describes a unit of context which users may wish to be notified about. Labelling a geometric region assigns it semantics and can help prune and improve place predictions. More importantly, a label directly represents the place's demographic, environmental, historic, personal, or commercial significance and is the desired abstraction for emerging proactive applications (Scott, 2004). A *place* is rigidly defined, since being on the street outside your house is very different from being inside your garage, despite being only metres apart. A room scale is another natural human scale, and might be defined by one's field of view indoors, or visual range outdoors. Obvious applications include finding people or objects, such as the nearest printer. The categorisation of location resolutions is useful in identifying groups of applications with similar location requirements. It is also useful in the design and classification of context-aware systems. This thesis seeks to automatically label places by inferring the user's context, based on what objects are discovered in that particular place. For example, a library is a *place* where bookshelves are discovered, an office is a *place* where a desk and filing cabinet are discovered, and a computing lab is a *place* where computers and printers are discovered;
- contextual information – defines all the sensor and device data being discovered;
- percept – also referred to as a *device* in this thesis, defines a device or object discovered in a context;
- device – also referred to as a *percept* in this thesis, defines objects (both technological and non-technological) that may be discovered in a context;
- record – also referred to as an *experience* in this thesis, defines contextual information for a selected context saved and stored in the application;
- experience – also referred to as a *record* in this thesis, defines contextual information for a selected context saved and stored in the application; and
- memory – defines what experiences are saved and recalled by the application.