Visualising and Interacting with a CAVE using Real-World Sensor Data

Anonymous for review

ABSTRACT

A CAVE (Cave Automatic Virtual Environment) is a 3D interactive environment that enables a user to be fully immersed in a virtual world and offers a unique way to visualise and interact with digital information. In this paper we build a foundation of CAVE interaction design by characterising generic affordances of such an environment and enumerating currently conceivable/implementable interaction mechanisms. In particular, we focus on how different aspects of the CAVE affordances relate to virtual worlds that are generated for the purpose of visualising and interacting with real-world sensor data. In support of this we present a case study which explores how the CAVE can be used to visualise and better understand data from 16 residential apartments where we have been collecting daily home usage information for the past year. We summarise our on-going work by demonstrating how the unique characteristics of the CAVE can be used to visualise this large quantity of heterogeneous data, which provides more novel possibilities than traditional data visualisation and interaction methods.

Categories and Subject Descriptors

H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*Artificial, augmented, and virtual realities*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Input devices and strategies*

General Terms

Design, Human Factors, Experimentation

Keywords

CAVE, Ambient Assisted Living, Interaction Design, Visualisation, Sensor Data, Virtual Reality

1. INTRODUCTION

As part of a research programme investigating novel Ambient Assisted Living (AAL) technologies, we are working directly with the elderly residents of the Great Northern Haven (GNH) apartments in Dundalk, Co. Louth, which is made up of 16 high-tech purpose-built two bedroom apartments. The building contains thousands of sensors and actuators as well as other interactive technologies. For example, sensors detect when a resident opens a window, starts a kettle, goes to a bedroom, etc. Our objective is to use the real-world data generated from these devices to develop new ways to promote greater levels of independence and well-being for older people. The data is also being used to develop behavioural models to promote early interventions and help decrease the rate of cognitive and physical decline as well as, ultimately, better enabling older people to monitor and manage their own health. To date, we have collected over 83 million records worth of daily usage information from the 16 apartments.

As part of this research we use a CAVE, an immersive virtual reality system, to develop and test new technological, environmental and architectural concepts and designs. The CAVE is a cube shaped room which offers a multi-person, multi-screen, high-resolution 3D video and audio interactive environment. The users enter the CAVE wearing 3D glasses and as they move and interact within the display boundaries, the correct 3D perspective is displayed in real-time to achieve a fully immersive experience. Using the CAVE we are able to virtually recreate the GNH apartments to-scale and in full 3D. The users can then navigate around the environment as if being there using different control methods and devices. Not only does this allow us to investigate different scenarios and sensors before they are implemented in the real world, it also allows us to feed the sensor data from the apartments directly into the CAVE. As we navigate around the environment we can visualise both the sensors and the data in a format that is most suitable to our needs.

This paper is organised as follows. In the next section, we briefly review ongoing CAVE related research. In Section 3, we describe a set of essential components that make up a typical CAVE setting, leading to Section 4 that lists methods and techniques to enhance the usability of the CAVE environment. Section 5 then goes on to describe various ways in which a CAVE user can interact with the environment. In Section 6 we summarise a case study where these design considerations were used to explore different ways to visualise and interact with a large amount of GNH sensor data. Finally Section 7 concludes the paper with our plans and future work.

2. RELATED WORK

Since the creation of the CAVE in the early 90's by Carolina Cruz-Neira, Thomas A. DeFanti and Daniel J. Sandin [7], the Electronic Visualisation Laboratory based at the University of Illinois Chicago has remained at the forefront of CAVE and virtual reality research. Intriguingly, in spite of all the advances in technology over the last twenty years, it is surprising how little CAVEs have changed since they were originally shown to the world at SIGGRAPH'92. Other laboratories such as the 3DI lab at Virginia Tech University under the direction of Doug A. Bowman are carrying out exciting research in the area of virtual environments and particularly in relation to how we interact with them [2, 1, 3]. Their research into the area of 3D user interfaces is particularly relevant not just for CAVE environments but also the re-emergence of interactive 3DTV and other 3D enabled devices today. Increasingly important in this context is Microsoft's Kinect[1]. While Microsoft did not set out to develop a technology for CAVE interactions, their Kinect technology is creating fascinating new ways to interact with virtual worlds. Similarly, the emergence of Google's liquid galaxy project[2] is likely to have an impact on CAVE-based research. The Virtual Reality Applications Center at Iowa State University is another laboratory which is highly influential in the field of virtual reality and immersive environments. They have been prolific in both their research and ability to create spin-off companies related to the VR field, one of which, Mechdyne[3], is a leader in the set-up and development of CAVEs such as ours. For all the research into CAVE environments, however, there is significant scope for more research into the area of visualisation strategies and interaction techniques within a CAVE [4]. Exhibiting very different affordances to conventional desktop interaction or more recent mobile interaction, the currently available design knowledge base for the CAVE is limited. In this context it is clear that understanding the special characteristics and affordances of CAVE interaction will be the key to successfully exploiting the benefits of such a platform.

3. CHARACTERISING THE CAVE

The typical CAVE build is made up of 4 screens laid out similar to Figure 1 with each screen measuring between 2 to 3 metres in both width and height [6]. Each of the 3 vertical walls is a projectable screen with a projector sitting some distance behind it (perhaps up to 10 metres or more). Placing the projectors behind the screens prevents shadows being cast by users or other objects onto the image being displayed. The use of rear projectors means that CAVE setups generally require significant amounts of space in which to be housed, usually many times greater than the space of the user interaction area itself. These distances, however, can be reduced through the use of mirrors. The floor screen is made up of a down projection screen which can be walked upon by the user. Due to the distance required to project onto the screens, the floor projector is generally attached to the CAVE frame and uses a mirror to project down onto the floor. This is the typical physical structure of a CAVE, however, CAVE builds can be made up of more or less screens up to a maximum of six screens, such as the one implemented by the Virtual Reality Applications Center (VRAC)[4] at Iowa State University.

[1]http://www.xbox.com/en-IE/kinect
[2]http://code.google.com/p/liquid-galaxy/
[3]http://www.mechdyne.com

Figure 1: The CAVE in operation

Of course, a key aspect of the CAVE is the ability to display full stereographic 3D images. This is typically provided through the use of stereoscopic 120Hz projectors, 3D shutter glasses worn by users and Infrared (IR) emitters to synchronize the glasses and screens. However, a similar effect can be achieved through the use of passive projectors and polarized glasses. Related to this, an important factor is ensuring that the CAVE is configured to display images across the screens as if on one large screen. This can be achieved through the use of a range of different software applications which are specifically developed for CAVE environments. These applications generally look after all aspects of the screen management as well as handling the virtual world object models and interaction, object tracking integration, controller integration and management of the different computing hardware involved.

Another key component of a CAVE is the use of head tracking in order to adjust the image displayed across the screens based on the position of the user's head in the CAVE. To demonstrate this, imagine standing in the centre of the CAVE looking at an image of a wall that is shoulder height on the screen in front of you. Behind the wall some distance away is the image of a tree. In the real world if we hunker down behind a wall, the trees behind it will be hidden from view. In the CAVE the head tracking replicates this scenario so as the user hunkers down in the CAVE the view will descend lower to the ground and the tree will no longer be visible behind the wall. In order to achieve this, our CAVE uses 10 IR cameras placed around the top edge of the 3 vertical screens. These are able to detect the exact position of 3 reflective balls that are placed on the user's 3D glasses providing the x, y and z co-ordinates of the user's head within the CAVE space. This is a widely used object tracking setup, however, technologies such as Microsoft's Xbox Kinect potentially provide an interesting and much lower cost alternative to traditional tracking implementations.

All CAVEs require a method of interacting with their virtual worlds. This is something we will address in greater detail

[4]http://www.vrac.iastate.edu/

in Section 5. However, to date, joystick or gamepad style controllers similar to those used by many popular videogame consoles as well as interactive gloves are the most common devices for interaction.

Finally, all CAVEs require the necessary processing power to produce smooth and graphically detailed virtual worlds and enable the users to interact with them in a usable and meaningful way. There is no standard implementation, but our CAVE uses 5 Xeon quad core workstations each possessing 16GB of RAM and an NVIDIA Quadro video card, as well as a standard desktop PC to handle the object tracking.

4. ENHANCING IMMERSION

The objective of a CAVE environment is to provide as *immersive* an experience as possible into a virtual world [11]. In other words it is about trying to create an environment that absorbs the user so that they become unaware of the physical CAVE and its surrounds and fully experience and "believe" the three dimensional world that is generated for them. This is mainly achieved by the typical CAVE set-up as described in Section 3, but there are other key elements that can be used to further enhance the immersive experience.

Photorealism. In the context of the paragraph above, one could be forgiven for assuming that the level of immersion generated in a CAVE is all about the level of photorealism [8] but in reality it is more to do with the sense of depth and space that the virtual world conveys. A cartoon or abstract world can be just as immersive as one that is photorealistic. Techniques such as ambient lighting or global illumination [9] within a 3D world can have a significant impact on the sense of depth and space conveyed to the user.

Reduced Ambient Light. Not to be confused with the use of ambient lighting within a 3D scene, the ambient light we are referring to here is any light that is not generated by light thrown on to the screens by the projectors. The more ambient light there is the more the user notices the physical CAVE surrounds and the less immersive the experience. Ideally CAVEs should be set-up in windowless rooms and where any artificial lighting can be fully controlled.

3D Projectors. Projectors that are capable of displaying crisp and clear stereo 3D images are one of the most important aspects in enhancing the immersive effect of the CAVE. It is the 3D projections that provide the overriding illusion of depth in the virtual world. The resolution provided by the projectors can also help to support that effect with higher resolutions providing more defined images.

3D Glasses. Flicker-free 3D shutter glasses or good quality polarized glasses help to ensure the illusion of depth is effective.

Screens and Seams. The number of screens in a CAVE environment impacts the level of immersion felt by the user. A standard 4 screen CAVE provides a high level of immersion, particularly when a user is facing the centre screen. This ensures that the user's entire horizontal field of view (which is generally anywhere between 160 to 180 degrees) will be fully covered by the CAVE's 3 vertical screens. It is only if the user positions their head a significant angle facing away from the centre screen or up towards the ceiling that there is a break in the screen coverage. CAVEs with 6 screens can provide complete coverage.

Another important aspect to the level of immersion is the way in which CAVE screens are linked together. Any material other than the projection screen itself will interfere with the three dimensional illusion generated by the CAVE, particularly in the low ambient light conditions recommended for CAVE use. Many CAVE environments do not use visible seams and instead join the seams behind the screens outside the path of a projector's beam of light.

Audio Set-up. Depending on the type of 3D world the user is interacting with, the type of audio set-up may be of greater or lesser importance to the user. Assume we have a 3D world which is designed to assess a user's ability to identify and locate traffic sounds. In this scenario we may present the user with an outdoor scene where we have traffic arriving from many different points in the 3D world. Having multiple speakers placed at different points behind the CAVE screen would allow for the generation of traffic noise from multiple points and help users to correctly identify which direction a vehicle is coming from before they see it just as in the real world.

Hardware and Software Specifications. Ensuring that the CAVE hardware and software are up to the job of handling the 3D world to be experienced is essential. A 3D world that jitters and struggles to keep up with the user's movements will not provide for a usable immersive experience.

Object Tracking. Object tracking can add to the immersive experience in different ways. Most often object tracking is used to determine the positional co-ordinates of a user's head within the CAVE. This data is then used to adjust the image that the user sees on the screen as outlined in the earlier example of hunkering down in front of a wall with a tree behind it.

CAVE Floor Area. One of the limitations of a CAVE is the fact that its usable area is generally only $2m^2$ to $3m^2$. This means that a user's movements are significantly restricted. Ironically, due to the highly effective immersive effect of most CAVEs, this can sometimes lead to users wandering towards the screens and becoming disorientated.

Interaction Methods. In addition to the key elements listed in this section, the mechanisms as to how a user interacts with the CAVEs virtual environments impacts on the level of immersion the user experiences which we detail in the next section.

5. INTERACTION IN THE CAVE

The methods we use to interact with CAVE environments is an extremely important and topical issue for many CAVE researchers [12, 10]. It is also one of the most exciting areas to work on particularly as novel but affordable interaction control mechanisms are starting to appear on the market. Study into the interaction strategies for the CAVE is important because it can fundamentally change how a user engages with the virtual world they are presented with and topical because new advances in technology such as Microsoft's Xbox Kinect and Google's Liquid Galaxy project open up a whole new range of possibilities. In this section we enumerate currently available or conceivable interaction methods in the CAVE and their characteristics. Before doing so, it is worth highlighting some key questions that should be answered to help assess which methods of interaction would be most suitable for a particular scenario:

- What is the purpose of the virtual world?
- Who is that world created for?
- What type of world is being generated?
- What level of interaction is required?

So, for example, if the purpose of the virtual world is to assess a user's natural movements within a scene then perhaps we need to ensure that the interaction method is a very natural one; if the world is specific to older people perhaps the necessary interaction method would be different than that for a young adult; if the world being generated is an abstract environment such as a data model then perhaps we should provide the user with a more sophisticated navigation control for 6 Degrees of Freedom (DoF); if we want to allow the user to manipulate objects in a scene then perhaps this will require a specific interaction method that provides features to point, select, drag and stretch in some way. Depending on the envisaged usage scenarios the ideal type of interaction strategies for the CAVE will be different. So what are the interaction methods options available to CAVE users? The following are some of the main options available but are by no means exclusive:

Gamepad/Playstation-style controller. Gamepad controllers are already a familiar device for video gamers, typically the user holds the device with both hands and a number of push buttons and thumbsticks are conveniently accessible under thumbs and forefingers. While these controllers tend to prove very effective for those accustomed to them (e.g. hard-core gamers), they can appear overly cluttered and non-intuitive to first-time users. The gamepad controller being used in our CAVE provides 2DoF (backwards, forwards, turning left and turning right) through the use of one of its analog thumbsticks. It does however have the capability of providing more DoF through the use of an additional thumbstick and digital directional pad. Some buttons are also used to reset views and toggle menus on and off.

3D mouse. A 3D Mouse is similar to a standard mouse for a desktop PC except that it is held freely in the hand and does not need to be placed on a surface. A 3D mouse uses a trackball and buttons and is typically compatible with tracking software giving it the possibility of 6DoF simply by turning, twisting and tilting the device. The relatively simple design of a 3D mouse may prove more intuitive to some users than a gamepad controller and is particularly suitable for interacting with abstract worlds where significant degrees of freedom are required.

Fixed position joystick. Should we require that the user does not move around the CAVE space then an integrated joystick and stand provides an excellent method for interacting with a virtual world, particularly if it is a built environment scene where only a few degrees of freedom are required. Using a simple joystick with perhaps one or two trigger buttons provides an easy-to-use and intuitive style of control. Its fixed position helps users to focus on the virtual environment and eliminates the need to hold and carry a cumbersome controller. However, this may limit its usefulness in many situations and the fact that the joystick is placed on a stand will also block some of the user's view of the floor screen and may cast additional shadows.

Dance pad/floor sensor. A dance pad is a flat electronic controller that is placed on the ground and provides the user with interaction through the placement of feet on specific areas of the pad. Dance pads are typically used at home and in games arcades for dance based videogames and can be easily replicated through the use of floor sensors attached to a microcontroller. The advantage of a dance pad in the CAVE is that it does not require the use of a hand-held or on-body controller and it has the potential for a simple and intuitive interaction. However, interacting using foot placement alone may be restrictive and tends to force the user to concentrate on their foot movements rather than the virtual environment.

Interactive glove. Interactive gloves often use a finite number of sensors at the finger joints and/or finger tips which will detect the bending or stretching of joints as well as pressure on finger tips and send the information back to the system. Most interactive gloves for CAVEs also use a tracking mechanism in order to enable the user to simulate their hand position in a virtual world. Using hand gestures the user can then navigate through the world and manipulate and interact with objects. Anyone who has seen the movie Minority Report[5] can realise the potential. The reality, unfortunately, is that most current glove controllers are not considered to have the level of functionality, accuracy and sensitivity required for commanding generic functions such as pointing and selecting a button or manipulating a virtual object. In spite of this, it is worth noting that interactive gloves are still widely used in CAVE environments.

Wand controllers. Wand controller is a somewhat generic term to define a particular style of controller device (of which there are many), here referring to devices such as the Playstation Move, Nintendo Wii-mote and Nintendo Wii Nunchuck. While each of these is quite different they also have key similarities: (i) they are all held or gripped in a similar way in one hand, (ii) they all contain accelerometers in order to determine the movement and rotation of the devices, and (iii) they all feature trigger buttons. While different devices have different additional features, they all have the potential to be used in a CAVE. We are able to integrate a Wii Nunchuck into our CAVE through the use of an Arduino[6] microcontroller. The advantage of the Wii Nunchuck is that it provides a very simple alternative to a

[5]Minority Report (Steven Spielberg, 2002), DreamWorks and 20th Century Fox

[6]Arduino is an open-source electronics prototyping platform based on flexible, easy-to-use hardware and software. It is intended for artists, designers, hobbyists, and anyone interested in creating interactive objects or environments.

standard gamepad with just one analog thumbstick and 2 trigger buttons. Yet it allows for the possibility of 6DoF meaning that it has the ability to be useful in both built environment and abstract data interactions (see Section 6.1 for more detail).

Voice control. Using voice commands in order to be able to control specific aspects of a virtual world adds a natural and intuitive interaction modality suitable for CAVE environments. Through the use of voice recognition software which can recognise words and sentences and then convert them into actions we have the possibility to better integrate users into the CAVE. Although voice control is generally not considered a primary interaction method it provides huge potential when combined with other more natural interaction methods such as controller-free gesture recognition or haptic/aural feedback.

Phone and tablet apps integration. In certain situations the use of a smartphone or tablet device could be a useful way to interact with a CAVE. Using custom-built apps, the device can be used to interact with a built environment situation in real-time in the CAVE. So, for example, we could use the device to simulate a home automation control panel (e.g. turning lights on or off, activating an alarm or opening and closing blinds). In other words if we want people to be able to interact with their homes with these devices in the real-world, then we can also do the same in the CAVE. This would give us the ability to easily simulate and test potential real-world scenarios. In terms of interacting with abstract data visualisation the device could also be used to feed data directly back into the CAVE and adjust the data in real-time rather than through a computer placed outside of the CAVE space.

Body sensor-based recognition. The possibility of being able to interact with a CAVE environment by thought alone is an exciting prospect. One type of body sensor device that offers the promise of thought based interaction are headset devices which collect neuro-signal information through multiple sensors placed on the users head. Devices such as Emotiv Systems EPOC neuroheadset[7] is one such controller device that is able to generate keystrokes from neuro-signal data. While this device can be easily integrated into a CAVE environment there remains a big question mark over its accuracy and ultimate usability in CAVEs.

There are many other types of body sensors (e.g. sensors for monitoring heart rate, body temperature, fall detection, etc.) which also offer the possibility of interacting with CAVE environments. So, for example, we may want to use a body sensor in the CAVE to help simulate a real-world scenario (such as a fall) which when detected activates an alarm in the virtual home.

How easily body sensors can be integrated into a CAVE will vary significantly from sensor to sensor and what technologies, protocols and connecting hardware and software they normally use, as well as the wearability and the level of comfort when the sensors are wired to the body.

[7]http://www.emotiv.com/

Computer vision-based gesture recognition. An alternative way of recognising human gesture in the CAVE environment is by using computer vision techniques to provide "controller-free" interaction whereby the user does not have to hold any physical device or sensor. In this regard, the Microsoft Xbox Kinect offers probably the most exciting new method of interaction for CAVE environments. The Kinect has already been successfully integrated into our CAVE and allows its users to interact with virtual environments using body gestures alone[8]. The Kinect is able to detect multiple large objects such as people's bodies and their exact position within a space, but also fine details such as individual finger movements at specific ranges. Using FAAST(Flexible Action and Articulated Skeleton Toolkit) we are able to generate a simple skeleton over a user's body image or points over a person's fingertips. Once this is done we can configure FAAST to detect specific skeletal or hand movements and tolerances (i.e. clearly detect detailed body actions) and translate them into actions in the virtual environment. All this is achieved through the low cost Kinect device which uses nothing more than an RGB camera, infrared depth sensor and microphone array.

Object tracking. When we think of object tracking we often think of head tracking which adjusts the image we see on the screens based on our head position. However, because we know the exact positional co-ordinates of our head or any other object that is being tracked within the CAVE space, we can also use this information to interact with the CAVE in other ways. One way is to use object tracking to generate a virtual pointer from a controller device allowing the user to select, move and manipulate objects and menus in the virtual world. Another way would be to trigger certain actions when a tracked object is within specific co-ordinate positions within a CAVE space. So, for example, we could trigger a light being switched on in a virtual room when our tracked object was below a certain height (this could simulate the scenario of a light being switched on if someone fell to the ground at night).

6. CASE STUDY: THE GREAT NORTHERN HAVEN SENSOR DATA

6.1 Data Visualisation in the CAVE

The core of our research is based around the real-world data being generated at the 16 apartments that make up the GNH complex in Dundalk, Co Louth. With thousands of sensors gathering vast amounts of data on all aspects of the apartments 24/7, we are in the enviable position of having access to information that is not only immense in quantity and detail but which is generated by real people living real lives. The entire building contains a total of 2,240 sensors and actuators with 100+ in each individual apartment in addition to other interactive technologies (such as internet televisions and touch screen tablets). The sensors include door and window contact sensors, presence sensors, temperature sensors, light sensors as well as water, heating and electricity usage sensors. Every time a sensor is triggered the data is instantly recorded in our central database. With this data we can better understand the day to day actions and activities of the older people living in the apartments and in the process help to develop, implement and better

[8]See a video demonstrating this at: [anonymous for review]

utilise AAL technologies. It is through these technologies that we hope to add tangible benefits to the overall health and well-being of older people and help them to live longer and more independently in their home of choice. The work is being developed in close partnership with the residents of the GNH to ensure that the technologies being developed and implemented provide a real benefit to them.

The CAVE provides us with the ability to work with the data gathered from the GNH and visualise and interact with it in very unique ways. Through the CAVE we have the potential to better understand and interpret the data being collected as well as replicate potential scenarios or even interact with or feed data back to residents in real time.

In terms of visualisation, our key focus in this paper is on how we view real-world sensor data in a CAVE environment. In this context we have concentrated on two approaches. The first, termed *Built Environment Data Integration Visualisation*, concerns how we can present data in a 3D representation of a real world environment. The second, termed *Abstract Data Visualisation* investigates how we can present data in an abstract way that provides us with useful and important information that we cannot obtain from standard two dimensional graphs and charts. To demonstrate each of these methods of visualisation leading to our apartment data, we give a simple example of how each of these might be implemented:

6.2 Built Environment Data Integration Visualisation

Imagine that a user is standing in the CAVE in the middle of a 3D model of a kitchen containing all the usual kitchen appliances such as fridge, cooker, kettle, toaster, microwave etc. Now let's assume that each of these electrical appliances represents an appliance in a real-world kitchen and each of these electrical appliances is connected to a sensor which monitors usage. Using the data collected from the real-world kitchen we can feed that information directly into our 3D model where we might for example see a label with each appliance informing us of its usage details or perhaps use heat mapping to indicate the appliances that are used most (or which use the most electricity). Perhaps we might even use the information to indicate movement and usage around a home which could be represented by an avatar moving from appliance to appliance. In each case we are feeding data directly into the 3D representation of the real world enabling us to interpret the data in completely different ways.

In this scenario we generate a three dimensional model or representation of one of the GNH apartments in the CAVE in full stereo 3D. This allows CAVE users to immerse themselves in the virtual apartment. Using one of the many interaction methods discussed in Section 5 we can then start to navigate and explore that world as if we were walking around the real apartment. This scenario on its own provides an excellent way to assess the suitability of building design for specific requirements. We have already engaged older volunteers to assess the GNH apartments where issues such as counter heights, shelf heights and turning cycles for wheelchair access were some of the specific topics identified and discussed. Once a particular issue is identified then the virtual model can be quickly modified and re-assessed.

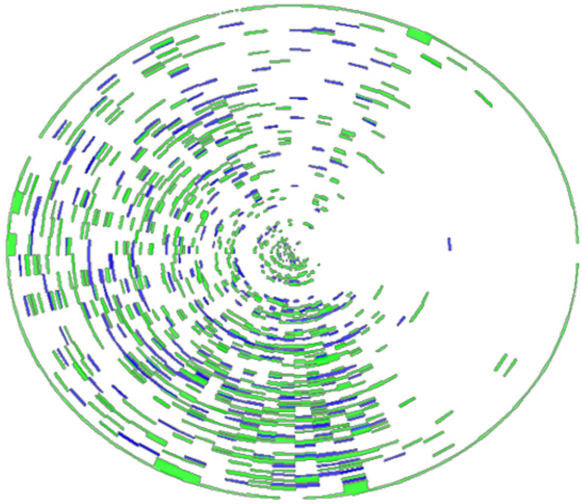

Figure 2: Clock plot displaying data from one individual sensor at the Great Northern Haven

Figure 3: Displaying Great Northern Haven PIR sensor data using a clock plot within the CAVE

The next step we have taken is to integrate sensor objects, sensor data, trigger points and actions within the virtual apartment. In this case we highlight different sensors throughout the virtual apartment with a red marking to make them easily visible. If a user navigates close to one of the markings it will trigger the generation of a clock plot graph using data from a real apartment. This will provide the user with a visual representation of sensor activity over a defined period of time (see Figures 2 and 3). Using this method the user can identify a sensor in its proper context and instantly view the data in a clearly interpretable and understandable way.

This Clock plot can be easily deciphered once we comprehend the elements that constitute it. Firstly, we need to realise that each colour marking on the clock plot represents a point in time when the sensor has been triggered (remember this clock plot is showing information for just one sensor in the home). Next we need to view the plot similar to how one would read a traditional clock face. In this case, however, a full 360 degree rotation of our "imaginary" hour hand represents 24 hours instead of 12. This allows us to visualise the activity of a sensor in a 24 hour period in a very clear manner. The next element that makes up the clock plot is the layering of these 24 hour rings starting from the centre point of the circle with each new ring wrapping itself around the outer edge of the last one. This allows us to view the activity of the sensor over a period of many months rather than just one day and enables us to identify and compare patterns of behaviour. Finally, we can use additional colour codings within the plot to further identify specific elements. In this case we have separated weekday and weekend data using blue and green colour markings.

Once we understand these elements we can start to see clear patterns of activity where the sensor is being triggered. So, in the case of the clock plot as illustrated in Figure 2, we can easily identify that the left-hand side of the clock face has a lot more activity than the right-hand side. The activity on the right-hand side of the plot relates to night time activity and the left-hand side relates to day-time activity. We also start to identify clear lines and clusters of data demonstrating particular activity periods as well as what appears to be a difference in patterns of behaviour between weekdays and weekends. This is just one example of how integrated clock plots can be a powerful and effective tool in helping us to understand and interpret large quantities of sensor data.

Now that we have the sensor data within the CAVE we can use different methods to learn even more about the activities within a particular apartment. Using heat mapping we have the ability to apply different colours to specific zones where sensors are being used. The more a sensor is being triggered, the hotter (or redder) that zone becomes. This allows complex sensor activity data to be displayed in a way that can be easily visualised by the CAVE user. Generating heat mapping on a model in a time lapsed manner could help in identifying interesting behaviours or patterns within an apartment.

Keeping in mind the idea of generating a time lapse sequence, the CAVE also offers us the possibility of playing back actions and activities of a particular apartment over varying time frames and speeds. So, for example, we can use an avatar to represent a person triggering sensors within an apartment and simulate any resulting actions such as doors and windows opening and closing, lights turning on and off, or any other identifiable electrical appliance being switched on or off at the same time that we ourselves are immersed in the environment observing it all happening. Integrating the temporal aspect within the spatially-oriented world can be designed in such a way as to explain and narrate a set of phenomena or stories in a coherent way, ultimately turning into a 3D CAVE version of a "visual confection", an assembly of many visual events, selected from various streams of story then brought together and juxtaposed [13].

Finally, we are able to integrate live video feeds into the virtual worlds generated within the CAVE. This opens up the possibility to do what can best be described as a "reverse augmented reality" i.e. superimposing a real world image onto a user's view of a computer-generated CAVE world.

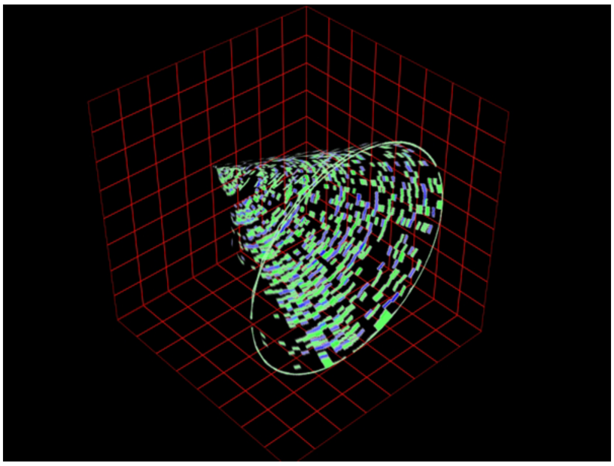

Figure 4: Three dimensional clock plot

This means we have the potential to immerse ourselves in the CAVE and at the same time engage in a live video and audio feed with one of the residents in their home at the GNH.

6.3 Abstract Data Visualisation

Using the same example of collecting data from the electrical appliances of a real world kitchen, instead of using a 3D representation of the kitchen in the CAVE to view that data, this time we generate three dimensional charts using the data. Let's assume we wish to monitor the usage of a kettle every day for a period of one month. A standard two dimensional graph can easily provide us with that information. However, using the CAVE we can now add depth to the data and transform this graph in to a three dimensional graph allowing us to show not only an x and y plane but also a z plane and in full stereo 3D. This means that we can now map the usage data to the x plane, the daily time data to the y plane and the monthly time data to the z plane producing something similar to a terrain map. Although this is a relatively simple example, the CAVE offers us the potential to visualise vastly more complex data and perhaps more clearly identify patterns within it. Due to the almost complete immersion provided by the CAVE we also have the ability to visualise and compare vast numbers of complex datasets set side by side, layer by layer in a way which would not be possible using traditional monitors or screens.

Although still at an early stage of our CAVE development, one of the first steps we have taken is to gather data from the apartment sensors and generate three dimensional versions of our two dimensional clock plots within the CAVE (see Figure 4). For example, if we take data from a living room presence sensor we are able to generate a three dimensional cone that shows us every time that sensor was triggered within a 24-hour period. This is then enhanced by presenting that data over a period of 6 months represented by colour markings positioned along the depth of the cone. Individual colour markings can then also highlight specific days of the week or a simple weekday/weekend mix. All this information displayed within an immersive three dimensional cone helps to better illustrate and interpret the breakdown of the resident's activity within the living room during the specified time.

Three dimensional clock plots are one of the primary methods we are currently using to visualise our data but we are also looking at many other suitable methods such as 3D surface maps, 3D scatter plots, 3D bar charts, etc. They all have the ability to teach us something new about the data within the immersive setting of the CAVE. 3D scatter plots using data from multiple sensors and which are colour-coded have the potential to provide us with similar information to heat maps within a built environment model but without the need to model a representation of the apartment where the sensors are installed.

Because the CAVE can generate what appears to be an infinite expanse to its users, we can display large numbers of three dimensional charts, plots and graphs at any one time and in any position within a three dimensional space. This allows users to easily compare and interact with large numbers of different datasets at any one time. We are currently working on a number of combinations of these while identifying and recording various visualisation and interaction issues arising from applying each of these to our CAVE environment.

7. CONCLUSION

Visualising and interacting with real-world sensor data in a CAVE can be a powerful tool in a research environment. It has the ability to immerse users and present data in truly unique ways. The oft-used phrase "it has to be seen to be believed" still has some merit in relation to the CAVE. Its ability to mix virtual and real-world elements in such an immersive way can be both fascinating and highly informative.

For this paper we have focused on a specific aspect of CAVE use. However, new technological developments such as the Google Liquid Galaxy project (which enables us to view Google Earth and Street View across screens), Xbox Kinect integration (which offers controller-free interaction) and the latest synchronizable WebGL browsers (enabling plugin free 3D) offer us many exciting new avenues for CAVE research.

While there is no doubt that CAVEs can be extremely beneficial they are costly (particularly in terms of the initial financial investment required), even if there are exceptions to the rule [5]. They also require a significant investment of time in order to be able to generate useful and valuable results. Challenges in terms of its usability include motion sickness and fatigue when used intensively or for an extended period of time, user training for the modalities and controllers provided, and the lack of current usage resulting in difficulty to find more practical but novel use scenarios today, all of which provide directions for future research.

We have been able to use sensor data in new ways that enhance our understanding of the lives of the GNH residents. This has been possible because of the unique data available to it and the unique resources such as the CAVE that the research team has at its disposal. Through the visualisation and interaction of real-world sensor data in the CAVE we are able to interpret information in entirely new ways and in the future better understand how we can use AAL technology to improve levels of independence and well-being for older people.

8. ACKNOWLEDGEMENTS

Anonymous for review

9. REFERENCES

[1] D. A. Bowman, J. Chen, C. A. Wingrave, J. F. Lucas, A. Ray, N. F. Polys, Q. Li, Y. Haciahmetoglu, J.-S. Kim, S. Kim, R. Boehringer, and T. Ni. New directions in 3d user interfaces. *IJVR*, 5(2):3–14, 2006.

[2] D. A. Bowman, E. Kruijff, J. J. LaViola, and I. Poupyrev. An introduction to 3-d user interface design. *Presence: Teleoper. Virtual Environ.*, 10:96–108, February 2001.

[3] D. A. Bowman, E. Kruijff, J. J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 2004.

[4] S. Bryson. Virtual reality in scientific visualization. *Commun. ACM*, 39:62–71, May 1996.

[5] D. C. Cliburn and K. Stormer. The hive: Hanover immersive virtual environment. *J. Comput. Small Coll.*, 20:6–12, April 2005.

[6] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: the design and implementation of the cave. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '93, pages 135–142, New York, NY, USA, 1993. ACM.

[7] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The cave: audio visual experience automatic virtual environment. *Commun. ACM*, 35:64–72, June 1992.

[8] F. Durand. An invitation to discuss computer depiction. In *Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering*, NPAR '02, pages 111–124, New York, NY, USA, 2002. ACM.

[9] J. Mortensen, P. Khanna, I. Yu, and M. Slater. Real-time global illumination in the cave. In *Proceedings of the 2007 ACM symposium on Virtual reality software and technology*, VRST '07, pages 145–148, New York, NY, USA, 2007. ACM.

[10] M. A. Otaduy, T. Igarashi, and J. J. LaViola, Jr. Interaction: interfaces, algorithms, and applications. In *ACM SIGGRAPH 2009 Courses*, SIGGRAPH '09, pages 14:1–14:66, New York, NY, USA, 2009. ACM.

[11] R. Pausch, D. Proffitt, and G. Williams. Quantifying immersion in virtual reality. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '97, pages 13–18, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[12] A. Sutcliffe, B. Gault, T. Fernando, and K. Tan. Investigating interaction in cave virtual environments. *ACM Trans. Comput.-Hum. Interact.*, 13:235–267, June 2006.

[13] E. Tufte. *Visual Explanations: Images and Quantities, Evidence and Narrative*. Graphics Press, Cheshire, Connecticut, U.S.A., 1997.