

Hybrid AR User Interfaces in Collaborative Gaming

Adrian Clark

November 15, 2004

Declaration

This report documents the research and work which I performed as part of the ARTankwars system. The system was an application developed for the SmartSpace project proposal which was written by Dr. Mark Billingham of the Human Interface Technology Laboratory NZ (HITLab NZ).

The ARTankwars development team consisted of three people including myself, Trond Nilsen, Julian Looser. Accordingly the system was divided in to separate components. Trond was responsible for the project management, and the design and development of the ARTankwars server and the game architecture. Julian was responsible for the development of the Augmented Reality client and the collaborative user support. As the title of this research states, my responsibilities in this project involved the design and implementation of the hybrid user interfaces, in particular the Tablet PC and Projector Interfaces. The contents of this report, including the literature review, observations and conclusions are all my own work.

Acknowledgements

I would like to thank Trond Nilsen and Julian Looser for their work on the ARTankwars project, and their support throughout the year. In particular Trond has been invaluable for his help keeping the project running smoothly and always providing me with suggestions and assistance when I needed help. I would like to thank Dr Mark Billingham for his assistance getting the project up and running and helping me to get set up in the HITLab. I would like to thank Dr Mukundan for supervising my project, organising meetings to make sure the project was on track, and offering suggestions for my progress reports and seminar. I would like to extend my thanks to Daniel Wissa also, who was kind enough to take the time to proofread this report and offer some feedback about it. I would also like to thank Dr Andy Cockburn and Mr Warwick Irwan for examining my report. Finally I would like to thank all the Postgrad Computer Science students, my Computer Science lecturers, friends and family for putting up with me in what has been a mentally and emotionally challenging year.

Contents

1	Abstract	3
2	Introduction	4
3	Background	5
3.1	Related Work	5
4	Technical Details	8
4.1	The Public Space	9
4.2	The Group Space	10
4.2.1	The Augmented Reality Interface	11
4.2.2	The Tablet PC Interface	13
4.3	The Private Space	15
4.3.1	The Private Augmented Reality Window	15
4.3.2	Transitional Interfaces	16
5	Results	18
5.1	Our Interaction with the System	18
5.2	Novice Users Interaction with the System	19
6	Discussion	20
6.1	The ability to represent multi-attribute data	20
6.2	Support for multiple views into the same dataset (ego and exocentric)	20
6.3	The use of real tools and natural interaction metaphors	20
6.4	Support for a shared workspace	21
7	Further Work	22
7.1	Markerless Tracking	22
7.2	SmartSpace Framework	22
7.3	Intelligent Projector Display	22
7.4	Outdoor AR Interface	23
7.5	Empirical evaluation	23
7.6	Rewrite Tablet PC Client	23
8	Conclusion	24

1 Abstract

This report documents the design and implementation of a number of hybrid user interfaces for the SmartSpace project. The SmartSpace project's aim was to explore "how hybrid user interfaces can be developed that employ multi-modal input and display techniques to enable users to visualize and interact with geospatial content"(Billingham 2003). A proposed application of such a project was for a Command and Control application, for example a military operation where commanding officers need to visualize real-time data from soldiers in the field. The system designed for this project was based on a simulation of a military scenario, where teams of two or more players compete against each other by controlling a number of tanks, which they can order to capture objectives or attack the other team's tanks. Because of the three dimensional characteristic of the environment in which this simulation takes place, an Augmented Reality interface was developed to allow users to view the data presented. Due to the non-tangible nature of Augmented Reality interfaces, a number of hybrid interfaces had to be developed to allow users to interact with the system.

2 Introduction

This research was part of the SmartSpace project, originally proposed by Mark Billinghurst of the Human Interface Technology Laboratory, New Zealand (HITLab NZ). The SmartSpace projects main focus was to “explore how hybrid user interfaces can be developed that employ multi-modal input and display techniques to enable users to visualize and interact with geospatial content”(Billinghurst 2003). The project uses multiple display and input interfaces to allow support for individual and group work. By combining the strengths of these various display and input interfaces, a seamless workspace could be created which supported three “Interaction Spaces”, as shown in Table 1.

Table 1: SmartSpace Display and Input Modalities (Billinghurst 2003)

Interaction Space	Display	Multimodal Input
<i>Private</i>	Tablet PC, Handheld display	Speech + pen, gaze/head tracking
<i>Group</i>	Tangible Map collaborative AR interface	Natural gesture + speech, tangible objects
<i>Public</i>	Wall Projection Display	Speech + laser point. PDA Control

A proposed application for SmartSpace was for a command and control scenario, for example a group of commanding officers co-ordinating a search and rescue or military team. Because of the three dimensional geospatial nature of the environment that the teams are in, a standard two dimensional interface would prove inadequate. SmartSpace uses the unique capabilities of different devices to seamlessly blend real and digital information spaces. In a situation where the safety of others is reliant on users of our system, we want those users to be as aware as possible of the conditions that the people they are responsible for are in. “[Increasing] the users ... level of awareness is more effectively achieved if the constraints of a conventional two-dimensional display are overcome. Systems with tangible interfaces or augmented reality technologies are examples of overcoming these constraints.”(Sugimoto, Hosoi & Hashizume 2004).

The requirements for a command and control application were observed as follows:

- The ability to represent multi-attribute data.
- Support for multiple views into the same dataset (ego and exocentric).
- The use of real tools and natural interaction metaphors.
- Support for a shared workspace.

As no single interface technology can easily satisfy these requirements, the SmartSpace projects aim was to develop hybrid interfaces which “employ multimodal input and display techniques to enable users to visualize and interact with spatial content in a far more efficient and intuitive manner than traditional interfaces.

3 Background

The SmartSpace proposal described a number of requirements, and gave some suggestions as to how to achieve them, but left the actual implementation open to any command and control application. The nature of command and control applications suggest a need for real time data flow, such that those using the system can issue their commands and get immediate feedback, allowing them to view the results of their actions and to refine ideas and plans. While the ideal implementation would be to have real information coming in from people out in the field via GPS or some other means so as to create a real scenario, this wasn't feasible at such an early point in the system design. Instead, we decided to make a simulation of an actual military operation in the form of a simple war game. This simulation would allow users to interact with the system in a realistic way, but would not require any potentially expensive running costs.

The game's dynamics were simple; two teams of two or more players have a number of tanks, and must gain points by destroying the opposing team's tanks or capturing "objectives". At the end of a time period, the team with the most points wins. As a real command and control situation may have people who are not interacting with system but are still observing it, such as a military adviser, the game should support spectators with the means of a publicly viewable interface. We felt that this simple scenario provided an accurate enough simulation so that the system could be designed and implemented as it would be in a real command and control situation. By building the game with a competitive nature, but with team collaboration, "the notion of public vs. private information, communication, and negotiation becomes a highly interesting field of research"(Magerkurth, Memisoglu, Engelke & Streit 2004). We believed that this public/private information span would provide interesting challenges in designing hybrid user interfaces, while the competition/collaboration aspects would encourage users to fully utilise the available interfaces.

3.1 Related Work

There is already a wealth of research and systems built around the idea of ubiquitous workspaces, including those using Augmented Reality. EMMIE, a system designed for collaborative augmented environments, was designed under the premise that "AR interfaces can enhance a ubiquitous computing environment by allowing certain parts of its hidden virtual layer to be visualized, as well as displaying personal information in a way that guarantees it remains private and customizable for each user." (Butz, Hllerer, Feiner, MacIntyre & Beshers 1999). In the EMMIE System, users are seated around a table, where they are able to use existing devices such as Laptop PCs to interact with the system. Optionally users can wear an optical see through Head Mounted Display, which allows them to see virtual three-dimensional content overlaid on the real world.

Because of this hybrid 2D/3D display environment, they had to develop a number of hybrid input interfaces. The devices used consisted of custom handheld devices which were tracked in three dimensional space by computer vision techniques. With the press of a button on these custom devices, users could drag and drop virtual content around in three dimensional space in a similar manner to how most two dimensional interfaces operate. The same devices could also be used for content privacy management. One such interface functioned in the virtual world as a spot light does in the real world. By moving the cone of light which was projected from this device over the virtual content, the visibility of the content could be changed, allowing it's users to choose which information they kept private and what would be shared.

The EMMIE system featured one type of interface which we considered using in the ARTankwars project. The interface ran on a Tablet PC or a Laptop if desired. The display space functioned as a

type of window between the real environment and the virtual environment. Using a camera positioned on the Tablet PC, it could display the real world with the virtual content overlaid on it. Our attempt at a similar interface is described in the Technical Details section.

One of the major problems which ubiquitous computing research has attempted to resolve is the stark physical boundaries between devices. One such example of this is the Augmented Surfaces project conducted by Rekimoto and Saitoh. Their aim was to overcome the problem that “When [information is] ... displayed on computer screens, information exchanges between computers often require tedious network settings or re-connection of computers”(Rekimoto & Saitoh 1999). The Augmented Surfaces project allowed devices such as Laptops to be used as they typically are, but removed the boundaries around them, by a technique named “Hyperdragging”. For example, a user may want to share some information they have on their laptop. A laptop is usually used as a private information space due to its small display area. In the Augmented Surfaces system, the user could share content on their laptop by simply dragging it off their laptop screen, and onto the desk in front of them. When the information reached the edge of the laptops screen, a projector mounted above the table would project the image onto the desk.

To help users identify what other participants are doing, when a cursor was displayed on any public space, they were “anchored” to their originating computer by a line draw from the cursor to computer. If any user wanted to view or copy information displayed in a public space, they could simply hyperdrag the information from the table onto their laptop. A camera mounted above the table recognises markers placed on the computers, and as such is able to detect their location so that the system could recognise when information had been dragged from a public space to a users computer. Users could also drag information across the table and up the wall due to another projector mounted horizontally. A server system runs silently behind it all transferring the information as required, creating a truly seamless virtual environment.

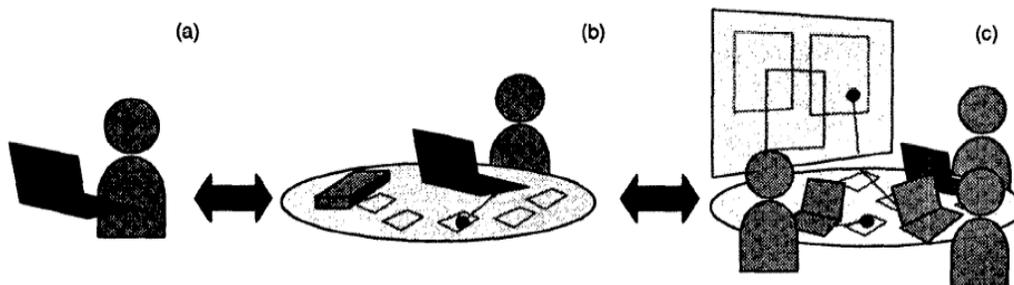


Figure 1: The Augmented Surfaces projects’ “Spatially Continuous Workspace”. A user can work privately on their own computer (a), then share content on the table top (b), or even display content on the wall (c)(Rekimoto & Saitoh 1999)

The Augmented Surfaces system focused on extending standard user interfaces into a ubiquitous collaborative environment, which was a considerably different domain from the ARTankwars project. However, some of the ideas used were common to both domains. The Augmented Surfaces project was able to work because of its server architecture, which co-ordinated all the clients, a concept which was especially useful for the ARTankwars system. By moving the responsibility for game logic and all non interface related control into the server, to change the system from using simulation data to a real command and control application would not require any major changes to the hybrid interfaces, only a change to the servers processing.

With the increasing frequency of mobile computing devices, such as laptops, PDA's, there is an increasing amount of interest in "the integration of various differently sized and shaped devices within ubiquitous computing environments"(Magerkurth, Stenzel & Prante 2003). This is especially becoming noticeable in Augmented Reality gaming areas. While there is a number of Augmented Reality games available, many focus on putting the player into the game, such as AR²Hockey(Ohshima, Satoh, Yamamoto & Tamura 1998), ARQuake(Thomas, Close, Donoghue, Squires, De Bondi & Piekarski 2002) or Human Pacman(Cheok, Goh, Liu, Farbiz, Fong, Teo, Li & Yang 2004). The STARS platform however was built for games where users control virtual characters, similar to the ARTankwars system. One such application built on the STARS platform is a collaborative tabletop role playing game, where users control a player around a virtual environment. In this application, each user has a tangible character model, similar to a small action figure, which they move around the table to control their character. These models are tracked by small transmitting devices built inside them, which can be detected by RF-ID antennae built inside the table. A roof mounted projector renders images onto the table top, to change the environment the players are in.

In the game, each player has a PDA which they can use to store and view private data about their player, and interact with the game world privately. Users can also use speech to interact with the game environment. Players can either work collaboratively or competitively to complete their goals. While the actual application is quite different between ARTankwars and the STARS role playing game, in both systems users are controlling virtual character(s). Both games also allow users to either work co-operatively or competitively and, as such, require the need for public and private interfaces to allow users to interact with the data accordingly.

4 Technical Details



Figure 2: The AR Interface and Projector Interface

There were two main sections in the ARTankwars project, the server and simulation logic, and the hybrid interfaces. While the internal architecture of the simulation is not the focus of this research, it is beneficial to understand some of the core design concepts, to provide some basis for the reasoning behind some of the interface design decisions.

The server, designed and implemented by Trond Nilsen, was written in the Java programming language. To minimise the low level networking details and allow us to concentrate on the application, the Internet Communications Engine¹ was used for all server and client networking. The Internet Communications Engine (ICE) allows processes to pass defined objects across the network, which minimises the amount of low level processing required, as well as enforcing a fixed data structure that communicating processes must adhere too. The ARToolkit² was used for all marker tracking and other Augmented Reality functions.

Upon start-up, the server loads a height map from a file which it uses to generate the terrain. It places all the tanks in two straight lines at opposite ends of the map, and objectives and buildings in the middle of the map. The exact locations of all these objects are described within the server's initialisation routines. Once these operations have been performed, it waits for interface clients to connect, adding them to the client queue when they do. When one of the clients issues the command to start the game, the counter is started. Several times a second the server will send an update to all clients informing them of any changes which have occurred since the last update, such as tank positions, objective owners, and any attacks that may have taken place. Clients provide input to the server such as tank destinations, which are then interpolated by the server to allow for a smooth animation of tank movements. The server is also responsible for details such as objective ownership, which is calculated by tank distance from objectives, collision detection, and whether tanks are within range of their opponents, so that they may fire upon them.

¹www.zeroc.com

²<http://www.hitl.washington.edu/artoolkit/>

Because all of this information is handled by the server, the system can be changed from a simulation to a live command and control application by simply removing the internal game logic, and passing data to and from the teams who are out in the field. No clients will need to be modified for this change to take place.

The following sections deal with the focus of this research, the hybrid interfaces. The sections have been divided up to match the “interaction spaces” as defined in original SmartSpace proposal.

4.1 The Public Space

The SmartSpace proposal described the display for the Public Interaction Space to be a wall projection display (see Figure 2). The input for the public space was to consist of Speech and Laser pointing, and PDA control. However, as a dedicated computer had to run the projector (since it would take up its whole display space), it was decided that the input functions could be built into the project display client itself. This meant we no longer needed an additional PDA to control the projector display, and could reduce the amount of hardware required by the system.

The main challenge for the projector display was deciding exactly what information was appropriate to make public. Due to the competitive nature of the simulation, not all information, such as the destination of tanks, could be made public; otherwise there would be no advantage for users to co-operate to out perform their opponents. Deciding how much information to show was one of the hardest areas of design for the public space, especially when taking into consideration the other interfaces. For example, originally we had planned for the Augmented Reality interface to only show each team the tanks which belonged to them, and what was in each tanks “field of view”. We felt that this would make for a more realistic simulation, as in a real command and control situation, those who are in command can only see information that the people in the field provide them. This would also allow another level to strategising, such that users could use buildings and other obstacles to obscure their tanks from their opponents. However, if this were to take place, then it would not be possible to display any information about tanks on the projector, as either team could look at the public display to find where their opponent’s tanks were. In contrast, if no information about the tanks were displayed on the public interface, then there would be no benefits of having a public interface at all, as it would display very little which could be considered useful.

After some planning and discussion, we decided that to make use of the public interface, the position of opponent’s units would be shown, and due to this all tank locations could be shown on the projectors display. We felt this was an acceptable compromise, as in a real command and control situation, there is often some information about uncontrollable circumstances, perhaps an estimation of the position of enemy units, or the location of a search and rescue teams target. In addition to the locations of all the units being displayed, we also decided to show who controlled each objectives, as it was something we felt would definitely be known in a real scenario, and information about the amount of damage sustained by each tank.

We decided that the best way to provide an overview of the entire scenario on a two dimensional interface was with a top down rendering of the terrain. While this perspective does not provide any information about the three-dimensional geospatial layout of the scenario, this does not impede the ability of any observers to gain an overview. A three-dimensional display would cause problems with occlusion, as well as perspective problems, such as deciding which area of the map should be closest to the screen.

In addition to the terrain, buildings and objectives are shown on the display as filled grey squares and filled yellow circles respectively. When a team captures an objective, the circle changes to the colour of the team which captured it, which is by default either red or blue. Tanks are shown on the

map in their appropriate locations, as well as their rotations. The turret of the tank rotates separately of the tanks body, which allows users to see both the direction the tank is travelling (shown by the direction of the body of the tank), and whatever tank it may be targeting (shown by the direction of the turret). To help observers identify the team which each tank belongs to, underneath each tank is a filled circle the colour of the team, or a grey circle for destroyed tanks. Tank fire is shown by lines drawn from the firing tanks turret to the location of the hit. If the hit is on another tank, a small explosion is drawn, to show that the shot was successful. In addition to this graphical display, there is also a grey panel which runs down the right hand side of the interface which shows each teams icons, their tank numbers and damage information, and any objectives they may have captured.

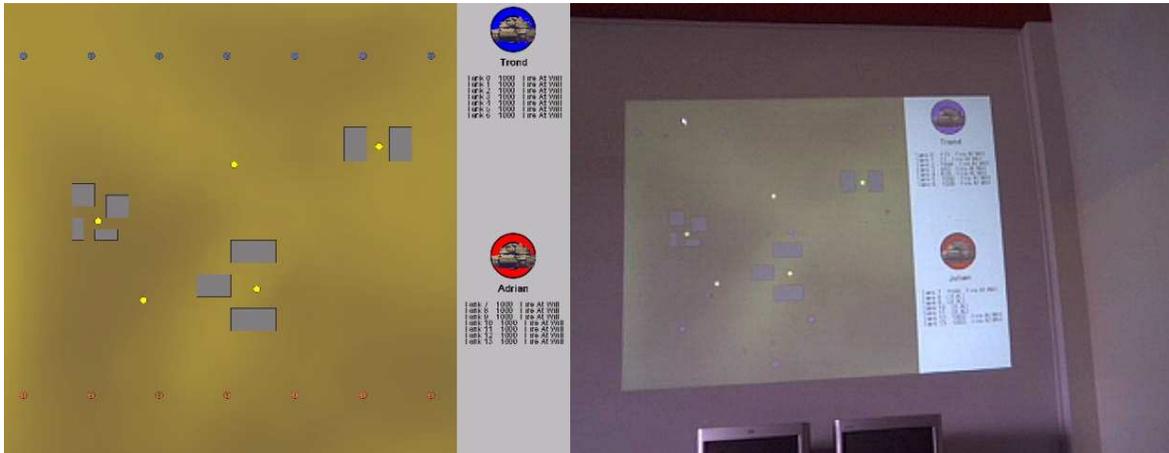


Figure 3: The Projector Interface, As a screenshot (Left) and projected on wall (Right)

Once the display section of the interaction space had been designed, it had to be decided how users could interact with the display. As the interface had to be able to provide an overview of the entire scenario, the level of detail was fairly low. We decided that the most appropriate way to increase the level of detail was to use a zooming interface, such that any area of interest, such as the location of a particular battle, could be made larger by increasing the zoom level. The interface could also be panned, so that the information space could be navigated at a higher level of detail.

For both the drawing and interaction, a toolkit called Piccolo³, which was designed for building two dimensional interfaces with support for zooming, was used. The projector client also used the ICE networking engine to communicate with the server. The projector client was written in the Java programming language, and as such could run on any platform which both Piccolo and ICE libraries ran on, which included Windows, Linux, MacOS and more. The client ran using the main computer display, so that it could run on a standard monitor or a projector.

4.2 The Group Space

The group space is where most of the interface design challenges took place. The SmartSpace project was to be centred on the Augmented Reality display. The major problem with Augmented Reality interfaces is that they're not tangible. While everything is displayed in a way that everyone can easily view and make sense of, it's difficult to interact with what you see, especially in a system such as

³<http://www.cs.umd.edu/hcil/jazz/>

ARTankwars. A lot of solutions which have been used in other Augmented Reality applications are not suitable in a situation like this, where users need to be able to execute a large number of commands in a short period of time. In addition to rapid interaction with the system, the competitive nature of ARTankwars requires these interactions to be private, which is not possible in many Augmented Reality interfaces. If users have to reach out or point at what they wish to interact with, then there is no privacy from their opponents.

We decided that for the group space there would be two modes of interaction, one for controlling individual tanks, and one which teams could use to plan strategies. The separation of these two interfaces is due to the different tasks involved. Planning strategies is something which users will work together on, and may take some time to complete. Controlling the individual tanks is something which users do separately, and in a lot of situations will need to be completed quickly for the results to be effective. In addition, the actual process involved in designing strategies is quite different to that of moving tanks, and due to this, it would be awkward to have one interface used to perform both.

The following sections discuss the two separate interfaces which were developed for the control of tanks, and the annotation of the map.

4.2.1 The Augmented Reality Interface

In ARTankwars, the Augmented Reality interface is the centre of focus. Participants get most of their visual feedback from this display, and because their concentration is on the display so much, it makes sense to try and include a means of interaction into this display space. One of the major advantages of Augmented Reality in Computer Supported Collaborative Work (CSCW) is that it supports separate views for each user. This means that while one member of a team is concentrating on one subset of the data, the other user can be working on a completely different subset concurrently. Separate views are also advantageous for competitive work in a localized area as well. In ARTankwars all users, even those on different teams, are working on the same workspace. In a typical situation, this means that any interaction with the workspace would be visible to all users. But with Augmented Reality, we can make sure that interactions and changes in the workspace are only visible to the people who should see them, such as the user who initiated the change, and their team members.

The SmartSpace proposal stated that Augmented Reality is excellent for displaying three-dimensional geospatial content. Most current command and control applications utilise two dimensional maps, however there is a lot of information in the third dimension which either gets discarded or requires graphical or textual annotations to the map, which requires a greater mental load for users to process, which in turn can lead to mistakes. A true three-dimensional representation can show important information about a real environment that a two dimensional representation may be unable to provide. Digital representations of environments have added benefits over real ones, due to the fact that they can be updated instantly with new information, and data examination techniques such as zooming and filtering can be achieved easily. Digital representations also present the ability for multiple views into the same data set, such as a top down view or a first person view from the perspective of someone in the field. Despite their strengths, digital representations on a monitor are still only displayed on a two dimensional plane, even if they are from a three-dimensional model. Augmented Reality resolves these problems, as it allows us to display digital representations in true 3D, as it moves the digital information out of the virtual space and into the real world.

To allow the users of the system to make effective and educated decisions, they need to know as much information about the scenario as much as possible. We decided the best way to provide a maximum amount of information in an easily understandable way was to provide users with a miniature replica of the environment. The terrain is displayed as a three dimensional model, as are all

the tanks, objectives and buildings. The Augmented Reality Interface uses some of the same design ideas as the Projector interface, such as displaying coloured circles under the tanks to show which team they belong to, or grey circles if they had been destroyed.

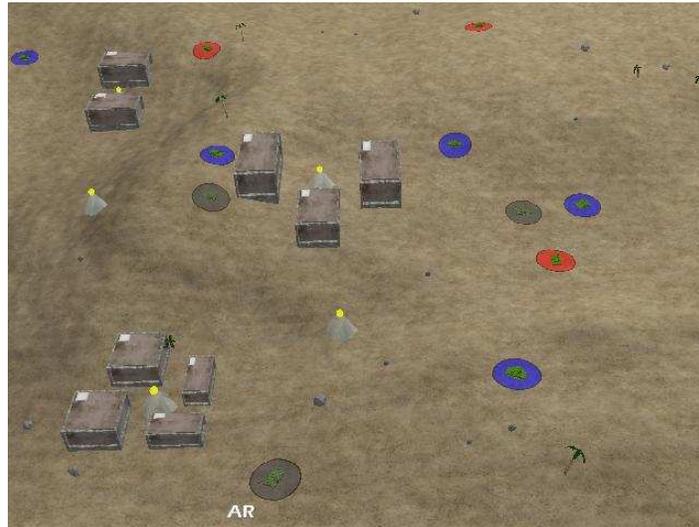


Figure 4: The Augmented Reality Interface

AR headsets cause issues when working in a collaborative environment, as they obscure the eyes of the user wearing them. When someone is talking about a subset of data in a workspace, it is often possible to determine what information they are referring to by following their gaze, something which isn't possible when the user's eyes are obscured. To resolve this problem, a "cone of vision" is displayed in the Augmented Reality interface. As shown in Figure 5, the cone of vision is essentially a translucent cone with its apex at centre of the users headset, and its base at what the headset is pointing at on the interface. By following a user's cone of vision to the environments display, we can tell where each user's focus lays. This was achieved by a feature in the ARToolkit which allows the calculation of the cameras location based on its orientation to the AR Markers.

The input interface for controlling the units was more challenging to design. We needed to provide a way for users to easily and quickly interact with a non tangible interface, while at the same time keeping these interactions private. The SmartSpace proposal suggested interfaces such as speech and natural gestures, as well as tangible objects. Pure speech input has problems when attempting to control a large number of similar objects at once. Without some other means of interaction, it is very difficult to communicate to the system which object you are wishing to interact with at any point in time. Natural gestures and tangible objects tend to be slow to use, and can often lead to problems in group situations where users are all working on the same workspace. An example of this occurs when a user wants to interact with information on the opposite side of the workspace to where they are located. To do this they must lean over to reach it, which may block other users from working. In doing so, users may also reveal the intention of their interaction, which gives competing teams knowledge about their private strategies.

One interface we were interested in was an Augmented Reality magnifying lens developed by Julian Looser. This interface acted like a normal magnifying lens, however the lens only magnified virtual content. With some extension this tool, we realised that it could also be used as a method for selection, such that with some input event, which ever object was in the centre of the lens would be

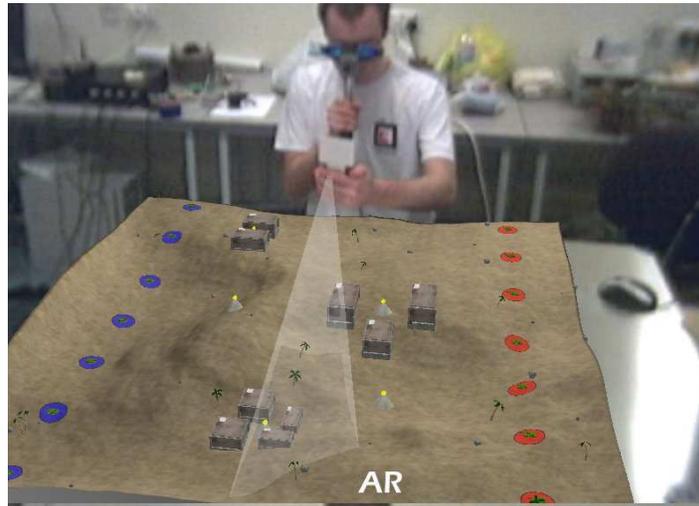


Figure 5: The Cone of Vision

selected. The lens was now dual function, allowing for both zooming and command input. The lens worked by calculating the position of an Augmented Reality marker in relation to the environment Augmented Reality markers, which were used for the system to calculate the location of the terrain. When this lens marker was attached to the top of a PDA, we now had a complete device which was capable of interacting with the Augmented Reality interface. The selection system removed any ambiguity which occurs when issuing commands via speech. By moving the lens over an object, users could issue commands by voice, and the system would be able to identify what object to apply the commands to. By attaching the lens to a PDA, users could use the PDA's buttons as input, allowing users to interact privately with the system.

An issue with Augmented Reality is that there is no depth ordering, so that the virtual content gets draw over all real content, regardless of distance. This caused problems when using the PDA and lens system, as using the lens for selection usually meant that the PDA would be obscured by virtual content. To solve this, the AR client was designed so that when it recognised the lens marker, it would not draw any virtual content in a PDA sized window below it. In addition, virtual labels were drawn which gave text descriptions of what each of the buttons on the PDA did (See Figure 8). Above the PDA was another virtual label which described what kind of input the system was currently expecting. This provided a clear explanation of the controls.

To avoid confusion, the PDA interface used a modal design, preventing users from entering invalid commands. For example, when no unit was selected, the only commands available were select, zoom in, and zoom out. Once a unit was selected, users could order a move command, an attack command, and a cancel command. If the user then issued move or attack, they would be presented with options to select their target, zoom in or zoom out, or cancel the previous command.

4.2.2 The Tablet PC Interface

Originally we had planned for the Tablet PC to be used in a similar fashion to the EMMIE system "magic mirror", allowing users to "open a portal from the real world to the virtual world for those EMMIE users who are not wearing head-worn displays and who otherwise cannot see the 3D virtual

objects”(Butz et al. 1999). By mounting a web camera on top of the Tablet PC, the ARToolkit could use the video feed from the camera to create the Augmented Reality scene. Using a second Augmented Reality interface would allow for a separate view, such as a wire-frame mode, allowing users to see through buildings and the terrain, removing any occlusion problems they may have with the textured interface.

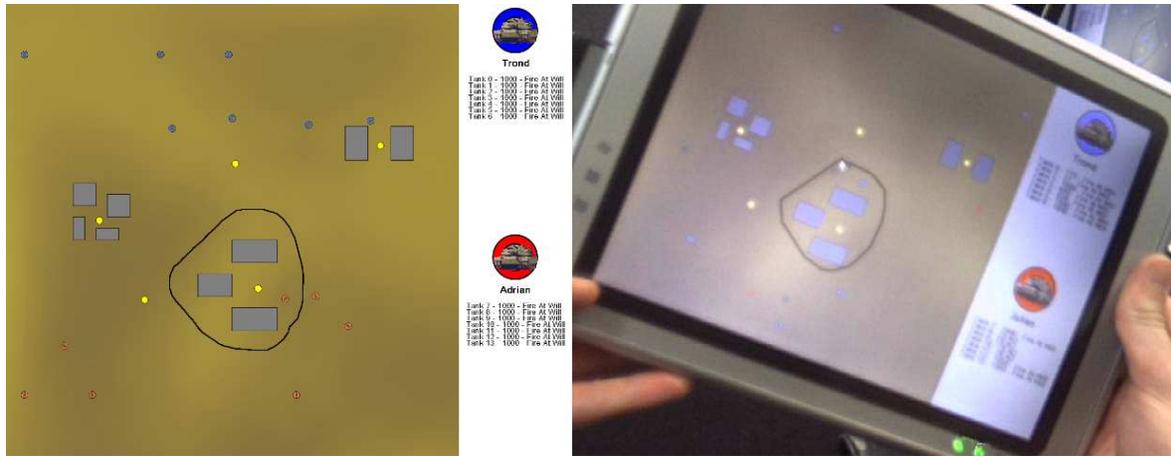


Figure 6: The Tablet Interface, As a screenshot (Left) and running on the Tablet PC(Right)

However, very early tests of the ARToolkit on the TabletPC showed that the hardware was just not powerful enough to deliver a real-time augmented video stream. There was some discussion about offline processing, by taking shots from the camera and sending them to a more powerful computer for processing, to be sent back to the Tablet PC for display, however this would increase bandwidth requirements, and was unlikely to offer much of a speed increase. Eventually we decided to discard the idea, as another Augmented Reality interface was superfluous to the requirements of the SmartSpace project. We decided to look for other possible uses for the TabletPC. We noticed that, while there were benefits for users to work together, the system design didn't support much collaboration between users on a team. We had decided early in the project that it would be useful if users could annotate the Augmented Reality view, much like a person might draw on a real map with a pen, so that users could communicate their ideas to each other. However, this would be difficult and time consuming if annotations were to be made with the lens and selection system. This would cause additional problems as "It is not easy to add annotations ... while another participant is presenting his data"(Rekimoto & Saitoh 1999), and during time it took one user to make annotations with the lens, there could be events taking place which required a user response which may invalidate the annotation altogether. If annotations were to be made in a real time environment such as this, they would need to be generated quickly and easily for them to be of any use, as "people are more concerned with their task at hand and would rather use collaborative tools only to the extent they aid them with their task and not interfere with it"(Pingali & Sukaviriya 2000).

The need for quick and easy annotation was perfect for the Tablet PC. With the Stylus based interaction method, Tablet PCs are ideal for drawing and making quick notes. While the PDA also includes a stylus, its small display size would limit the amount of annotations which could be made, and it would make it difficult to highlight smaller areas. We decided the best way to make annotations on the interface would be from a top down perspective, as the location of the objects on the two dimensional plane was more important than their vertical position, especially all the objects of interest

were positioned on the ground. A simple two dimensional interface was also practical as it would not be as demanding of the Tablet PCs hardware as a three-dimensional display, and, as the Projectors display was in the same format, we could use the same display system for both clients and only the interaction functionality would require change.



Figure 7: The AR Interface, showing Annotations in light blue

In the final system, users can use the stylus of the Tablet PC (or a standard mouse attached to the Tablet PC if they so choose) to draw on the Tablet PC's screen. As shown in Figure 6, when a user makes an annotation, a thick black curve appears on the Tablet PC's interface. Users could then share the Tablet PC with other users to discuss the annotations made. In addition to this, when users draw on the Tablet PC's screen, the markings appear in the Augmented Reality view as a thick line, the colour of the team which made the annotation. As we can see in Figure 7, the area highlighted in Figure 6 is shown as a thick blue line in the Augmented Reality interface. To keep these annotations private, users could only see the annotations which their own team had made.

4.3 The Private Space

The Private Space also had two main interfaces which allowed users to view and interact with the data, the Private Augmented Reality window, and Transitional Interfaces.

4.3.1 The Private Augmented Reality Window

The SmartSpace proposal suggested that the Tablet PC be used for the Private Space, in addition to any handheld displays. We believed however that the Tablet PC was more useful as a device for annotating the Group Space, due to its stylus based input and reasonably large display size. The PDA was being used as a controller for the tanks, but it had a screen which was not in use, which could be used as a private display space. This seemed like a good idea, users could select a tank with the PDA, and the PDA's display could provide a list of information about that tank which the user could see. Users could also make private notes on the PDA, which they could then review later.

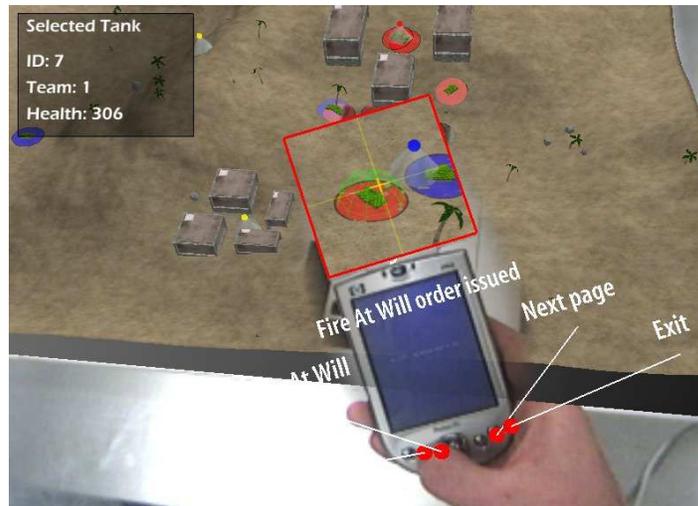


Figure 8: The Private Space Window and PDA Interface

However, when we actually tried viewing the PDAs screen through an AR headset, the low resolution of the headsets camera made the screen very difficult to read, unless the PDA was very close to the headsets camera. However, the headsets camera is not mounted exactly between the eyes where we expect our perspective to be, but rather slightly above the eyes and angled down. When combined with the low resolution of the cameras, this caused significant problems when trying to read the screen. Upon not being able to read it, a user's natural instinct would be to bring it closer to their eyes. However, as the camera was mounted above the eyes, doing so would move the PDA out of the frame. When testing the PDA's display initially it was not uncommon to see users waving the PDA in front of their face trying to find the perfect position to hold it so they could actually view the display. With users having so much trouble trying to view the screen, we realised that the display size was just too small to use in Augmented Reality, and that it would be even more difficult to write on with the stylus when the user couldn't view the display properly.

The idea behind the Private Space is that individual users could have their own information, which no one else, not even their own team members, could view. After some discussion, we realised that we already had a display which supported this, the Augmented Reality display. The AR Display already supported information hiding, so that teams could make annotations which weren't visible to the other team, all that was required was to extend this idea further so that users could also have individual information, which would not be shown to anyone. A small translucent window was added to the top left of the display, which provided all the details which we had planned on implementing into the private interface. This display was still controlled by use of the PDA, such that only relevant information would be shown. When a tank was selected, the window would display details about its id number, health and other relevant information. When nothing was selected the window was empty. The only drawback to this interface was that there was no way for users to make their own amendments to it.

4.3.2 Transitional Interfaces

The final concept which was implemented from the SmartSpace proposal was Transitional Interfaces. Computer interfaces can be placed along a Reality-Virtuality Continuum based on how much of the

user's environment is computer generated, as shown in Figure 9.

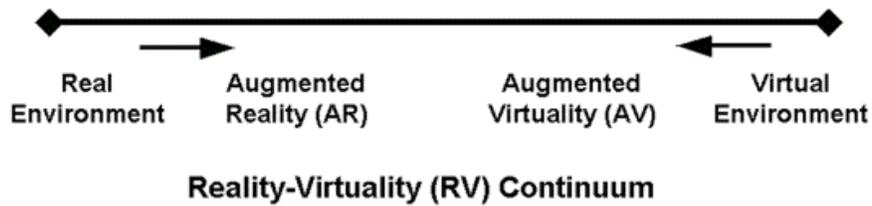


Figure 9: Milgrams Reality-Virtuality Continuum (Milgram & Kishino 1994)

As we can see, a shift towards the right on this continuum occurs when more of the user's environment is replaced by virtual content. Our Augmented Reality display is shown towards the left of the continuum. A transitional interface is one which moves along the continuum. For the ARTankwars project, we thought it would be beneficial if users could transition themselves from a giants eye view overlooking the map, to being down on the map with the units they are controlling. This would allow users to change their perspective to that of a person in the field, allowing them a better understanding of the situation the people in the field are facing.

This change from an Augmented Reality view to a Virtual Reality view is a transition across the entire Continuum. To make this change as smooth as possible for the users, they are “flown” into the scene, which gradually fades from being augmented reality to virtual reality, as shown in Figure 10. Users can initiate this transition by selecting the location they wish to travel to with the PDA lens, then pressing the appropriate button. When in virtual reality mode, a virtual avatar displayed on the map to all members of their team, allowing team members to see where a user is viewing the map from.

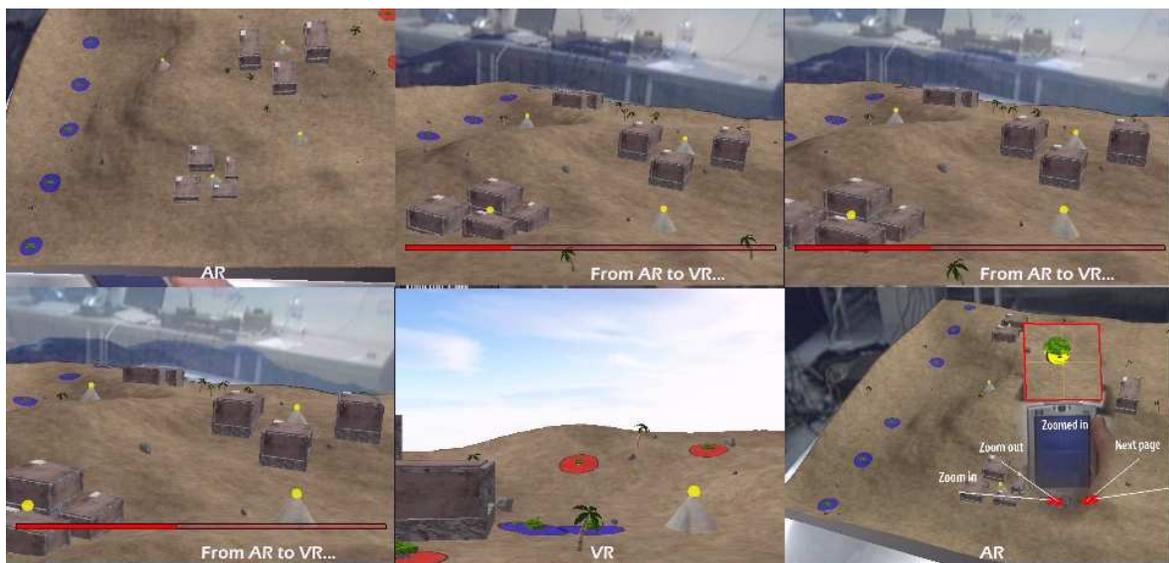


Figure 10: From Left to Right, Top to Bottom, a transition from Augmented Reality to Virtual Reality. The bottom right image shows a Virtual Avatar.

5 Results

Due to the nature of this research, it is very difficult to get any meaningful quantitative values as results. While we have not conducted an in-depth empirical study of the system yet, following the systems completion the HITLab NZ had an Open House, where members of the public were encouraged to come in and try some of the projects available, including ARTankwars. This gave us a valuable insight into how novice users interacted with the system. Also discussed are our own experiences with the system throughout development and testing.

5.1 Our Interaction with the System

During the development phase of this project, considerable testing was conducted to make sure the interfaces worked correctly and to look for any problems. The projector interface was one of the first interface clients to be usable. The zooming interaction offered by the Piccolo library was initially unusual, as it allows zooming to be performed by a right mouse click where the user wants the focal point of the zoom, followed by a mouse motion to the right to zoom in, or to the left to zoom out. After some use of the interface, it became apparent that it was very difficult to zoom in on an area, and then zoom back out to the original overview of the scene. To resolve this, another function was added, where by pressing the space bar, the entire scene would return to its default zoom level and location.

We found that it took some time to become competent with the lens selection system. As the ARToolkit is susceptible to occlusion, if the PDA interface covered too much of the AR markers representing the map, all the Augmented Reality display would disappear. After some practise though it became easier to position the PDA interface such that both the Lens marker and the map markers were visible. A proposed solution to this problem is described in Section 7.1

The investors in the SmartSpace project were based in Singapore, and towards the end of the development of the system, Trond and Julian went to Singapore to give a demonstration of the ARTankwars system. In transit, both of the PDAs which were being used for the system were damaged and as a result could no longer function. While this was an inconvenience, the decision to not use the PDA's screen as a display space meant that there was no need for a device with a screen, in fact any device with buttons could do the job just as well. We were able to secure access to some wireless game pads, and after moving the input handling code from the PDA to the computer which the game pads were connected to, the game pads entirely replaced the PDAs. The game pads had an advantage over the PDAs as their design was better suited to having their buttons pushed rapidly. After we had modified the system to use game pads, we decided they were better suited to the ARTankwars application than the PDAs had been.

After completion of the Tablet PC interface, it became apparent that the lack of computing power that the Tablet PC suffered from was more of a problem than first anticipated. The Tablet PC's interface was originally developed on an ordinary PC, using a mouse pointer instead of a stylus, and ran at a decent speed. However, when running on the Tablet PC, the program was quite slow to respond to interaction made by the stylus. While the interface worked correctly, it was frustrating to have to wait a few seconds for the annotations made to show up on its display. We attribute this lack of speed to the operation of Java and the Piccolo library on the Tablet. In future we plan to develop a C++ and OpenGL version of the Tablet PC client, as this lower level language and graphical library should provide a considerable speed increase for the client.

5.2 Novice Users Interaction with the System

The HITLab NZ Open House provided us with an excellent opportunity to get novice users feedback on the system, as well as the opportunity to observe how users interacted with the system to look for any possible problems. Unfortunately, due to large numbers of people trying the system for only a short period of time, users of the system did not have an opportunity to fully experiment with all aspects of the project. Because users would only be having a short use of the system, we decided not to include the Tablet PC in this trial. There were technical problems with one of the AR Headsets was not displaying an image, so it was used as a camera which displayed the Augmented Reality view on a regular PC monitor. This meant that only one user could be interacting with the system at a time, but this did not impact the trials considerably as most users only had enough time to try the interface as opposed to actually competing with other people in the simulation.

The range of users who tried the system varied greatly, all different ages, ethnicities, genders and backgrounds. Most users were fairly inexperienced with Augmented Reality interfaces, and only some had ever been in a Command and Control simulation such as a computer based war game. When a user first put on the headset, they were given a very brief overview about how to use the interfaces provided. Users were then allowed to interact with the system by themselves, but given feedback and assistance if they ever became confused. It was common for one user to be trying the system while others observed, watching the monitor displaying what the user saw.

Many users had problems initially with the AR Interface. This was partially because most people are not used to wearing the headsets, and also because of tracking problems with the Augmented Reality markers. Users would often move their head suddenly, causing a loss of tracking and the display to disappear. Nevertheless, most users were enthusiastic about the interface, though almost all agreed that the Augmented Reality marker tracking needed improvement.

Despite our change from the PDA to the easier game pad controls, users still had trouble controlling the selection interface. This was in part due to the layout of the controls, as some were attached to “shoulder buttons” on the underside of the game pad. Some users also had trouble understanding how the selection system worked, and had to be shown how to centre a tank under the lens and push the select button. However, by far the biggest problem users had with this interface was the occlusion problem mentioned earlier. Initially almost all users had problems positioning the game pad in such a way that both the lens AR Marker and the environment AR markers were visible enough that the ARToolkit could recognise them.

All users were impressed with the system, many expressing interest in further development to a commercial quality product. Some users even stayed round for a long period after trailing the system to discuss its features and development, and observe other users interacting with it. Even those users who had some trouble using the interfaces at the start were able to achieve a reasonable amount of competence with the system with the aid of some verbal guidance. A young male participant was even able to complete the scenario, by capturing all the objectives and destroying all the opponents’ tanks. We felt these positive reactions were a good sign that our interfaces were usable with some practise, even to novices.

6 Discussion

The SmartSpace project proposal listed four requirements which command and control applications, like ARTankwars, would need to fulfil. The following sections describe each of these requirements and how they were handled in ARTankwars.

6.1 The ability to represent multi-attribute data

In a command and control system, the users need to be able to make sense of a large amount of data to be able to perform effectively. The system needs to present all the data in an efficient way to lower the mental workload of these users. To do this, users should be able to get an overview of the entire scenario, so that at a glance they can get enough information to react to any important issues. It is also important that users can also get more information when required.

In ARTankwars, there is a large amount of data available to users. By default, the two main display interfaces, the Augmented Reality and Projector display, only display a minimum amount of data about terrain, tanks and objectives. This is all presented in a form which is instantly understandable, from either a three dimensional giants perspective, or a two dimensional top down perspective. This minimal overview also relates to the public nature of these information spaces; the amount of information shown on these displays allows participants and observers to understand the scenario without exposing any private information of individual teams or users.

When users wish to access more information, they can use the Tablet PC interface to make annotations or strategies with other team members. When users select one of their tanks with the lens selection system, they are presented with in-depth information about that tank in their own Augmented Reality view. In addition, users may choose to use the transitional interfaces to change their perspective from that of someone overseeing the field, to someone in the field. This allows users to obtain a better understanding of the situation the people they are commanding are facing.

6.2 Support for multiple views into the same dataset (ego and exocentric)

With a large amount of data available, it is useful to be able to have different views of that data, so as to highlight certain features. This is essentially a form of filtering the data, allowing users to view what information is relevant to them at that point in time. In ARTankwars, each display interface represents a separate view into the data set. The AR display shows a three dimensional geospatial view from an exocentric perspective. However, with the push of a button, users can “fly-in” to the display with the help of the transitional interface, and get an egocentric view of the same data set. The Projector and Tablet PC interface also offer a separate view of the dataset, providing a top down view of the map, providing participants and spectators with an overview, without inundating the viewer with information

6.3 The use of real tools and natural interaction metaphors

One of Augmented Reality’s strengths is its ability to introduce Virtual content into the real world, allowing a logical view into what could possibly be a complex data set. The drawback of this is that it is difficult to interact with virtual data using standard computing practises. Concepts need to be borrowed from real domains which can be applied to the Augmented Reality interface so that interaction with the virtual content is as logical and straight forward as possible.

The ARTankwars project uses a lens selection system, where objects which are within the centre of a virtual magnifying lens are selected. The concepts behind using a magnifying glass are natural

to most people, and so we felt a similar interaction metaphor in this project would be easy for users to understand. In addition to this selection method, users can use speech, one of the most natural interaction methods humans have, to communicate with the system.

When users wish to make annotations or plans with other users, they can use the Tablet PC and stylus to draw on the top down view of the map. This concept of drawing on a two dimensional information space is an action which most people are very proficient in, and by including the annotations on the Augmented Reality view, provides users with a natural way of interacting with this non tangible interface.

6.4 Support for a shared workspace

In typical collaborative environment where the focus is not on the computer, such as a business meeting, the users' task space and communication space overlap each other, as shown in Figure 11. For example, if people are sitting round a table with some information on it, when they are working on the information they are still facing the people they are collaborating with, which allows those people to pick up on certain cues, such as gaze, facial expressions, and other body language. However, in computer supported collaborative work when the workspace is on a computer screen, the participants focus is on the computer where the work is taking place, instead of the people they are collaborating with. "When players gather together around these consoles the players generally sit side-by-side and interact with the screen, not directly with each other. Thus, their interactions occur via the interface"(Mandryk, Maranan & Inkpen 2002). In addition, typical computer interaction methods are only designed for one user at a time, meaning that true collaboration is not taking place.



Figure 11: Left: Users working around a computer, the task space and communication space are separate. Right: Users working around a table, the task space and communication space overlap ⁵

In this project, the Augmented Reality display brought the virtual content into the real world, bringing the workspace and the communication space together, similar to that business meeting example. "When using a collaboration-table the main focus of attention lies on the public surface of the table that holds a shared artefact such as an architectural plan or a geographical map [for example, a] 'commander's table'"(Magerkurth & Tandler 2002). All the users are positioned around the same workspace, and are able to see each other. Even though the AR headsets obscure people's faces, the "cone of vision" provides information about a users gaze. While users are all working on the same dataset on the same workspace, the augmented reality display also provides the functionality to allow users to have their own private data and team data displayed in their field of view.

7 Further Work

The ARTankwars system development had deadlines set which had to be reached, and there were a certain number of requirements which had to be met in the time available. Because of this, there were a lot of features which we wanted to add into the system but were unable to in the time available. This section describes a number of features which we would like to implement into the ARTankwars system to improve it.

7.1 Markerless Tracking

As mentioned by most users who trialled the ARTankwars system, the lens and selection interface has significant problems, as it tends to occlude the environmental AR markers, causing the ARToolkit to lose tracking, and the loss of the Augmented Reality display. It takes some practise to use an AR headset also, as any sudden movements can cause loss of tracking as well. For optimal results, users must make slow steady movements with their head. These problems can be minimized by a technique called Markerless Tracking. The ARToolkit uses black and white markers with a pattern inside to determine where to draw virtual content in real space. Markerless Tracking uses existing objects, such as maps, as markers. An algorithm is run on the object, which scans the object for areas of interest, for example lines, corners and gradients. It generates a library of these areas, which are then used for object identification by the ARToolkit at runtime. Because there are a large number of these areas of interest, as opposed to a single AR marker, the ARToolkit has more objects which it can track, resulting in more accurate tracking. The added benefit is that the objects used for tracking can be relevant to the Augmented Reality display. For example, in ARTankwars, the Augmented Reality display shows a map. With a markerless tracking algorithm, the object tracked could be the corresponding two-dimensional map. This would mean that even users without an AR Headset would still be able to see some form of visual representation of the domain.

7.2 SmartSpace Framework

The research into hybrid interfaces in this project resulting in some interesting solutions to interaction problems. Ideally a framework could be build into the ARTankwars system which would increase its flexibility. This framework could include the ability to add more teams, change display and input interfaces, or even remove the entire game logic as mentioned earlier in this report, so that application can be change to a real command and control system, allowing the system to be used for military or search and rescue operations.

7.3 Intelligent Projector Display

The main problem with the projector interface was that it required a user to perform any zooming operations. It would be useful if some algorithms could be developed for this interface so that the system could automatically choose areas of interest, perhaps when there was a battle occurring between a number of tanks, and zoom in on this area of the environment. Another field of possible investigation would be to explore other methods of interaction with the projector interface. Some examples would include using computer vision techniques to track a laser pointer across the surface which the projector display is displayed on, or using voice commands to control the interfaces zooming and panning abilities.

7.4 Outdoor AR Interface

There has been a wide range of research into uses of Augmented Reality into outdoor environments. In a real command and control scenario using this system, it would greatly aid the people in the field if they were able to have information sent to them from their commanding officers through the system. An example of this would be if a commanding officer were to make annotations on the map using the Tablet PC interface, and then these annotations could be sent and displayed to the people in the field. In addition, those people could be tracked using GPS, and their positions updated on the display immediately, allowing those in the commanding situation to get immediate feedback about all their personnel.

7.5 Empirical evaluation

While the trials we had with novice users showed promising results and that users liked the interface, there is no empirical evidence to prove whether the system has any significant effect on the performance of users of the system. While it is difficult to extract any relevant quantitative results from a study of a system such as this, it would still prove useful to observe a group of experts in the command and control domain, to see how they interact with the system. These experts would also be able to provide feedback in regards to how useful they believe the system to be, and any strengths and weaknesses they observe. In the future we hope to conduct empirical evaluations of the system with people who are very familiar with command and control situations, such as Table Top war gamers, Civil Defence and Search and Rescue teams. The feedback from these sessions will allow us to find the strengths and weaknesses of the system in a real event, and to improve the system accordingly.

7.6 Rewrite Tablet PC Client

As noted in the Results section, the Tablet PC's hardware had problems dealing with the Java client and Piccolo image library. To really get the most out of the Tablet PC, the client will need to be rewritten in a language that the Tablet PC can execute faster. Already work has begun on a C++ client using OpenGL, which should decrease the response time of the Tablet PC client so that users can interact with it without an obvious delay.

8 Conclusion

The ARTankwars project was a system designed to fulfil the requirements of the SmartSpace proposal. The SmartSpace proposals main focus was to “explore how hybrid user interfaces can be developed that employ multi-modal input and display techniques to enable users to visualize and interact with geospatial content”. In particular the applications suggested for such a system was that of a Command and Control system, such as for military or search and rescue purposes.

ARTankwars is a war game simulation developed to represent a Command and Control scenario. It encourages a collaborative effort between team members, and competition between teams. This presented a number of interesting challenges for data privacy and data sharing across the interfaces. The main workspace is a three dimensional environment implemented in Augmented Reality. To allow users to all share this same workspace, some new interaction techniques needed to be designed and developed to allow users to easily work with this virtual content.

The interfaces developed allowed users to work by themselves or in teams in ways that were easy for users unfamiliar with the system to learn. In trials with novice users most of them had some problems initially with the interfaces, as we expected given that they were quite different to interfaces most people were familiar with. However, after some explanation and practise, all users were able to use the interfaces provided to interact with the workspace. Many found that once they had some experience with the interfaces that they were quite natural and efficient to use.

The interfaces developed needed to support Public, Group and Private work. The Public workspace was shown on a projector which displayed a two dimensional top down overview of the map, showing tanks, objectives and buildings, and included a side panel describing the current state of each tank, and the owners of the objectives. The Group workspace was displayed as a three dimensional miniature environment in Augmented Reality, and a Tablet PC interface with a similar display as the projector interface allowed users to make annotations on the Augmented Reality display, by simply drawing on the Tablet PC with the Stylus. A custom designed Virtual Magnifying Lens and Selection interface, which could be used in combination to a game pad or PDA, allowing users to control Augmented Reality content by focusing the lens on it, then pressing buttons on the game pad or PDA. The private space was displayed in each users Augmented Reality view, as a translucent window which displayed relevant information about objects they had selected with the Lens and Selection device. Users could also “fly” into the environment, allowing a first person perspective into the world.

The system was completed and fulfilled all the requirements of the SmartSpace proposal. However, there are still a number of features which we, the authors, wish to implement, to provide a richer and more stable experience. While no official empirical studies have been conducted, an Open House allowed members of the public to come and trial the system. Many users enjoyed using the system and, while they were aware of a number of problems with the system, believed that it had a lot of potential for further application.

References

- Billinghurst, M. (2003), 'Smartspace: A hybrid interface for ubiquitous computing'. The proposal for the SmartSpace project. Unpublished.
- Butz, A., Hiller, T., Feiner, S., MacIntyre, B. & Beshers, C. (1999), Enveloping Users and Computers in a Collaborative 3D Augmented Reality, *in* 'Proceedings of IWAR'99 (International Workshop on Augmented Reality) San Francisco, CA, 20–21 October', IEEE Press, pp. 35–44.
- Cheok, A. D., Goh, K. H., Liu, W., Farbiz, F., Fong, S. W., Teo, S. L., Li, Y. & Yang, X. (2004), 'Human pacman: a mobile, wide-area entertainment system based on physical, social, and ubiquitous computing', *Personal Ubiquitous Comput.* **8**(2), 71–81.
- Magerkurth, C., Memisoglu, M., Engelke, T. & Streitz, N. (2004), Towards the Next Generation of Tabletop Gaming Experiences, *in* 'Proceedings of Graphics Interface 2004 London, Ontario, 17–19 May', Canadian Human-Computer Communications Society, pp. 73–80.
- Magerkurth, C., Stenzel, R. & Prante, T. (2003), STARS - A Ubiquitous Computing Platform for Computer Augmented Tabletop Games, *in* 'Video Track and Adjunct Proceedings of the Fifth International Conference on Ubiquitous Computing (UBICOMP'03) Seattle, Washington, 12–15 October'.
- Magerkurth, C. & Tandler, P. (2002), Augmenting Tabletop Design for Computer-Supported Cooperative Work, *in* 'ACM CSCW 2002: Workshop on Co-located Tabletop Collaboration: Technologies and Directions, November 16. New Orleans, Louisiana.', ACM Press.
- Mandryk, R., Maranan, D. & Inkpen, K. (2002), False Prophets: Exploring Hybrid Board/Video Games, *in* 'Proceedings of CHI'2002 Conference on Human Factors in Computing Systems. CHI Letters 4(1). Minneapolis, Minnesota, 20–25 April', ACM Press, pp. 640–641.
- Milgram, P. & Kishino, F. (1994), A Taxonomy of Mixed Reality Visual Displays, *in* 'IECE Trans. on Information and Systems (Special Issue on Networked Reality)', pp. 1321–1329.
- Ohshima, T., Satoh, K., Yamamoto, H. & Tamura, H. (1998), Ar2 hockey: A case study of collaborative augmented reality, *in* 'Proceedings of the Virtual Reality Annual International Symposium', IEEE Computer Society, p. 268.
- Pingali, G. & Sukaviriya, N. (2000), Augmented Collaborative Spaces, *in* 'Proceedings of CHI'2000 Conference on Human Factors in Computing Systems The Hague, The Netherlands, April 1–6', ACM Press, pp. 13–20.
- Rekimoto, J. & Saitoh, M. (1999), Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments, *in* 'Proceedings of CHI'99 Conference on Human Factors in Computing Systems Pittsburgh, May 15–20', ACM Press, pp. 378–385.
- Sugimoto, M., Hosoi, K. & Hashizume, H. (2004), Caretta: A System for Supporting Face-to-Face Collaboration by Integrating Personal and Shared Spaces, *in* 'Proceedings of CHI'2004 Conference on Human Factors in Computing Systems Vienna, Austria 24–29 April', ACM Press, pp. 41–48.

Thomas, B., Close, B., Donoghue, J., Squires, J., De Bondi, P. & Piekarski, W. (2002), 'First person indoor/outdoor augmented reality application: Arquake', *Personal Ubiquitous Comput.* **6**(1), 75–86.

Appendices

Appendix A: Additional ARTankwars Images

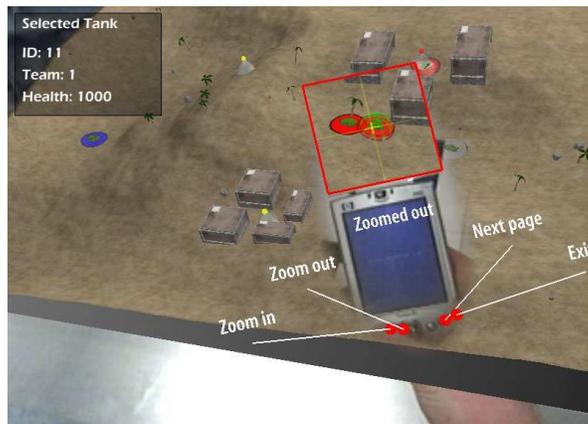


Figure 12: The lens after zooming out to a 1:1 zoom ratio

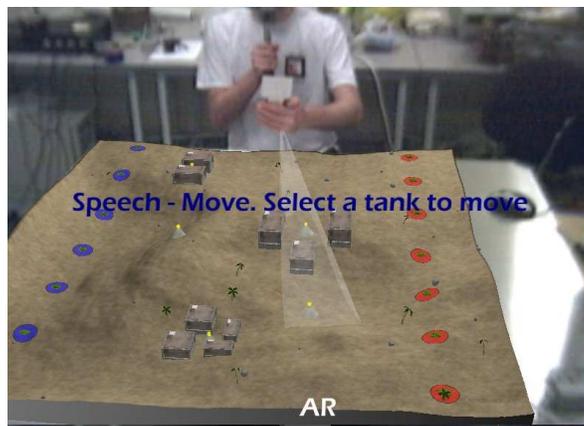


Figure 13: The speech recognition system in use

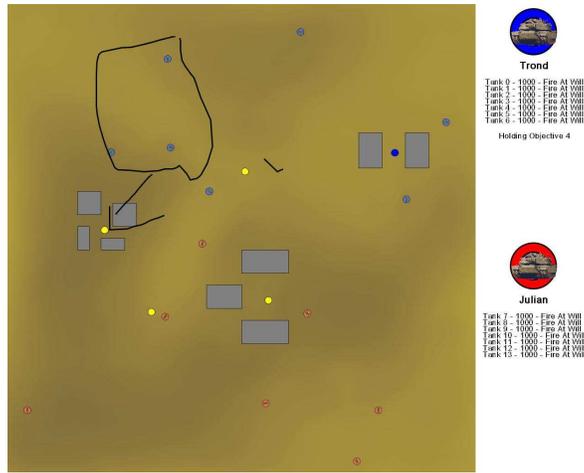


Figure 14: The Tablet PC interface after annotations

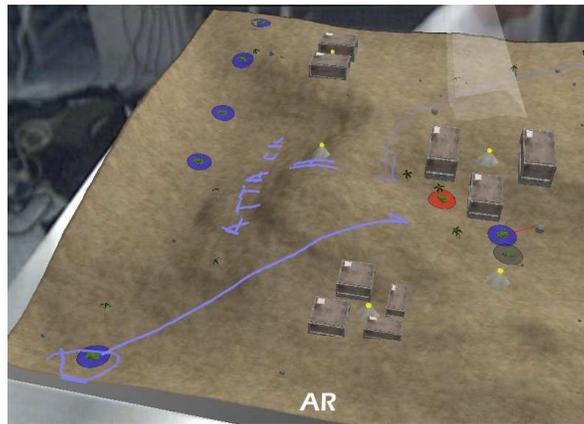


Figure 15: Annotations to the Augmented Reality interface

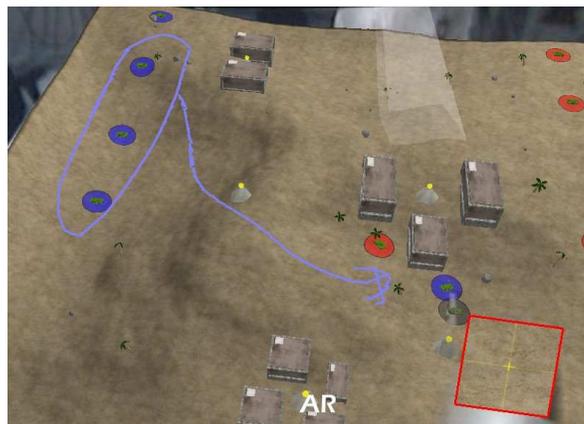


Figure 16: Annotations to the Augmented Reality interface

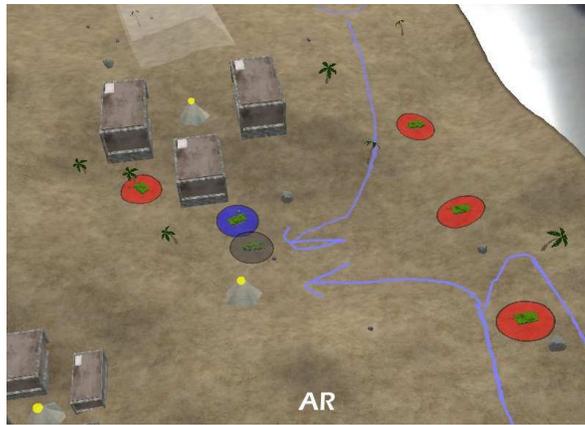


Figure 17: Annotations to the Augmented Reality interface

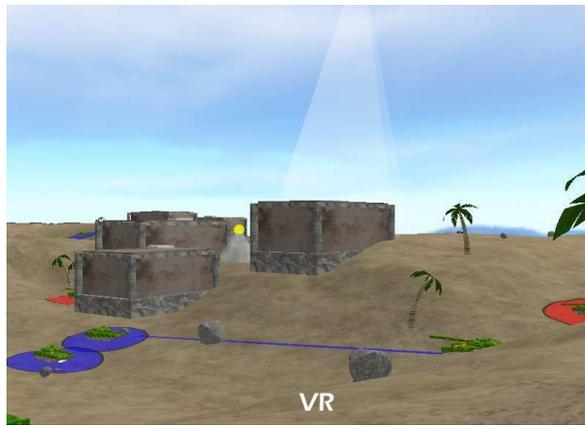


Figure 18: The Virtual Reality interface, showing tanks attacking (as shown by blue line)

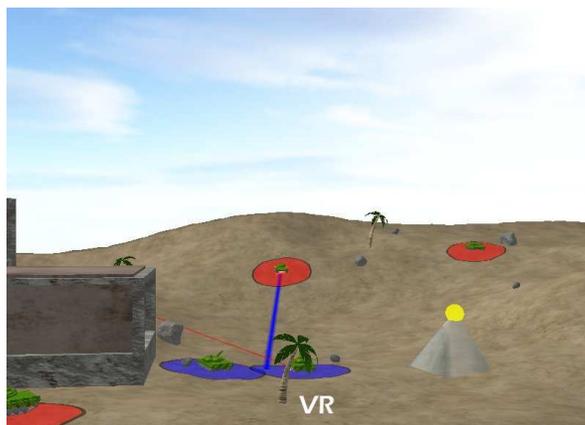


Figure 19: The Virtual Reality interface, showing tanks attacking (as shown by blue line)