Augmenting Interaction and Cognition using Agent Architectures and Technology Inspired by Psychology and Social Worlds

Steve Goschnick and Connor Graham

University of Melbourne
Interaction Design Lab
Department of Information Systems
111 Barry St, Carlton
Victoria, 3010, Australia
+61 8344 1498
stevenbg@unimelb.edu.au | cgraham@unimelb.edu.au
http://idealab.dis.unimelb.edu.au

Abstract. Intelligent agents can play a pivotal role in providing both software systems and their augmented interfaces, to individuals from all walks of life, to productively interact with the Internet 24 hours a day, seven days a week (24x7) and with each other, over both wireless and broadband infrastructures. This paper describes a novel intelligent agent architecture called ShadowBoard with an equally novel agent-oriented user interface named ShadowFaces. Both the architecture and interface draw upon concepts from psychology. The architecture has been instantiated as a software framework also called ShadowBoard. We illustrate how aspects of user *cognition* were outsourced using the ShadowBoard framework, in developing a prototype lab manager agent - a 'Digital Self' - for booking services and providing advice to potential users of the Interaction Design Lab at the University of Melbourne. We then take the Locales framework from Computer Supported Co-operative Work, using it to both understand the problematic aspects of *interaction* involved in complex social spaces, and to identify specific needs for technology intervention in such social spaces. Guided by the Locales framework, we are adapting the sophisticated individual agent technology demonstrated in the Digital Self, into a multi-user, agent-augmented system dubbed ShadowPlaces. Our aim with ShadowPlaces is to outsource some of the *interaction* necessary, for a group of mobile individuals to interact cooperatively and effectively in a Social World supported by wireless networks, and backed by broadband Internet services. An overview of the architecture and methodology (ShadowBoard), the interface (ShadowFaces), the resulting framework (ShadowBoard) and a user implementation (a Digital Self), are all presented, and then progress on ShadowPlaces - the multi-user version - is outlined.

1. INTRODUCTION

A compelling vision of the future is one with pervasive and ubiquitous computing spanning our public and private lives. As broadband access to the Internet becomes widely available, and wireless networks extend our access to it, individuals require new levels of interaction with the ever-present system. At the outset of the ShadowBoard project - which is central to this paper - we envisaged having a software system on call 24 hours a day, seven days a week (24x7) to enable an individual user to have an effective and productive Internet presence, regardless of their current circumstance.

Agent-oriented analysis, design and programming has emerged as a paradigm designed to facilitate higher level modelling of user needs and the transformation of those needs into flexible and diverse systems [Wooldridge 1995]. Along with concepts from Psychology, our project adopted the agent-oriented paradigm to form a novel architecture for building a *Digital Self* - a continuous representation of the user via the Internet, which *monitors*, *summaries*, *alerts* and *notifies* the user appropriately, not over taxing their attention. Features of this architecture named ShadowBoard [Goschnick 2001] include:

- decomposition of a user's multiplicity of *roles* into a hierarchy of sub-agency
- relaxing of the autonomy of the sub-agents under control of an autonomous agent – the *Aware Ego Agent* – which is autonomous in the usual agent sense
- wrapping of external services and agencies including *web services* and utilising them as if they were internal sub-agents
- ability to apply *ontologies* at a localised level.

We have instantiated the architecture into a framework also simply known as ShadowBoard [Goschnick 2003a]. I.e. We are using the name ShadowBoard interchangeably, between the architecture, the accompanying methodology, and the enacted framework of tools.

Features of the accompanying methodology, which is inclusive within the user interface (named ShadowFaces), include:

- a generic range of sub-agent *types* with enough breadth to harness sub-agencies capable of servicing both work and recreational interests of an individual user 24x7, engaged in multiple social worlds.
- generic envelopes-of-capability within the agent types, which help the user identify, collect and organise a collection of sub-agents into useful groupings for further thought, discussion, design and implementation of a Digital Self.

With this methodology and framework of tools, we now have a development environment for rapidly building custom multi-agent applications from constituent sub-agents and services, often accessed using Internet protocols, for an individual user when connected to the Internet.

With the advent of wireless networks and the subsequent expansion of the Internet's reach into potentially every corner of an individual's life, we saw the need to expand our system and methodology to address the mobile aspects of a typical user's life.

Mobile systems pose serious challenges to the system designer. [Graham 2003] describe these as threefold: means of input are limited, screens are small and use contexts are dynamic. Of these three issues, highly dynamic use contexts are a challenge for the designer. The users of mobile systems are particularly challenged in their interaction with the device when the world about them is “fluid” [Kakihara 2002]. In fluid interactions “neither boundaries nor relations mark the difference between one place and another” [Mol 1994]. Boundaries and relations among places are permeable, transitory and transmutable: “…sometimes boundaries come and go, allow leakage and disappear altogether while relations transform themselves without fracture. Sometimes, then, social space behaves like a fluid” [Ibid].

This fluid interaction can also be understood as the user moving among multiple social worlds [Strauss 1976] or among orbits of collective action. In a social world a body of users is united in a shared activity and symbolisation. *Travelling*, *wandering* and *visiting* have been identified as three key activities that are pertinent to current work-related mobility [Kristoffersen 1999]. Travelling is “the process of going from one place to another in a vehicle”, visiting is “spending time in a place for a temporal period of time before moving on to another place”, and wandering is “extensive local mobility in a building or local area” [Ibid]. The shared symbolisation in these activities could be a common understanding of language.

The challenge posed by fluidity is heightened by the vision of pervasive and ubiquitous computing we subscribe to, where online and offline activities involving others is layered on top of the individual's activities. Being on call and collaborating with others constantly is technologically feasible, but humanly impossible. Here we present a technology that facilitates the user having an effective and productive Internet presence continuously and in changing use contexts involving others. The central theme in this approach is the outsourcing of some interaction and cognition. This involves handing over responsibility for particular tasks to software systems, namely the agents and web services we build into a Digital Self using ShadowBoard. We are now extending the framework of tools to cover mobile technology, enabling the mobile user to be engaged in fluid, cooperative interaction, while moving between different social worlds. In maintaining the naming theme, we are calling this expanded framework, ShadowPlaces. Interestingly, there is a one-to-one map between the hierarchy of an individual's *roles* as modelled in *ShadowBoard*, and the hierarchy of multiple *social worlds* as modelled by Strauss that the user moves between, in the course of their daily life. This conveniently makes the ShadowBoard agent architecture and methodology eminently suitable for building systems that cross into the mobile context, down at the theory level.

Section 2 of the paper is an overview of the *ShadowBoard architecture*, the blueprint of all that we have done, along with an introduction to the foundational psychological theory upon which it sits.

Section 3 presents the associated *ShadowBoard methodology* which allows developers/users themselves, to build agent-based applications realising a Digital Self for a given user.

Section 4 describes aspects of the software framework that has been developed, including the *ShadowSpaces* middleware, the *ShadowFaces user interface*, and a constraint logic programming language named *CoLoG*. - which together constitute an instantiation of the ShadowBoard architecture.

Section 5 puts the concepts together by sketching an example *Digital Self* that wraps web services to support the lab management of the Interaction Design Lab at the University of Melbourne.

In Section 6 we draw upon concepts from the well known *Locale Framework*, itself based on Social Worlds theory, to outline what we now need to do, to take the ShadowBoard system into mobile contexts. The resulting enhanced architecture *ShadowPlaces* is then depicted.

Section 7 uses the Locale Framework to *evaluate* ShadowPlaces as a viable multi-user environment suitable for groups of mobile users.

2. THEORY AND ARCHITECTURE FOR A COMLEX INDIVIDUAL AGENT

Figure 1 below is a graphic overview of the ShadowBoard architecture, collectively representing an individual *whole agent* made up of numerous sub-components – the structural implications are inferred from the theory based on the *Psychology of Subselves* psychology, covered extensively elsewhere [Goschnick 2001]. In the centre of the agent is the *Aware Ego Agent* – the dominant sub-agent in the whole cluster of sub-agents. In the figure, the Aware Ego Agent is surrounded by eight first-level sub-agents, diagrammatically drawn as circles the same size as the Aware Ego Agent. Five of these example sub-agents are not nested any deeper (i.e. sub-agents can be clustered recursively), while the other three have clusters of circles within them, representing a second-level of sub-agents, grouped into numerous *envelopes-of-capability*.

Each *envelope-of-capability*, of which there are arbitrarily eight in the figure, represent different areas of expertise that a particular whole agent embodies. The whole agent can perform a number of consecutive and diverse tasks, depending on what *goals* via what *roles* it has taken on in the outer world, on behalf of the user.

Each envelope-of-capability (EoC) contains a number of sub-agents with similar capabilities, but each different from it neighbour's abilities in some specialised way. At the two far ends of an EoC are two diametrically skilled agents - in Figure 1 the two are adjacent, one is white the other is black, separated by a dotted radial line. One of these two extremes is the *basic reactive sub-agent* a purely reactive agent [Brooks 1987], with a hard-wired, rule-action mechanism with no *deliberative* capability. The basic reactive sub-agent is usually called upon within a particular EoC when *time* or other *resources* are severely constrained, and a rule-of-thumb to deliver an action, is better than no action at all.

At the other end of the scale is an *archetypal sub-agent*, one that has maximum deliberative ability, used when time and other resources are plentiful allowing some *automated reasoning* to be incorporated.

The *Aware Ego Agent,* as well as each of those other sub-agents that are shaded as spheres in Figure 1 (*delegation sub-agents*), have *knowledge about the capabilities* of sub-agents in their respective EoCs. They use this knowledge to select the appropriate sub-agent to achieve the particular goal that has been sent their way, from higher up the recursive hierarchy.

When a sub-agent is found lacking in capability to achieve a specific *intention*, or when an external (and available) agent matches the particular specialty better than any internal sub-agents, an external agent can be called upon as if it were an internal sub-agent. This process is termed *disowning a sub-agent* and outsourcing via an external agency. In Section 4.4 we use this mechanism to *wrap external web services* as sub-agents.

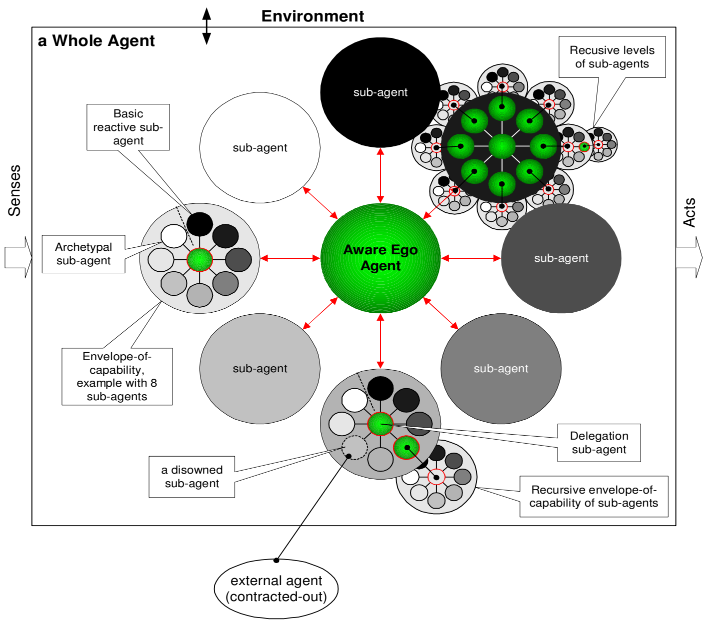

Figure 1. The ShadowBoard Architecture.

2.1. Sub-Agents as Subselves

The mentalistic notion of a *subself* at work within the psyche of an individual, is a metaphor for the *sub-agent* of a ShadowBoard agent. To broadly place this work in context of research upon multi-agent systems (MAS) – the closest comparative field of study - most multi-agent systems can be described as *inter*-agent systems. In contrast, the ShadowBoard theory and architecture is an *intra*-agent system, one enabling the incorporation of many components that together represent one *whole* agent, albeit a very sophisticated entity. Such a *whole agent* built upon the ShadowBoard architecture – e.g. the *Digital Self* – seen from without, should be seen as a fully *autonomous* individual agent, one compatible with contemporary definitions of agency, such as that of [Wooldridge 1995].

Unlike the autonomous whole agent, the inner sub-agents are semi-autonomous or even totally subservient to the *Aware Ego Agent* - the executive controller within a ShadowBoard agent. Sub-agents may themselves be sophisticated agents capable of their own semi-autonomous work, or they may be conventional application programs, expert systems, or even wrapped web services as we shall see further down. This is a significant relaxation on the need in most MAS systems for each individual agent to be a fully functioning autonomous agent. This flexibility in capability within a ShadowBoard agent, is aimed at augmenting human skills and attention ranges, rather than for building synthetic entities, and it is attained by making all sub-agents ultimately subordinate to the *Aware Ego Agent*. Our relaxation of autonomy is comparable at some level with the work of [Scerri 2002] on *adjustable autonomy* on their *electric elves* project.

ShadowBoard Sub-agent Types	
	(cont)
DecisionMaker	**Logician**
Manager	**Player**
Protector	**Teacher**
PersonalAssistant	**Mentor**
Adviser	**Engineer**
Critic	**Artist**
Initiator	**Intuitive**
Adventurer	**ContextSituator**
KnowledgeSeeker	**Intrapersonal**

Table 1. ShadowBoard Sub-agent Types.

The inclusion of sub-agents and envelopes-of-capability of sub-agents within the ShadowBoard architecture, allows it to be populated with all manner of domain-specific and types of sub-agents, allowing for an open-ended expansion of capability and knowledge, depending on an individuals needs and the groups they interact with, in the social worlds they participate in.

3. METHODOLOGY

The aim of building a Digital Self, involves building a comprehensive *user-proxy* - something that not only augments the user's ability and attention span but which also represents the user at times when they are *off-line* – technically or otherwise. In order to do so for the range and diversity of individuals that we are aiming to cater for – the mass-market user - we have formulated a methodology. It begins with the definition of a generic range of sub-agent *types* which are most likely to encapsulate the sorts of envelopes-of-capability any given user may want to build into their Digital Self. These types are a starting point only, used for individuals to articulate and accumulate their necessary sub-agents. This list of ShadowBoard sub-agent *types* appears in Table 1 above.

Most often these high level sub-agent types will represent an EoC, which in turn have a range of like-functioned sub-agents that cover a range of capabilities and degrees of deliberation. Example generic names for sub-agents within two of the types from Table 1, appear in Table 2 below. There are extended generic names for other sub-agents types that similarly drill-down the other Envelope-of-capability types in Table 1, fully detailed in [Goschnick 2001].

It is unlikely that a given individual will have sub-agents of all these types in their Digital Self. This is the full generic range, the purpose of which is to help *any* user identify, collect and organise a combination of sub-agents into a common grouping, about which they can discuss issues with other users/developers, or obtain guidance from further automated processes, as they evolve. Section 5 outlines an example Digital Self, using this ShadowBoard *methodology* to identify candidate sub-agencies.

Example 1	**Example 2**
PROTECTOR **(envelope-of-capability)** **Safety Officer *(archetypal)*** **Defender** **Risk Analyst** **Environmentalist** **Pacifier** **Doctor** **Exit Strategist *(reactive)***	**MANAGER** **(envelope-of-capability)** **Benevolent *(archetypal)*** **Conciliatory Manager** **Planner** **Scheduler** **Coordinator** **Recycler** **Decisive (*reactive*)**

Table 2. ShadowBoard Sub-agent Types.

Furthermore, it should be noted that other envelopes-of-capability outside this starting range, drawing upon agent ontologies compiled by other people, can occupy any level within the recursive hierarchy of ShadowBoard. For example [Hristozova 2002] specified a range of Middle Agents - *Matchmaker Agent, Recommender Agent, Mediator Agent, Facilitator agent, Broker agent* – which establish, maintain and complete communications with end agents, in a graduated range of capability. In other words, the methodology for building a Digital Self, is very much about utilising third party agents and ranges of agents, as sub-agents within the ShadowBoard architecture framework.

4. IMPLEMENTATION OF SHADOWBOARD FRAMEOWRK

Enacting the ShadowBoard architecture has been hastened by the use of several existing technologies at our disposal, which we have modified as necessary, and added to with new technology resulting in an integrated, cohesive system.

There are several levels of technology. At the base is a system called *ShadowSpaces* an enhancement of earlier technology which is a dynamic model of the hierarchy of sub-agents which are later intertwined with logic programming, to enact a Digital Self. There is an interface system we have named *ShadowFaces* (Section 4.2) which the user uses as both: an editor to *create* and *modify* their Digital Self. ShadowFaces remains running at all times, as the face of the Digital Self, receiving notifications, filtered information and details of decisions that the Digital Self has made, or referring decisions back to the user for their executive decision when necessary. It displays visual alerts for the human user's attention.

There is an underlying language parser and computation system, which uses constraint logic programming [Marriott 1998] to bind together the various sub-agents into a more powerful configuration. We have named the language *CoLoG*. For further details see [Goschnick 2003b].

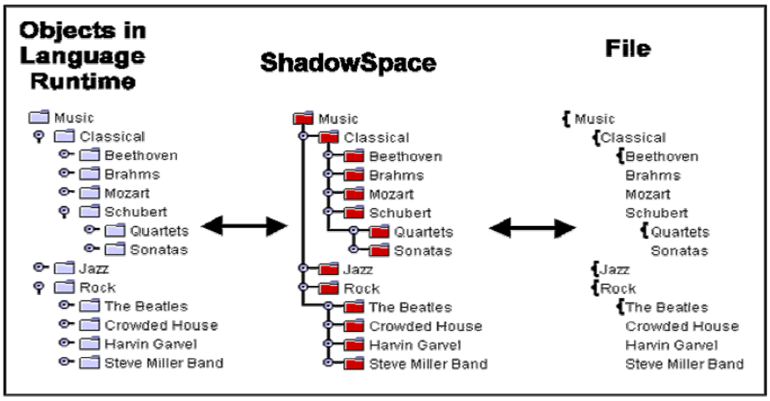

Figure 2. Context of a ShadowSpace.

4.1. ShadowSpaces: a Dynamic Object System

ShadowSpaces is not unlike the W3C DOM (Document Object Model), and with some limited effort could be made to be DOM compliant [WWW 2001]. It is used to dynamically maintain a hierarchy of objects, independent of the way they are instantiated and stored in the language the systems is created in, i.e. Java. See Figure 2 above.

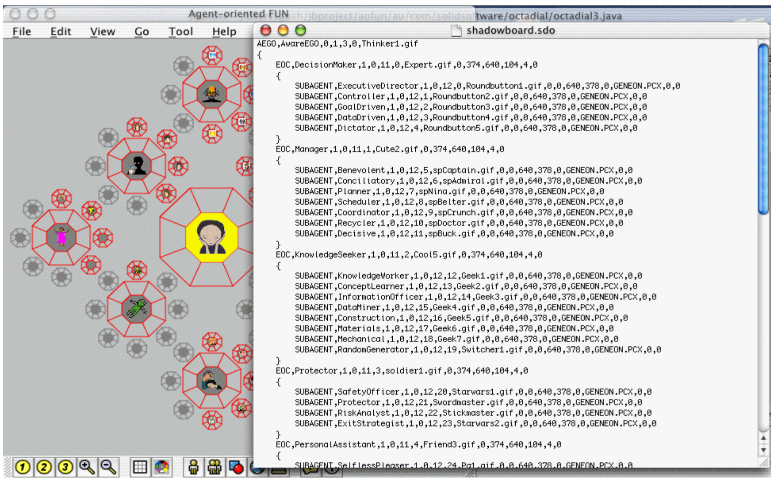

Figure 3. The ShadowSpace representing a Digital Self.

The ShadowSpaces is initialised via a human and computer readable file as depicted on the right-hand-side of both Figures 2 and 3. On the left-hand-side of Figure 2, is the equivalent runtime object hierarchy within the Java language. The hierarchy in the middle of Figure 2 represents the independent ShadowSpace hierarchy of tuples, which are a light-weight fingerprint of the objects as implemented in Java. As one small example of its usefulness, if for some reason the instantiated objects get dropped from memory, they are easily reinstantiated via the fingerprint in the ShadowSpace. As with the W3C DOM model, the ShadowSpace tree is *mutable* - able to deal with events that modify the structure of the tree.

The human-readable configuration file of a ShadowSpace for a given Digital Self, is initially created using the ShadowFace interface in *editor mode*. Afterwards, when the system is running as a continual 24x7 process, the ShadowSpace represents a snapshot of the *system state*. The ShadowSpace is also visually rendered in the interface (as per left-hand-side of Figure 3), both as part of the editor, and for the purpose of alerting the user about some change in the Digital Self, which the user has indicated an interest in, via earlier configuration.

4.2. ShadowFaces

The interface we have built adopts an agent-metaphor (in addition to it being an agent technology) rather than the desktop metaphor – which we consider to be a severely challenged metaphor in the 24x7 space. It is worth noting that throughout the era of GUI document-centric operating systems (MS Windows, Mac OS), in common applications such as the humble word processor, there has been some functionality that has followed an agent-metaphor. In products such as Microsoft Word and Wordperfect, the *find/search* and *replace* functions, follow an agent metaphor style of operation – as opposed to the direct-manipulation operation of say, the *marking* and *bolding* of text in a document [Lewis 1998].

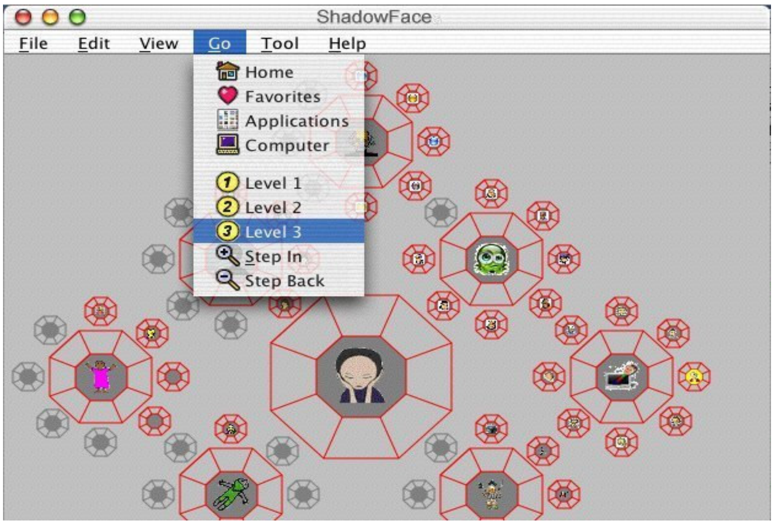

Figure 4. ShadowFace the ShadowBoard Interface.

ShadowFaces - the ShadowBoard interface - is an enhanced version of our earlier technology called FUN – Friendly User Navigation. It displays three levels (four, counting the leaves) of the hierarchy of sub-agents at any one time, and can be envisaged as a *information lens* that the user can use to navigate through the recursive structure of the ShadowBoard architecture. See Figure 4 above.

The advantages of this interface are numerous:

It is based on an agent-metaphor rather than the document-centric interface of the desktop metaphor of most current day operating systems and many of the applications that run upon them.

Up to 589 individual items in the hierarchy can be displayed and are selectable at the one time, versus the use of a tree manager approach (such as Microsoft Windows Explorer, or the Finder in Mac OSX), which display no more than about 50 items in a typical display.

The shape echoes the graphic depiction of the ShadowBoard architecture, reinforcing the *system image* in the user's mind – useful when introducing a new software interface [Shneiderman 1997].

The interface can be expanded to full-screen or reduced down to the size of a handheld screen (as low as 160x160 pixels), making it an ideal interface for a number of different consumer devices.

It has an eight-way navigational action making it equally suitable to: the mouse; the keypad, function-key pads, and game console joystick interfaces – again making it ideal for 24x7 operation via suitability to different types of consumer devices, each applicable to different locales.

The same interface can be used to: zoom, navigate, filter and get details-on-demand (see Figure 5 below, editing an Envelope-of-capability).

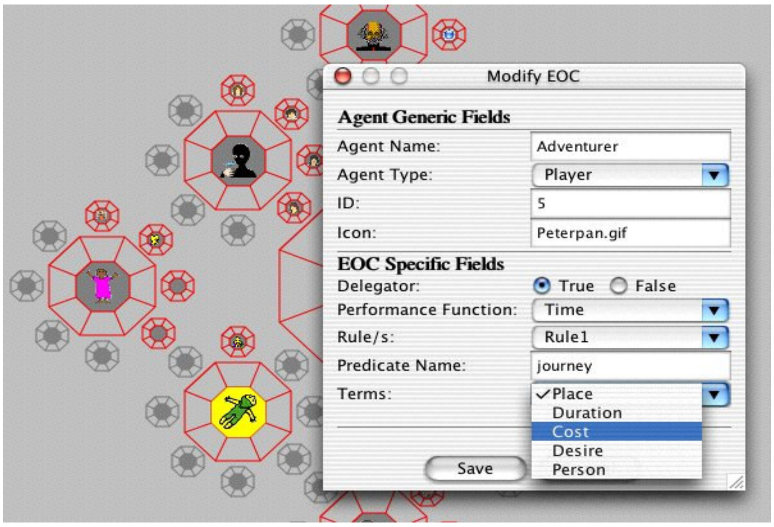

Figure 5. Editing the current *Envelope of Capability*.

There are a number of issues with regard to usability of this novel interface, which are currently the subject of an ongoing HCI research project, to make it increasingly usable to the uninitiated.

4.3. The Computational Engine: CoLoG

The computation engine is based on a *distributed constraint logic* program called *CoLoG* that we have developed as a part of the ShadowBoard system. We have derived CoLoG from a Java implementation of a single-threaded logic parser, added some constraint commands, and then have embedded it into the ShadowSpaces hierarchy. Each Envelope-of-Capability node within ShadowBoard is capable of its own thread of execution, and the primary execution in such a node is a CoLoG program. This aspect of the implementation is covered elsewhere [Goschnick 2003a].

5. AN EXAMPLE DIGITAL SELF

For the first useful implementation of a Digital Self, we chose to analyse, specify and then implement one representing some of the *roles* and *tasks* undertaken by the manager of the IDEA Lab [IDEA Lab 2004], a usability lab at the University of Melbourne. This illustrates how the ShadowBoard system can integrate otherwise disparate web services into a collaborative, dynamic system. To understand the process, it is first necessary to look a little into the day-to-day operations of the IDEA Lab.

The IDEA Lab is a usability laboratory in which people (researchers, usability engineers, web designers, students, etc) test software, web sites, mobile devices and other devices and procedures, from the point of view of their *user functionality*. For example, a *web-site usability test* involves booking two rooms, separated by a one-way mirror, with hardware and people on both sides, performing different but related functions. In the 'test room' is the user ('the subject'), working upon a computer connected to the web-site being evaluated. The subject is trying to perform a number of tasks, which the web-site has been designed and built to do. Focused upon the subject, are a number of video cameras and microphones and a screen video capture card. Behind the mirror in the 'observation room' there is typically a usability engineer, and/or members of the web-site design team. They are monitoring all actions, and the whole session is being recorded to tape or hard-disk, i.e. the video (multiple cameras) and sound. Back in the test room sitting beside the user, is a facilitator, a person who keeps the process flowing smoothly without adversely affecting the test results.

The IDEA Lab is a large facility with four possible combinations of test room-plus-observation room, and a multitude of video and sound recording devices. A typical usability test involves the sourcing of between 5 and 15 representative users of the web-site currently being tested, who usually have no previous experience of the IDEA Lab or its immediate locality. The IDEA Lab hires the facilities to a number of organisations, including other university departments and industry.

Managing the whole process of bookings, and orchestrating people and processes is multi-variant and lends itself well to being augmented with a Digital Self (representing the IDEA Lab Manager) loaded with an orchestration of web services and other software agents as necessary.

The tasks and processes in need of agency as listed below, were teased out using the ShadowBoard methodology. Where a particular agency represents a ShadowBoard sub-agent type, the type name is added in square brackets, e.g. [**Adviser**]. A number of the web services that enact several of the listed agencies, are currently available in the OBY1 (Online Booking System V1) booking system developed for the IDEA Lab [Jin 2003]. In the list below, such existing web services are noted with '(OBY1)' after them.

These tasks and processes are best broken down into the following temporal divisions:

- Before Booking
- During Booking
- After Booking

However, several tasks and processes from the following list, may be appropriate across two or more of those temporal divisions.

Before lab booking:

- Choices of rooms, equipment and experts, many possible combinations (OBY1).

- Cost comparisons of IDEA Lab options (OBY1).
- Cost comparisons with other usability labs.
- See an appropriate usability test demo video using the combination of equipment of interest. [**Server/Marketer**]
- Retrieve profile information on the experts being considered [**KnowledgeSeeker/InformationAgent**].

During lab booking:

- Coordinate possible schedules of main players (facilitator, usability engineer/s, designer/s) [**Manager/Planner**]

After lab booking:

- Source a number of appropriate test subjects (online bidding process) [**Server/Networker**]
- Advise public transport options to test subjects and others. Many test subjects are university students and use public transport. [**Adventurer/Traveller**]
- Car parking locations and related advice (street parking is at a premium nearby) – location, time limits, costs. [**Adviser/ParkingDomain**]
- Food outlet locations available for the test subject and others to source meals. [**Advisor/Café-Restaurant**]
- Nearest ATM/bank terminals. [**Adviser/BankingDomain**]
- Online manuals, interactive demonstration and/or self-paced lesson on the use of selected equipment. [**Teacher**]
- Suggested sources for appropriate media types – DV digital video tapes, CDR, etc – the IDEA Lab is not a media store. [**Adviser/MediaDomain**]
- Advice on how many tapes and other media will be needed for the booked sessions. [**Adviser/Materials**]
- Advice on professional photographers – many lab users use some content from their testing procedures in the promotion of the resultant product or service. [**Artist/Photographer**]
- Advice and booking of onsite mobile caterers. The IDEA Lab includes one lab large enough to run 30 people seminars on usability and related topics. [**Adviser/Catering**]
- Access a tutors database. The large lab can be booked for tutorials and labs which involve a number of supervising course tutors. [**Advisor/Tutors**]
- Onsite phone numbers and security advice. Because of the need to coordinate a number of people from different working backgrounds, usability sessions are often run outside of usual business hours. [**Protector/SafetyOfficer**]
- Video editing advice. The IDEA Lab includes video editing equipment, most often used to produce a 'highlights tape' – usually a 5 to 10 minute polished overview, of what was discovered/encountered during the many hours of testing. The final distribution of the 'highlights tape' can be by tape, CD, Internet, or other media. Selection of appropriate video file formats, resolutions and storage sizes, is fertile ground for a video editing agent/expert system. [**Adviser/VideoFormats]**
- Generate an operational plan detailing sequencing of people, equipment, timing and location – covering the whole usability test. [**Manager/Coordinator**]

Many of these individual agencies are provided via individual web services, and others are still being built as such. However, the orchestration of these services using the ShadowBoard system, brings new synergies of service to the user. Once the time schedule of the test subjects is know, the various adviser agents can customise their advice accordingly. e.g. The public transport advisor can then be automatically refined to just those services which are relevant to the specific test subject, booked-in for a specific time period. Likewise for car parking options. If the person needs to get money from an ATM machine before catching public transport, the direction of the nearest ATM can be taken into account when suggesting a tram stop. The available food outlets can be customised for both time and for the test subjects preferred type of outlet.

As these auxiliary web services come online and are orchestrated with the existing services via the ShadowBoard system, there are transactional issues to be addressed in order to maintain the functional integrity of the operational system.

These transactional issues are explored in research discussed in [Jin 2003].

6. SHADOWPLACES: MOVING BETWEEN SOCIAL WORLDS

Taking the extensive work we have done on sophisticated individual agents based on the Psychology of Subselves, we now move forward to insert these agencies, into a technological framework appropriately called *ShadowPlaces* - a framework designed to allow multi-ShadowBoard users to participate in networked groups which reflect some of the social worlds they participate in.

The Locale Framework developed by [Fitzpatrick 1996, 2003] was put forward as a way to identify and articulate the '*sites and means*' of a *Social World* [Strauss 1976] in the context of CSCW (computer supported cooperative work) systems. Where a *locale* is defined as a *place* rather than just a space, which includes both the space and the resources (hence the term: *site and means*) used by a group of people or a social world, meeting their needs in pursuing some common purpose. Note: A *social world* as defined by Strauss is a group of individuals (or groups) bonded by a common goal.

Fitzpatrick believed that the Locale Framework could play a communicative role in understanding key activities in CSCW, via the use of the set of abstract language terms that it defines – an understanding between the technical and social areas of research that were feeding into CSCW. She used and recommends the Locale Frameworks as a significant contribution in nurturing understanding between the technical and social, particularly with respect to CSCW.

This framework is precisely what we needed to guide us forward. While we are not just interested in social worlds that necessarily support *work*, but also on social groups forming and splintering through the disruption to communications and a changing context as the mobile user moves through real space and time, and other changing contextual parameters.

We are intent here to use the Locale Framework to improve both: our understanding of applying the albeit sophisticated single-user agent ShadowBoard architecture to a *multi-user* situation, and to promote the understanding of the ShadowBoard agent architecture to a wider audience. It is worth noting here, that although the ShadowBoard agent architecture was originally envisaged as a redefinition of a fine-grained individual software agent based on analytical psychology, it was always envisaged [Goschnick 2000] that the newly sophisticated agency built upon it, would then plug back into a multi-agent system (MAS). What we have already gained from the Locales Framework adaption of social world concepts, is that rather than model it on a MAS of software agents, we are developing a system that incorporates the user in *multiple social worlds* using their digital user-proxy, what we have termed the Digital Self.

Given that the ShadowBoard agent architecture reflects the Psychology of Subselves, which in turn reflections the complex roles and multi-facited identity of a typical, busy, modern individual, it ought not be surprising that it complements very well, Stauss' original multiple Social World view of an individual's multitude of actions. i.e. Social Worlds may contain sub-worlds, and those sub-worlds may contain further sub-worlds and so forth. Correspondingly, ShadowBoard can accommodate sub-selves (supported by sub-agents) and those sub-selves may contain further sub-selves, as apparent for the Individual performing Role C in Social World 1, in Figure 6 below. Very often the subselves equate to roles and sub-roles that an individual is performing, within the realms of work, leisure and family activities – i.e. their social worlds.

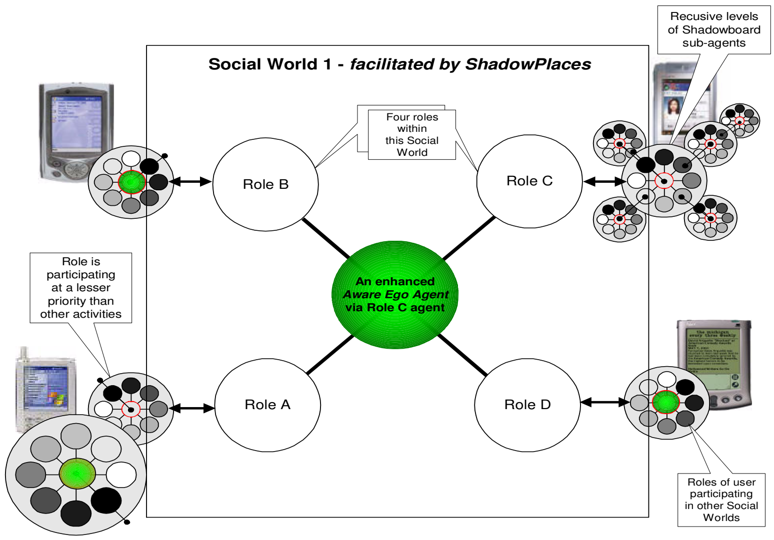

Figure 6. ShadowPlaces Architecture supporting a Social World using networked ShadowBoard agent enabled devices.

Our need to spontaneously support a network of individuals via their handheld PDAs and mobile phones – each with ShadowBoard agency built into them - is technically achievable with either Bluetooth and/or WiFi technologies. However, the social interactions that it is built to support, are more clearly defined and understood in the language of social worlds, particularly as articulated in the abstract language of Fitzpatrick's *Locale Framework*.

In Figure 6 above there are four people participating in a social world. The roles in ShadowBoard facilitate membership in a social world. All four of these participating individuals have a multitude of roles in their lives, but just one of those roles is currently included in Social World 1. They may be participating in other social worlds via other roles in those social worlds, either simultaneously or sequentially.

The sub-agents built into their ShadowBoard systems help them participate in *multiple* social worlds simultaneous, by dealing with some of the less-urgent requests and actions directed at, and expected of, the individual. In addition, it helps them to participate in social worlds in 24x7 time – monitoring, registering, queuing, notifying, and alerting the individual about appropriate events (actions taken, required actions) in their social worlds, at appropriate times, as deemed by their individual configurations of the Digital Self. As such the use of the agent system *expands their range* of participation in the social world to 24x7 time.

The individual who is performing Role A in Social World 1 above, is doing so at a sub-role level, down the hierarchy of sub-agents, at a lower priority of interaction than the other three individuals. This too, matches observation in the social worlds theory and which the Locale Framework alludes to. Fitzpatrick observes that "*People engage in a range of work activities with different social worlds and hence different locale views. … They move seemlessly and often unconsciously between activities, maintaining dynamically varying levels of focus and participation in the different locales, from full intense focus, as if being in only one locale, to having background awareness, to being 'out of sight'* ".

7. MEASURING SHADOWPLACES WITH THE LOCALE FRAMEWORK

In addition to aiding the understanding of complex systems across the technical and social divide, [Fitzpatrick 1996] also recommended the Locale Framework as a general model against which existing systems can be interpreted as contributing to one or more aspects of the framework. i.e. How well they *measure up*, is a direct measure of how well they support CSCW. The following five points paraphrase the fuller definitions of the five aspects of their Locales Framework:

1. *Locale foundations* define the basic locale structures.
2. *Mutuality* describes the way in which interactions between members are supported.
3. *Individual views* describe the way in which individuals construct personalised views of the multiple social worlds they participate in.
4. *Interaction trajectories* describe the temporal dimensions of interaction.
5. *Civic structures* define the relationship of locales into public spheres of interaction.

While our interests go beyond just *work*, work-oriented social worlds nonetheless are envisaged as a large generator of the interactions in the mobile-device-plus-Internet system we are building – as portrayed with the Lab Manager digital self in Section 5 above.

7.1. *Locale Foundations* in ShadowPlaces

In this part of the Locales Framework, the group, or the social world of interest is investigated and described in terms of the '*spaces, objects, tools, media and other resources used to support their interactions*' [Fitzpatrick 2003]. In short, the '*site and means*' that become the locale, ought to be described and discussed.

Our group of people of interest is any social world in which there are individuals with mobile devices in the form of PDAs or mobile phones that are also PDAs, that support either Bluetooth [Bluetooth 2003] or WiFi [Wi-Fi Forum 2003] wireless networking.

Strauss defined the limits of communication as representing the boundary of the social world – Bluetooth has a very limited range, while WiFi too has a limited range, but if it is within proximity of a base-station it will connect the user to the Internet thereby giving them global range. WiFi is now seen as the pre-eminent technology to help cover '*the last mile*' in connecting every one, every where to broadband Internet.

Because we are designing for users who will fluctuate in and out of range of base-stations, we need to consider them still within the group, even when they are out-of-range. The situation when users know they are out of range but are consciously willing to be active, is covered further down in the next section on *Mutuality*.

Incorporating software agency into this system, effectively extends the boundary of the social group, beyond what would normally be considered a limit of communication. Asynchronous communications such as store-and-forward email similarly extends the boundary of communication in a similar but less dynamic manner.

The particular software agency we use (ShadowBoard agents) also highlights a difference between 'site and means' that the agent system will use, compared with that the human user will use directly. In the case of accessing Internet-based sources of information, the human user uses a web browser in the PDA while the sub-agents employed within their Digital Self will use *web services* – a programmatic interface to the Internet – to source the same or similar information.

In the case of a network with Bluetooth-only devices, one of the devices has to take on a master/server role, while the other Bluetooth devices are slave/clients – this is a function of spontaneous Bluetooth networking. In Figure 6 above, the device enabling Role C has taken on this server role, by effectively facilitating the ShadowPlace which supports the Social World 1. i.e. The shaded sphere in the very centre of the figure represents the Aware Ego Agent of Role C, but one that is enhanced with the extra functionality to facilitate the Bluetooth connections, and other ShadowPlace functionality discussed further down. If the four people representing Roles A, B, C and D remain in proximity, all is well. If A, B or D drop out of range, the others may continue to interact with one another. However, if C drops out of range of the rest, they will all lose their connections and another master/server has to be instantiated drawn from within those group members remaining in range.

In a WiFi network the Role C agent facilitating the social world, would best be located on a server on the Internet, rather than on a PDA, negating the need to swap master/server roles amongst participants of the social world. However, if a social world is formed randomly, faciliated by people in a common location – such as customers at a restaurant or café – and is short-lived, then the need for the ShadowPlaces facilitating agent to be on an Internet server, might not be necessary nor desirable.

The negotiation of joining and leaving a social space, and the level of participation in it, are aspects highlighted in the Locale Framework that we will need to focus considerable attention on in the ShadowPlaces enhancements.

7.2. *Mutuality* in ShadowPlaces

Mutuality deals with the *who, what, when, where* and *how* of interactivity. The presence-awareness dialectic refers to: the individual's presence being made evident to the group (usually via a visual icon in an interface screen); and the individual's awareness that their presence is being presented to the group. As with most CSCW systems, our endeavours with ShadowPlaces includes the need of visual icons that indicate the presence of individual members of the social world.

This functionality is over-and-above that currently in ShadowFaces and so it is currently being enhanced in the ShadowPlaces version of the interface. The 24x7 nature of the system means that all members will be *present,* though some at a *reduced state of attentiveness* – including asleep. The capabilities and choice options of such individuals are dependent on the capabilities in their Digital Selves, and on the levels of autonomy they are allowed to take. Monitoring, alerting and notifying the individual appropriately, is an area of personal assistant agent research currently receiving a lot of attention.

7.3. *Individual Views* in ShadowPlaces

Individual views in the Locale Framework, is the view or perspective of an individual across multiple locales based upon their membership of multiple social worlds. Certainly, the ShadowBoard agent system supports *Individual Views* well beyond an individual's representation in a single social world in which they are participating. The recursive structure of sub-selves in the underlying ShadowBoard architecture, echos the recursive structure of social worlds and sub-worlds – after all, subselves in the Psychology drawn upon, equate most often to roles in a life, the same roles upon which social worlds aggregate their membership. Just as the membership of social worlds can be drawn from formal persistent roles through to highly informal and transient roles, so too the subselves in the underlying psychology of ShadowBoard.

Where the social world acts as a centre to attract individuals and resources to achieve its collective goals, the Digital Self built upon a ShadowBoard architecture is designed to enable the individual to be a centre, one that participates in multiple social worlds. This degree of refinement in a ShadowBoard agent is what sets it apart from other currently available agent architectures, and is why it is such a good candidate to build sophisticated agent-enabled interactive interfaces.

7.4. *Interaction Trajectories* in ShadowPlaces

Members of a social world perform actions and interact to accomplish collective goals. The *interactants* in social worlds are usually people alone. In systems such as ShadowPlaces – with *agents-in-the-loop* – many of the interactants are software agents. Whether the actions within the system are performed by humans or by member sub-agents, the underlying ShadowBoard technology is capable of recording past actions from the *action stack*. This is analogous to the *trajectory phase* of the Locale Framework – it is a related sequence of interactions occurring over time.

The Locale Framework's *Trajectory projection* is the vision of the expected course of action. Agent systems describe these as *goals* which are enacted via *intensions*.

The *trajectory scheme* in the Locale Framework is the consciously developed plan given the projection. Software agents have an equivalent in the selected *Plan* as they rationally deliberate (via the CoLoG language in our case).

8. CONCLUSION

The ShadowBoard theory, methodology, and the instantiated framework which includes the ShadowFaces interface are well advanced, are currently usable for building prototypes - as we did for the Lab Manager Digital Self in Section 5. These prototypes are proving very useful in building diverse and generic systems. However, we are finding that the web services that are currently available to the general Internet population (as candidate sub-agents), are extremely limited in both number, and in imaginative uses of knowledge in most domain areas. This problem will hopefully subside as web services continue to gain in popularity, as a way to provide distributed functionality via programmatic interfaces.

We are only now beginning to extend the system functionality (ShadowPlaces) into the realm spontaneous networks as needed for our ambitions in mobile applications, an area where the approach holds significant promise.

Prior to our investigation into the Locale Framework, multi-agent systems (MAS) looked like the appropriate direction for us to take. MAS have been successfully applied to robot soccer teams and to numerous eCommerce applications. However, our goal of building an individual agent with significant complexity is aimed at *augmenting* a human user in 24x7, not at building a fully autonomous synthetic agent. The resultant software agency in a ShadowBoard agent is one that very nicely fits into social worlds occupied by humans as described by Strauss, but one supplemented with technological wherewithal. As noted in Section 7.2 above, the *Psychology of Subselves* behind ShadowBoard and the sociology behind *Shadow Worlds*, are two sides of the same coin - people take on roles, formal or otherwise, as they negotiate the fluidity of multiple social worlds.

The ShadowBoard technology investment has given us the appropriate technology, while the Locale Framework has given us an appropriate tool to guide and extract the requirements for the functionality we need and are now placing in the ShadowPlaces system, to make it a *Locale* for mobile and Internet users, in 24x7 time.

9. ACKNOWLEDGEMENTS

This research is now been applied in the ongoing building of *The Digital Self Project*, as an open source project, funded by a Telstra Broadband Fund Development Grant, from Telstra Ltd, Australia.

REFERENCES

[Wooldridge 1995] Wooldridge, M. and Jennings, N. Intelligent Agents: Theory and Practice. *Knowledge Engineering Review*, 10(2), 1995, pp 115-152.

[Goschnick 2000] Goschnick S.B. ShadowBoard: A Whole-Agent Architecture that draws Abstractions from Analytical Psychology. In *Proceedings of PRIMA 2000*, Melbourne, Australia, August, 2000.

[Goschnick 2001] Goschnick, S.B. *ShadowBoard: an Agent Architecture for enabling a sophisticated Digital Self*. Thesis, Dept. of Computer Science, University of Melbourne, 2001.

[Goschnick 2003a] Goschnick, S.B. Enacting and an Agent-based Digital Self in a 24x7 Web Services World. In *Proceeding ISMIS-2003 the 14th International Symposium on Methodologies for Intelligent Systems*, Maebashi, Japan, 2003.

[Goschnick 2003b] Goschnick, S.B. and Sterling, L. Enacting and Interacting with an Agent-based Digital Self in a 24x7 Web Services World. In *Proceedings, IEEE joint conference on Web Intelligence and Intelligent Agent Technology (WI/IAT)*, Halifax, Canada, 2003.

[Graham 2003] Graham C. and Kjeldskov J. Indexical Representations for Context-Aware Mobile Devices. In *Proceedings of The IADIS e-Society Conference 2003*, Portugal, IADIS Press, 2003, pp 373-380.

[Kakihara 2002] Kakihara, M. & Sørensen, C. Mobility: An Extended Perspective. In *Proceedings of the 35th Hawai International Conference on System Sciences (HICSS-35).* Hawaii, USA, IEEE, 2002.

[Mol 1994] Mol, A. & Law, J. Regions, Networks and Fluids: Anaemia and Social Topology. *Social Studies of Science*, *24*(4), 1994, pp 641-671.

[Strauss 1976] Strauss, A. Life Styles and Urban Space. In Proshansky, H.M., Ittelson, W.H. & Rivlin, L.G. (Eds), *Environmental Psychology: People and Their Physical Settings*, New York, USA: Holt, Rinehart & Winston, 1976, pp 528-537.

[Strauss 1993] Strauss, A. *Continual Permutations of Action.* New York, USA: Aldine De Gruyter, 1993.

[Kristoffersen 1999] Kristoffersen, S. and Ljungberg, F. Mobile Use of IT. In *Proceedings of the 22nd Information Systems Research Seminar in Scandinavia (IRIS22)*, Jyväskylä, Finland: Jyväskylä University Printing House, 1999, pp 271-284.

[Brooks 1987] Brooks, R.A. Intelligence Without Representation. In Workshop in Foundations of Artificial Intelligence, Endicott House, Dedham, Massachusetts, USA, June, 1987.

[Scerri 2002] Scerri, P., Pynadath, D., and Tambe, M. Why the Elf Acted Autonomously: Towards a Theory of Adjustable Autonomy. In International Conference on Autonomous Agents and Multi-Agent Systems, 2002.

[Hristozova 2002] Hristozova, M. Middle Agents. Presentation, Intelligent Agent Lab, Department of Computer Science, University of Melbourne, 2002.

[Marriott 1998] Marriott, K. and Stuckey, P.J. *Programming with Constraints: an Introduction*, MIT Press, 1998.

[WWW 2001] World Wide Web Consortium Document Object Model (DOM), 2001. Documentation available at <http://www.w3.org/DOM/>

[Lewis 1998] Lewis, M. Designing for Human-Agent Interaction. *AI Magazine*, Summer 1998, pp 67-78.

[Shneiderman 1997] Shneiderman, B. *Designing the User Interface, Strategies for Effective Human-Computer Interaction*, third edition, Addison-Wesley, 1997.

[Jin 2003] Jin, T. & Goschnick, S.B. An Agent-based Transaction Model in Web Services. In *Proceedings Web-services and agent-based engineering, workshop at AAMAS*, 2003.

[IDEA Lab 2004] Interaction Design Lab, University of Melbourne, 2004. Documentation available at <http://www.idealab.dis.unimelb.edu.au>

[Fitzpatrick 1996] Fitzpatrick G., Mansfield T. & Kaplan S. Locales Framework: Exploring Foundations for Collaboration Support. In *Proceedings of the OzCHI'96: Sixth Australian Conference on Computer Human Interaction*, New Zealand: IEEE CS Press, pp 34-41.

[Fitzpatrick 2003] Fitzpatrick, G. *The Locale Framework: Understanding and Designing for Wicked Problems*. London: Kluwer Academic Publications.

[Bluetooth 2003] Bluetooth - the Official Bluetooth Website, 2003. Documentation available at <*http://www.bluetooth.com*>

[Wi-Fi Forum 2003] Wi-Fi Technology Forum, 2003. Documentation available at <*http://www.wi-fitechnology.com/*>