Synthetic Air Assets in the Live Domain

Mr Matthew Franklin
QinetiQ
Farnborough, Hampshire, UK
MFRANKLIN@QinetiQ.com

ABSTRACT

The purpose of this paper is to provide an overview of research undertaken on behalf of the UK MoD, investigating the maturity and application of Augmented Reality (AR) capabilities. The research has been broadly split into 4 phases consisting of the stakeholder analysis, technology analysis, test-bed development, and exploitation phases. This paper will focus on the technology analysis, and test-bed development phases.

Analysis of AR technology has been undertaken to provide an evaluation of the current state of play and possible future trends that may lead to a lowering of the barriers to entry for AR. The results of this analysis for the display, tracking and software elements of an AR system are presented. In order to explore the technical issues associated with AR technology and convey the issues to the range of MoD stakeholders a test-bed AR system has been developed. Justification for the design of the testbed is presented, and details are given on the individual hardware and software components that are being used.

The research culminated in a series of experimental field trials followed by a demonstration in March 05. The focus for the demonstration was providing training to Forward Air Controllers in a live/virtual Close Air Support exercise. Experiences from the trials and feedback from the stakeholders is also presented.

ABOUT THE AUTHORS

Matthew Franklin is a senior research scientist within the Simulation and Training group at QinetiQ. He has been working within the group for five years, since graduating with a first class MEng in computer science from Imperial College London. Matthew has been involved in a number of research programmes covering diverse fields such as, tactical engagement simulation, embedded training, virtual dismounted infantry training, virtual and live interoperability, vehicle dynamics modelling, distance learning, and capability gap analysis.

Synthetic Air Assets in the Live Domain

Mr Matthew Franklin
QinetiQ
Farnborough, Hampshire, UK
MFRANKLIN@QinetiQ.com

Introduction

The "Augmented Reality to represent synthetic air assets in the live domain" research was begun in August 2003 under funding from the UK MOD's Directorate of Analysis Experimentation and Simulation (DAES). The primary aim of the research was to assess the maturity of the techniques and technologies for the stimulation of the live domain with appropriate representations of simulated air assets.

The impetus for the research arose from studies in which Live, Virtual and Constructive (LVC) techniques were used to supplement live training exercises with a synthetic wrap including virtual aircraft. The virtual aircraft could see and interact with the live participants as a result of live instrumentation data being used to create virtual entities. However, the live participants had no representation of the virtual world. Although geospatial information about the location of the virtual entities could be passed to the live domain via the Command Information System (CIS) and other data links, there was no way of presenting that information in the modality that the users were accustomed to i.e. visual, aural, or via the platform sensors. Fig 1 gives an overview of the interactions and capabilities of each of the LVC elements.

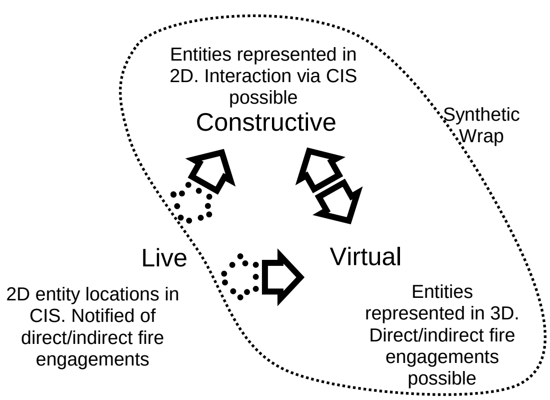

Figure 1. Connections and capabilities of an application of synthetic wrap

It was realised that to support more complex synthetic wrap architectures, needed for vertical training exercises such as those required for time sensitive target training, a method of stimulating the live domain was needed. For this reason research into Augmented Reality (AR) was begun. The term AR is now widely used to refer to the point along the Mixed Reality (MR) continuum, whereby the majority of the information presented to the user is from their live surroundings (Milgram, Takemura, Utsumi, Kishino, 1994). Strictly speaking the synthetic information should be presented in a three dimensional manner that correlates with the space around the user (Azuma, 1995). However, many AR systems simply provide two-dimensional information that is cued to the orientation and location of the user.

The main focus of the research has been on the analysis and development of individual stimulation techniques, whereby each participant in the live domain is provided with an AR system with which they can receive stimulation from the virtual domain. However, studies have also been conducted into the provision of environmental stimulation techniques whereby the environment is augmented and the live participants require minimal or no additional equipment.

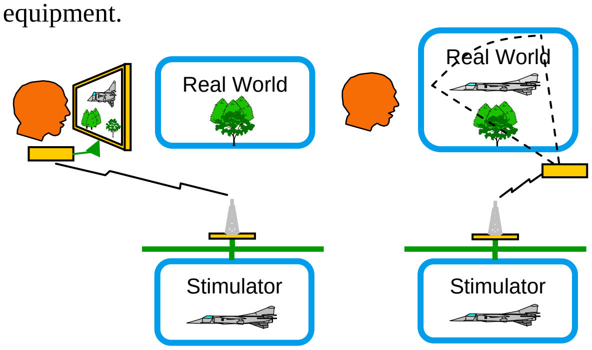

Figure 2: Individual and environment stimulation techniques

In addition to the technical studies the research has also investigated other potential application areas that could benefit from the use of AR. By involving stakeholders from each of these domains from the very beginning it has been possible to understand the scope of the required technical research such that its findings can benefit the broadest range of end users.

In order to fully explore the technical issues associated with AR technology and convey the issues to the range of potential stakeholders a test-bed AR system has been developed. The design of the test-bed was carefully chosen to achieve the following goals:

- Explore as many technical issues as possible.
- Be as adaptable and extensible as possible.
- Show stakeholders a practical application of AR.

Stakeholder and Domain Analysis

As indicated above the purpose of the stakeholder and domain analysis study was to gain an understanding of the domains in which AR systems might be used in the future and to understand the specific requirements of the stakeholders in those domains. Representatives from the land, sea, and air domains were approached to gain their thoughts on the possible applications of AR and the implications on the design of an AR system to meet their needs. More detailed information on the stakeholder analysis phase can be found in (Franklin, 2005).

Many applications of AR were found to be common across the range of stakeholder domains. The primary differences being the environmental constraints and the information that needed to be presented to the user. Many stakeholders were enthused about the possibilities of presenting stimulus for training applications as well as operational purposes. However, there were also reservations as to the cost effectiveness of the technology.

By engaging with the stakeholders it was possible to derive the following set of common requirements for a test-bed that could be used to further investigate the application of AR across the range of stakeholder domains as well as form the basis for exploring technical issues associated with the technology. The system should:

- provide naked eye stimulation as well as some form of magnified sight stimulation,
- support detection, recognition and identification up to a range of 5km,
- show that correlated sensor and/or CIS stimulation is possible,
- be man portable if possible, but also demonstrate applicability in mounted roles and an ability to be incorporated as part of a platform integrated system.

Underpinning Technology Analysis

The aim of the underpinning technology analysis study was to assess the current state of play of AR and assess the capabilities of the fundamental technology such that recommendations could be made as to its current and future feasibility. It is clear from existing AR systems, such as the BARS (Gabbard, et al., 2003), that current technology is capable of providing a functioning AR system. However, since AR is untested in many of the applications suggested by stakeholders the precise fidelity of the stimulation required is unknown. Therefore the implication on the capabilities of the underlying technology is reasonably unbounded. Parallels could be drawn with the fidelity of virtual simulators but there is currently little evidence for doing this.

The four primary areas that were addressed in the research were:

- visual stimulation systems,
- tracking systems,
- simulation/virtual data systems,
- integration software.

However, consideration was also given to the more practical elements of building an AR system such as aural stimulation, power supply, communications, and mobile computing.

Visual Stimulation Systems

It was found that there were many different types of display hardware that could be used for AR. These included video see through, optical combination, transmissive, and retinal projection systems, as illustrated in figure 3. Each display type has a range of advantages and disadvantages. For any one application several types of display may be applicable. However, no one display is best suited to all applications.

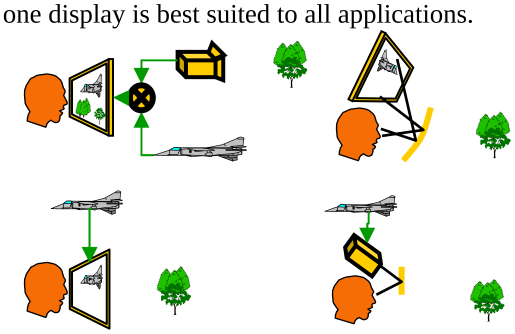

Figure 3: Video see-through, optical combination, transmissive, and retinal projection displays

The major factors that were considered as part of our research were; resolution, luminance, colour, power consumption, weight, lifetime and cost. For the type of applications considered by the stakeholders it was necessary for the system to match the resolution of the naked eye (~2 cycles/mrad. i.e. two transitions from black to white across the angle subtend by the eye see figure 4), be able to replicate the full colour spectrum (application and user dependent) and be used in bright sunlight (~1000 cd/m^2). The technology behind each of the display types (CRT, LCD, OLED, Laser) does not theoretically limit its ability to reach these fidelity levels but none are currently able to provide such stimulation with a reasonable lifetime. Non transmissive LCD displays are currently the closest to achieving the fidelity levels whilst remaining reasonably light and power efficient. However, future investment is most likely to be in OLED technology which currently cannot support transmissive displays.

Tracking Systems

Accurate tracking systems are the lynch pin of AR systems. They can be found in two main variants. The first variant is absolute, whereby the position of the AR system is tracked in real world co-ordinates (Inclinometers, Compasses, Global Positing System (GPS)). The second is relative, whereby the system is tracked relative to its direct environment (Radio Frequency (RF), Infra-Red (IR), Acoustic, Laser, Ultra Wide Band (UWB), Electromagnetic, Inertial, Visual).

Any AR system may use one or more tracking systems depending on its current environment. As with display technology there are a number of factors that need to be considered when selecting a tracking system. These include range, resolution, accuracy (position and orientation), latency, interference, weight, power consumption and cost. In order to accurately place an object such that its relative movement in relation to the ground is imperceptible to the naked eye the tracking system would need to have an update rate of at least 50Hz and an accuracy of approximately +/- 1/8th mrad (Calculated from the maximum error that can be tolerated by the human eye such that a white to black transition is perceived in a static location (Assuming that the normal visual acuity of the human eye is around 2 cycles per mrad). In actual fact the visual acuity of the human eye can be as high as 3.4 cycles per mrad. See figure 4). The higher the magnification the greater the accuracy needs to be. The positional accuracy of the tracking system depends on the range of the virtual objects from the user. Assuming the object is 5 km from the AR system the positional accuracy needs to be of the order of +/-0.625m (Calculated on a similar basis as above but subtended over the range. See figure 4), but for an object that is only 1m away the positional resolution needs to be of the order +/- 0.15 mm.

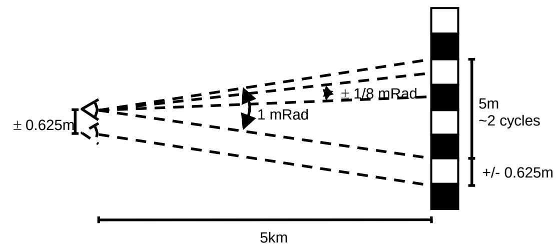

Figure 4: Calculating the position and orientation accuracy using eye limiting resolution

For the type of applications considered by the stakeholders only absolute tracking systems were appropriate because no infrastructure would be present on which to mount the relative tracking system. Currently only aircraft grade tracking systems come close to the kind of accuracy needed to deceive the naked eye, let alone a magnified sight. Unfortunately aircraft grade equipment does not tend to be suitably small, light, or cheap enough to make an AR system cost efficient. Currently the most cost effective approach is to use lower grade GPS and Inertial Navigation Systems (INS), combined together using predictive Kalman filtering. However, future investment is most likely to be in Micro Electro Mechnical (MEMs) technology which promises much smaller, lighter and more accurate tracking capabilities.

Simulation/Virtual Data Systems

Fundamental to the calculation of the placement and occlusion of virtual information in an AR system is a model of the real world. This includes the geometry of objects (static and dynamic), weather effects, lighting, and sound. Real time generation of 3D geometry is feasible by using stereo graphic imagery but the range of such systems is restricted by the separation and resolution of the cameras used. Therefore such an approach is currently limited to AR applications where the distance between the user and virtual information is less than 30m. Alternatively geometry can be gathered from systems such as UWB radar or Light Detection and Ranging (LIDAR). However, such systems tend to be bulky and power hungry and therefore also not suitable for most AR applications. Current AR systems usually rely on pre-defined geometry which can be modified in real time based on the known location of objects or markers (as is used in AR Toolkit (AR Toolkit website, 2005)).

For the applications considered by the stakeholders the majority of occlusion is performed by the terrain and the prevailing weather conditions. Terrain elevation data is readily available for many areas of the world in Digital Terrain Elevation Data (DTED) Level 2 resolution (around 25 m for the UK). Although this is less than the 1m resolution indicated above it seems likely that 1m resolution elevation data will become readily available in the near future due to commercial investment. Such an approach does not adequately deal with dynamic objects in the real world, such as trees or vehicles, but does provide a baseline for manual modifications or automated modification using techniques such as those adopted by the AR Toolkit.

Monitoring of environmental effects such as weather (fog, cloud, wind), lighting (natural and artificial) and sound remains difficult to achieve but reasonably easy to recreate in aural and visual simulations. Therefore at this time manual calibration of the AR system is the most effective means of achieving the required effect. For both visual and aural simulation multi-pass ray-tracing algorithms could be used to produce the most realistic results. However, current real-time solutions rely on approximating this process using line of sight calculations.

Integration Software

The integration software is also seen as a key component of an AR system. It was found that many applications of AR shared the same unpinning sub-components. The only difference being the configuration in which the sub-components were used. For example the display and tracking systems of a dismounted AR system may be significantly different from one mounted in a vehicle, but the information to be displayed to the user and the geometry of the real world environment are identical. The key elements that were identified as being fundamental to AR software design were:

- ability to synchronise and correlate the tracking information with the synthetic information to be presented to the user,
- ability to support the range of tracking, display, and sound hardware components that could be used across the range of AR applications,
- ability to support dynamic generation and calibration of the real world model,
- ability to present the various types of synthetic information to the user,
- support configurable update rates depending on application and fidelity.

Test-bed Design

As discussed above a test-bed was required to fully explore the technical issues associated with AR technology and convey the issues to a range of potential stakeholders. An underlying principle in the design of the test-bed was to avoid the development of a prototype system tailored to the needs of one particular stakeholder. The test-bed needed to be designed in such a way that allowed it to be used in the future to assess the functionality and fidelity required from AR systems across the range of stakeholder domains.

Comparing the needs of the stakeholders against the feasible technical approaches it was decided to base the test-bed around a pair of modified virtual binoculars. Information to be displayed in the binoculars would be provided by a remote server, as would be expected within a synthetic wrap environment on a training area or a range, or within a Command Information System (CIS) environment during operations. For the purposes of a demonstration these binoculars would be used by a Forward Air Controller (FAC) as part of a Close Air Support (CAS) scenario in which a virtual aircraft, controlled by a man in the loop simulator, would be talked onto live and virtual targets in the real world. The remote server in the demonstration would be represented by a fixed wing flight simulator backed up by a Semi Automated Forces (SAF) client (see figure 5).

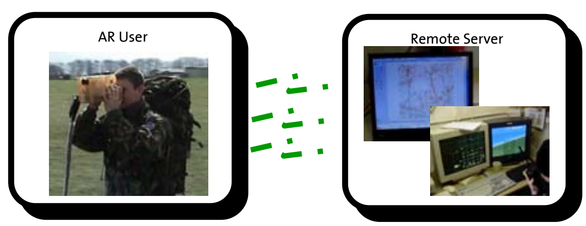

Figure 5. The Augmented Reality testbed components

Visual Stimulation System

Binoculars provided a happy medium between platform-mounted optics and dismounted naked eye stimulation and was a common tool used across all the stakeholder domains. Whilst binoculars do not provide the field of view of the naked eye they make it practical to use a video see-through approach, whereby the live image is captured by a camera and presented to the user via a pair of displays. This approach had four major advantages, the first being the ability to use full colour SXGA (1280x1024) resolution displays which approach the resolution of the human eye when used in the field of view of the eye piece of a pair of binoculars. The second advantage is that it avoided problems associated with generating images with sufficient luminance to be used in daylight. This is due to the combination of capturing the images via a camera, which limits the incoming light levels, and the cowl of the eyepieces blocking out much of the background light thus allowing the displays to be seen clearly. The third advantage is the ability to change the lens of the camera to allow the experimentation with different levels of magnification. Finally the casing of the binoculars provides a stable platform for mounting sensors and speakers. For the purpose of the testbed the displays were taken from a pair of NVIS virtual binoculars. For the applications indicated by stakeholders the virtual stimulation would be placed at ranges greater than 50m therefore only one camera was required by the test-bed since stereographic stimulation will provide limited benefit. The camera used in the test-bed was provided by JAI and supplied 1392x1040 resolution images at 24fps. Although this was less than the 50 fps indicated above it was the maximum that could be achieved at the time within the size and weight constraints of a pair of binoculars.

Tracking System

Analysis of the options for tracking the position and orientation of the binoculars indicated that no off the shelf capability could provide the level of accuracy that was needed to perform the tasks required by the stakeholders. For this reason a bespoke capability was developed based on an LN200 inertial sensor combined with an off-the-shelf DGPS system. Proprietary Kalman filtering techniques are used to provide orientation accuracies of up to +/- 1 mrad and positional accuracy of approximately +/- 0.1 m. Average accuracies are closer to +/- 10 mrad which enables the system to place objects at +/- 50m at a range of 5km. Processing for this level of accuracy required the use of a separate PC104 based computer.

Simulation/Virtual Data Systems

Within the constraints of a binocular system it was not possible to use sensors to build dynamic databases of the users environment. By leveraging tools such as Terrex Terra Vista and Multigen-Paradigm Creator it is possible to dynamically load terrain databases from a wide range of sources including among others, DTED, Digital Feature Analysis Data (DFAD), Vector Map (VMAP), Light Detection and Ranging (LIDAR) data, and Millimetric Wave Radar data. This provides sufficient accuracy for most applications of AR.

To generate the virtual imagery a COTS laptop was used. A smaller portable computer could have been used. However, at the time no portable system had the sufficient physical interfaces or processing performance to produce SXGA resolution images in real time. Size was not so much of an issue as it was apparent that the user would need to wear a backpack to carry the tracking computer, batteries and interface boxes for the camera and displays. As a general rule, size and weight of the AR system was compromised to support flexibility and functionality, with the understanding that in time the size and weigh of the components would drastically decrease.

For the purposes of the test-bed, communication between the binocular system and the remote server is performed by data radio modems or WIFI. However, this communication link could be replaced by whatever data dissemination network is available in a particular environment. Distributed Interactive Simulation (DIS) was chosen as the protocol of choice to support the flexibility needed for trials using the AR system as well as to take advantage of the plethora of DIS interfaces and tools currently available. However, the type of information passed to the AR system is largely irrelevant as long as the image generator understands how to display the information to the user. The advantage of using DIS over a CIS based protocol is that it supported a large range of visual effects such as entity status, explosions and weather effects. The SAF used within the test-bed provided a simple yet powerful ability to dynamically adjust the weather effects to be visualised in the AR system.

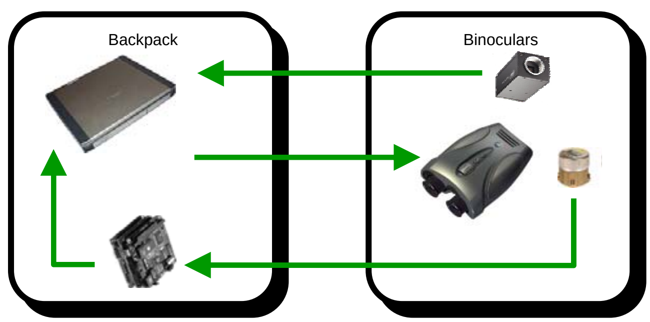

Figure 6. The Augmented Reality User system components

Integration Software

Rather than use an external hardware genlock to combine the virtual and the live image it was decided that greater flexibility could be afforded by performing the blending in software. This approach allows greater control over the synchronisation of the live and virtual information, and also allows the introduction of image manipulation capabilities needed to adapt the system to different types of lens'/optics and displays. The image manipulation software used within the testbed is able to arbitrarily morph input and output images as well as adjust colour levels to suit the application. However, if required, the software blending component of the system can be removed to support hardware blending applications.

Unfortunately, at the time of designing the AR system the only readily available AR software, the AR Toolkit, could not provide sufficient flexibility to support all of the requirements outlined in the analysis above. The software used on the laptop is based on the Synthetic Environment Application Layer (SEAL) (Smith, Kelly, 2001). This software is derived from the Synthetic Environment community and allows the rapid configuration of SE type applications. It operates in a HLA type manner whereby objects publish and subscribe their attributes but the federates can be tightly coupled to support real time applications such as that required by the testbed. It provided a flexible and controllable basis for exploring different configurations and synchronisation mechanisms of the AR architecture. A number of bespoke libraries were built in order to capture images, manipulate them and combine them with the virtual information. In particular additional capabilities were added to the image generator (based on OpenSceneGraph) in order to support mixed reality functionality. The remainder of the libraries used in the software remain common with existing SE type applications. In fact the core system is based on a stealth viewer that simply attaches to the tracking interface component rather than an entity. Being based on a standard stealth viewer means that the system supports all the generic 2D and 3D visualisation functionality that can be expected within a stealth. The type of information that can be displayed and its format is therefore only bounded by the functionality of the stealth and the limits of the graphics hardware.

The current AR software is illustrated in figure 7. However, the modules can be configured in any order and in any quantity. Through the SEAL architecture it is possible to configure the connections between each of the modules. Thus, the camera interface module can be removed in its entirety, i.e. for a transparent display AR system, without the need to reconfigure the rest of the software. In the configuration in figure 7 the input image morph module is used to flatten the camera image to remove any lens distortion effects. In parallel the image generator builds a 3D image by rendering to an off-screen buffer and the 3D sound module generates the necessary sounds based on the location returned by the tracker module. The flattened camera image is then combined with the generated 3D image but before being displayed it is distorted once again to suit the type of display on which it is to be presented. The entire chain is processor intensive thus elements, such as the image morphing stages can be removed to enable faster responses on slower processors or in AR systems that do not require them. Another key module that can be removed is the software genlock stage as the camera image and live image can be combined in the 3D image generation stage. However, this approach provides slightly less flexibility over the combination process than in using a software genlock. It also introduces a potentially higher latency as the image generation process cannot be performed in parallel with the image morphing process.

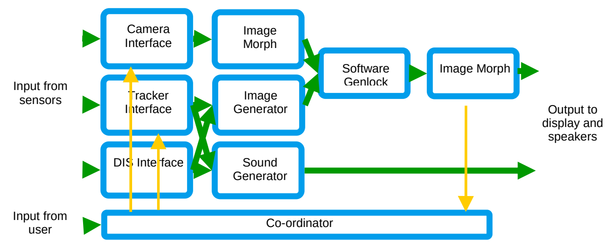

Figure 7: AR test-bed software architecture

Temporal and spatial co-ordination of the system is also equally flexible such that output of the system can be tied to the latest tracking information, latest input image, or even the interpolated position at the end of the chain. This functionality also provides control over the update rate and fidelity required from each of the components thus allowing the system to be configured to suit the end user application.

Trials and stakeholder feedback

The most problematic element of the system was configuring the camera interface to provide a consistent image, against which a virtual image could be generated. Despite automated software to control the exposure and colour balance, the resulting image could not be guaranteed to meet its desired limits. To cope with this variability the image generation software was modified to cope with per frame brightness, contrast and colour balance variances.

Generation of a realistic occlusion database also proved to be problematic. Many of the sources of elevation and feature data proved to be inconsistent. This was exacerbated by errors in the co-ordinate conversion and database projection algorithms. The later problems could be easily overcome by either modifying the source data or modifying the co-ordinate conversion process. The former proved more difficult and required manual intervention to modify the source data. The AR system itself proved useful for this process as it could be used to visually verify the data correlation and direct modifications to the source data.

It was found that rendering to an off screen buffer then using a software genlock took longer than sending the camera image to the image generator and combining the two during the image generation process. This is probably due to the image generator not being designed to render to an off-screen buffer thus operating much slower than would normally be expected.

Sound was a critical indicator of the location of virtual entities. Without the sound component many operators of the AR equipment struggled to detect virtual air assets. However, comparison with live air assets indicated that the system provided comparable levels of visual stimulation to the naked eye.

Weather effects had a major impact on the realism of the virtual objects. In particular shadows of clouds would cause the virtual entity to stand out against the darker background. Rain and fog was found to be inconsistent across the field of view of the system also leading to the unrealistic representation of virtual entities.

Another factor that significantly effected the realism of the virtual entities in a live environment was the fidelity of the models, textures, and behaviours taken from the virtual environment domain. These can be easily improved using existing tools but will require significant improvement if they are to be believable in the real world.

The size and weight of the test-bed was a concern of many stakeholders. However, it should be noted that many of the sub-components of the test-bed are likely to rapidly reduce in size as a result of commercial influences in the sensors, computing, and displays markets.

Stakeholders had reservations as to the fair fight principle of blending live a virtual in a training environment. In particularly the FAC in the scenario chosen for the demonstration had more information about the environment than was provided to the flight simulator. Key factors such as cloud positions and wind effects were not faithfully translated into the virtual domain.

Finally, many AR training applications would require all systems stimulation if they are to be effective. I.e. all participants in an exercise will require the same stimulus through all the systems that they would expect stimulation including sensors, optics, and naked eye. For example, the commander of a vehicle may spot a target that will need to be engaged by the gunner. Unless the two crew members visualise the same target in the same location, correlation issues will arise. Likewise, anyone else in the exercise would need to be able to visualise the effect of the engagement on the target.

Conclusions

In March 2005 the test-bed system was successfully used by an FAC to conduct a CAS exercise. Whilst the original plan was to only provide synthetic air asset stimulus the system was able to support the generation of ground based assets and effects (see figure 8). During the exercise the FAC was able to direct a virtual aircraft onto live and virtual targets placed on the ground at a distance of up to 3 km.

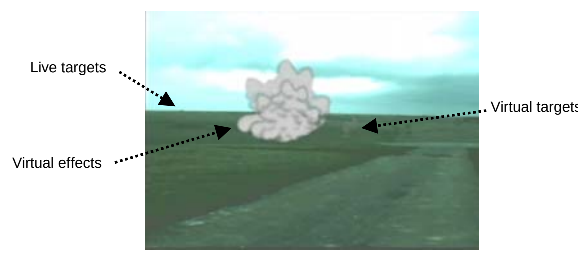

Figure 8. Example image captured from the AR binoculars

With current technology, applications of AR technology need to be carefully considered to ensure that all participants are receiving the necessary and correct information in order to perform their task. The inability to dynamically sense the real world and update the synthetic environment accordingly drastically reduces the realism of virtual entities in the real world and reduces the realism of the exercise for any virtual participants. However, the use of AR techniques for presenting operationally relevant information does not require such high fidelity occlusion models and so is probably a more practical use of the technology at this stage.

It is clear from the research that tracking technology is the main technical hurdle for the building of AR systems. However, many applications of AR do not require the level of accuracy needed to fool the human eye and MEMS technology is quickly maturing. Stimulating the naked human eye remains challenging but technologies such as laser retinal scanners may hold the key to providing such stimulation in the near future. Software for AR is also quickly maturing but there is a need for a standardised approach to the integration of mixed reality capabilities into image generators. The impetus for the inclusion of video into 3D rendering systems may come from other sources, such as the entertainment industry, and it is expected to lower the barriers to entry for building practical AR systems. The only area that remains to be adequately addressed is the rapid development of real world models. Techniques such as extraction from stereo imagery are likely to steadily improve as the resolution of image capture devices improves, but these techniques will remain secondary to predefined databases for the foreseeable future.

Concerns over the size weight and cost of AR remain but are likely to rapidly dissipate as a result of commercial development in the underpinning technology. In the mean time the AR testbed developed under this research will continue to be used to investigate the fidelity and functionality requirements of stakeholders across a broad range of domains. By using the testbed as a high fidelity exemplar it is hoped that it will be possible to design future AR systems to better suit the needs of the end user application and thus be much smaller, lighter and fit for purpose.

ACKNOWLEDGEMENTS

Thanks go to to DAES for sponsoring this work, and to the JASO and JFACTSU for their support.

REFERENCES

Milgram P., Takemura H., Utsumi A., Kishino F. (1994). *AR: A class of displays of the Mixed Reality Continuum*. Telepresence and Telemanipulation, 1994.

Azuma R. (1995). *A Survey of AR. Presence: Teleoperators and Virtual Environments.* ACM SIGGRAPH 1995.

Franklin M. (2005), *Augmented Reality for live training*, ITEC 2005

Gabbard J., Brown D., Livingston M., Baillot Y., Rosenlum L,, Julier S., Swan J,, Hix D., *An Augmented Reality System for Military Operations in Urban Terrrain*, IITSEC 2003

AR Toolkit (2005), *AR Toolkit*. Retrieved Jan 2005, from http://www.hitl.washington.edu/artoolkit/

Smith N., Kelly M., (2001). *Using the Synthetic Envrionment Application Layer (SEAL) for the Development of a Low Cost Human-in-the-Loop Land Platform Training Simulation*. ITEC 2001.