CENTRE *for* DOCTORAL TRAINING *in*
CYBER SECURITY

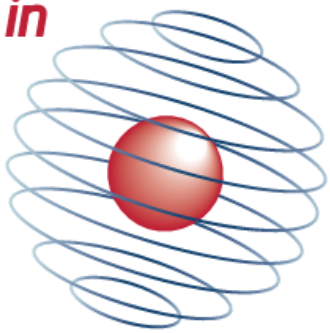

CDT Technical Paper

13/15

Stereoscopic Cyber Visualisations

Alastair Janse van Rensburg

Stereoscopic Cyber Security Visualisations

Alastair Janse van Rensburg

I. Introduction

Text-based tools are the primary tools of cyber-analysts, despite the potential visual tools have in this field [1]. Currently, analysts rely on command line tools which are favoured for their interoperability and flexibility. While many visualisations for cyber security data exist, they suffer from lack of adoption, due to not fitting in with the workflow of their users [2]. Some of the key challenges for security analysis are well-suited for visual solutions. Security analysts are commonly presented with large quantities of data to process, from many distinct data sources [3]. Using this data, analysts must obtain situational awareness of their networks in order to spot anomalous patterns as they occur.

With this in mind, the aim of this project was to explore new techniques that could have applications in cyber security visualisations. Specifically, the project aimed to explore the use of stereoscopic displays in cyber visualisation. It was hoped that visualisations based on stereoscopic technology would solve many problems for analysts, enabling a greater variety of techniques and putting them in a position where data can be easily presented to them.

First, an examination of existing work on stereoscopic visualisation was undertaken. Second, a collection of techniques was identified that could be utilised in a cyber visualisation. Third, a cyber dashboard proof-of-concept was built, consisting of a number of visualisations that explored the use of the identified techniques. Finally, a short pilot study was conducted to explore what potential the techniques could have in the future.

Despite problems with current hardware and with the designed visualisations, pilot study participants were broadly positive of their experience with the developed system and all felt that the techniques have potential.

II. Related Work

A. Security Visualisations

Security analysts often work in *Security Operations Centres (SOC)*. SOCs provide a centralised location for analysts work in and process data. In doing so, they aid in situational awareness, enabling analysts to discover and respond to threats on their networks [4]. SOCs may be supplied with information from across a number of networks, and may have access to information such as [5]:

- Firewalls
- Network flow data
- System and application logs
- Performance data
- Vulnerability scanner results
- IDS alert history
- File integrity checking tools

A lot of this data must be sorted through in order to find valuable information: in some instances up to 99% of intrusion detection alarms are false positives [6].

Much research has been done into the use of visualisations in cyber security. Overall usage of visualisations is low. Analysts mistrust current visualisations and see them as being responsible for over-aggregation, for hiding data and for being difficult to use in conjunction with other tools [2].

Despite this, visualisation has been shown to increase accuracy and speed of analysts [7] and system administrators can be enthusiastic about the use of visual tools [1].

Fig. 1. An example security analyst workflow

A typical workflow [5] begins with a monitoring stage, where analysts seek to understand their network and catch unusual or suspect events which require further investigation. This is followed up with an analysis stage. At this point, analysts seek to understand what caused their initial concerns and build a detailed understanding of the activity. Finally, the process moves to a recovery phase which will seek to undo the effects of the intrusion.

The monitoring stage is particularly important, and analysts have reported that one of the most critical requirements for cyber tools is the need for situational awareness - that is, understanding of the status of their networks as a whole [8]. This requirement stemmed in part from the need to continuously monitor the entire network.

B. Stereoscopic Displays

Stereoscopic displays are not a new invention. The first such device, using two mirrors to display two different panels to the observer, was devised in 1838 [9]. Recently there has been a surge in interest in creating head-mounted stereoscopic devices, as modern screen technology enables high resolutions displays to be mounted directly onto the headset. As a result, many companies are developing headsets that are set to come to the market from 2016. In particular, the Oculus Rift has been widely credited with leading the way in consumer-oriented virtual reality headsets [10], with the first commercial release planned for early 2016.

As a result of this significant investment into the technology behind the headsets, this project does not aim to examine the quality of the headset itself, and instead makes the assumption that many of the technical issues connected with current headsets (such as poor resolution or low refresh rates) will be temporary and resolved as new devices become available.

These headsets consist of two primary components. First, they have a display system designed to enable each of the user's eyes to see a distinct screen, or section of screen. Second, they have some capability to monitor the rotation (and often position) of the device, which is fed back into the virtual environment, normally to position the viewpoint.

While the vast majority of applications of VR headsets involve a full 3-dimensional world, there are some examples of stereoscopic vision being used in 2-dimensional data visualisation.

The technique of *stereoscopic highlighting* involves displaying the primary visualisation on a plane in front of the viewer, and highlighting some elements by moving them closer. This gives them the illusion of floating closer to the user and aims to draw attention to them, or make them look more significant. This technique was shown to be promising [11].

With similar technique, it has been shown that displaying labels at different distances from the viewer significantly aids decision time [12]. This technique involves placing overlapping labels at different positions in order to separate them from each other and make them clearer to read.

Stereoscopic vision is also used more commonly in the display of 3D models, such as bio-imagery [13], medicine [14] and scientific visualisation [15].

Images can also be presented to the eye dichoptically, that is, so that each eye receives images which are distinct and are not intended to create the illusion of depth [16]. These techniques can be used to

highlight data or to otherwise enhance the display of information. In particular, one of these techniques involves varying the lightness of a displayed object between the two eyes. This was used to provide a method to aid colour-blind users, which was found to be fast and accurate. This binocular lustre effect was found to be comfortably distinguishable for up to three discrete levels [17].

C. Motivation

The use of a VR headset in cyber visualisation has a number of potential benefits.

The soon-to-be widely available headsets will provide a relatively easy and cheap way to set up analysts with such a system. It is not inconceivable that in near-future everyday users will be accustomed to interacting with virtual reality displays. This makes them an easy and available tool.

By using VR headsets, visualisation designers have access to additional techniques, such as those described above, which aid them in displaying information to analysts.

Further, the problems some analysts report involving their requirements for large workspaces [2] are mitigated by the use of virtual reality, where a user's equivalent screen space can be as large as required, without significantly increasing the cost of such a workstation.

III. Techniques

The main purpose of this project is to collect and present a collection of techniques that can be used in future visualisations. These techniques serve to provide designers with additional methods of highlighting and displaying data, in ways which can work alongside traditional techniques such as the use of colour, size and position.

A. Stereoscopic Highlighting

Stereoscopic highlighting involves the use of the 3D environment to move some elements closer to the user. Our intention is to use this mechanism to guide users to aspects of interest in the visualization itself. See Figure 2 for a generic example of how apples on a tree will stand out more if they are closer to the user.

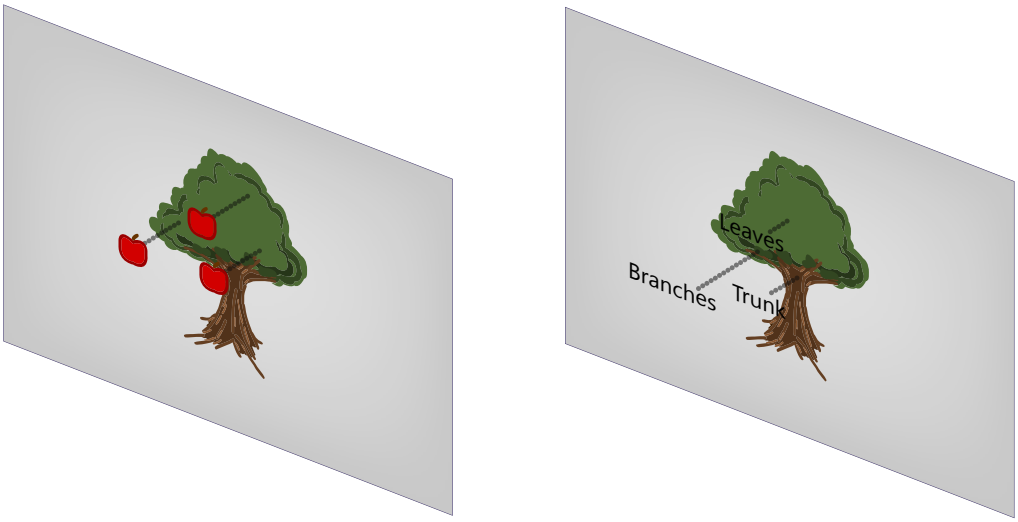

Fig. 2. In the left image, stereoscopic highlighting. The apples are lifted off the plane containing the image of the tree and moved closer to the position of the viewer. In the right image, stereoscopic separation. The labels are placed at different distances from the viewer in order to separate them.

This technique causes the elements to appear more significant and seeks to aid users in quickly identifying and locating parts of the display. As this technique does not affect colour or shape, it can be used without concern for conflicts. However, it may be unsuitable for applications where position is important, as the exact position of stereoscopically highlighted objects is not as clear. Because this effect is quite subtle, it may be less suitable for drawing attention.

This technique has been previously in the visualisation of graphs [11].

B. Stereoscopic Data Presentation

More generally, data can be presented at different distances from the user to separate it or differentiate it. See Figure 2 for a generic example of how labels can be placed at different positions to give them separation.

This can be done to present information to the user. For example, when intrusion detection alerts are placed closer to the user based on severity, to give emphasis and separate them. It could also be done to aid clarity, where information is placed on different levels to leverage stereoscopic effects and make them appear to no longer collide. Especially when coupled with positional head tracking, this enables users to "see around" elements of the visualisation in order to see data behind and prevent occlusion.

This technique has been used to make visualisations with many overlapping labels more readable [12].

C. Dichoptic Highlighting

Aside from the illusion of depth, a VR headset can be used to highlight data dichoptically, that is, by presenting different images to each eye. By changing the way a single element of the visualisation is displayed, a strong focus can be made in a way that should attract the attention of users. See Figure 3 for an example where the images shown to each eye are drawn side-by-side.

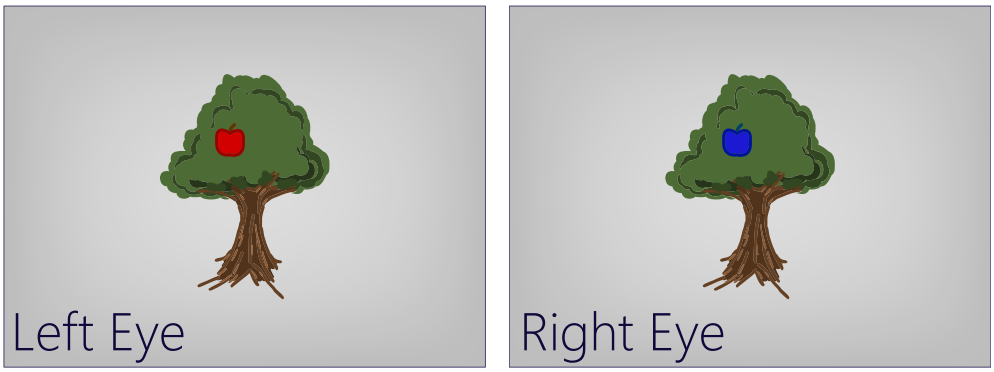

Fig. 3. Dichoptic separation. The apple in the tree is coloured differently between the two eyes.

While some applications of this technique may conflict with existing uses of colour in the visualisation, it can be applied subtly alongside colours. For instance, by varying the colour slightly (e.g. making it lighter) an effect known as binocular lustre can be achieved.

Aside from colour variations, texture can be presented differently to achieve the same effect.

This technique can be used to grab the attention of the user, as dichoptic images appear quite unusual.

An existing example of this technique can be found in [17]. Participants shown dichoptic highlighting in this study described it as either a positive effect: "highlighted", "blinking" or negatively as "distracting" or "annoying". While in most cases, "distracting" and "annoying" visual elements are considered problematic, there may be value in short-lived elements of this nature to rapidly and reliably draw the user's attention to something, such as a critical alert. This example also found that the highlighting could be applied without compromising the image composition by using two a pair of colours which differ from each other but not exceptionally from the background.

D. Dichoptic Data Presentation

Data can also be varied when presented to each eye. For example, some data points could be coloured when presented to one eye in order to present additional information to the user.

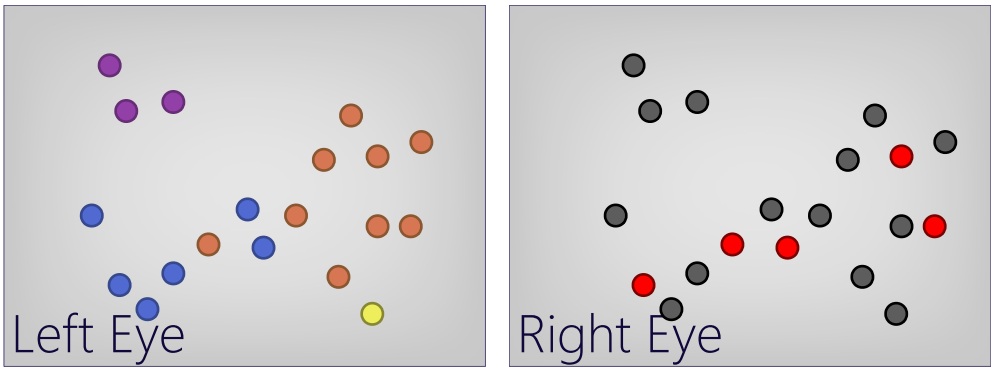

Fig. 4. Dichoptic separation. Data points are coloured according to groups on the left eye, but colouring is used to highlight data points on the right eye.

Alternatively, some information could be removed from one eye to prevent it from occluding the rest of the information.

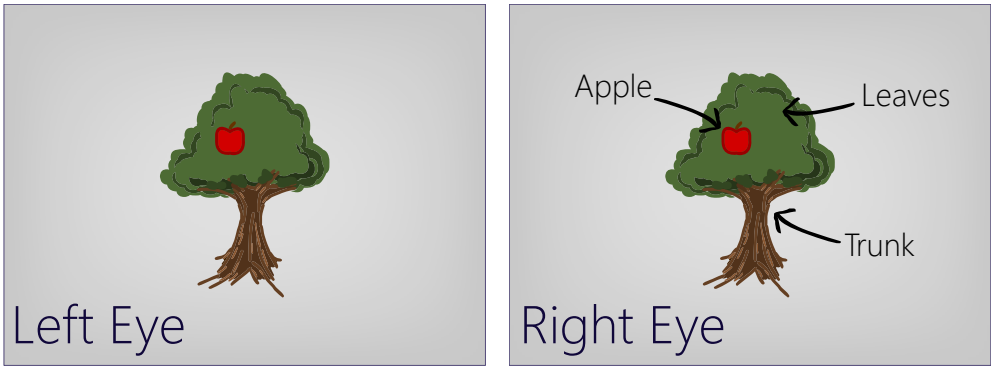

Fig. 5. Dichoptic separation. The labels and arrows are only displayed to the right eye to give the left eye an unobstructed view of the image.

In these dichoptic techniques, the intention is that the user can perceive the visualisation without extra interaction. However, there is the possibility that the user can close one eye to enhance their perception of the other data set.

E. Cyber Security Applications

It is hoped these techniques will be useful when applied to cyber visualisations. In particular, the requirement on cyber analysts to process large amounts of data from many sources, and to gain complete situational awareness over their network by doing so, suggests a need for techniques that enable them to interpret and understand visualisations quickly. In this regard, the techniques above should assist them by guiding their attention subconsciously and allowing them to view more complex visualisations without suffering from problems of occlusion or losing clarity.

To explore this more formally, we created a proof-of-concept system to examine the hypothesis: Is it possible to use these techniques to communicate cyber security data to cyber security experts?

IV. Tool Design and Implementation

To explore both these techniques and the use of VR headsets, an example system was created. This consists of a set of visualisations together with a 3D scene to present them within.

The use-case for the system is an analyst in a SOC environment. As a result, the visualisations display data that may be presented to a SOC operator. The system is intended for use on the Oculus Rift (specifically, the DK2 version).

The primary purpose of this system is to provide a context within which the techniques can be used. As a result, it is not intended to be usable independently.

A. Requirements

- A front-end which produces a display output.
  The output must be rendered for the Oculus Rift, so that the rendering takes into account lens distortion and chromatic aberration.
  The front-end must also receive input from the Oculus Rift in the form of head rotational tracking data, and use this to position the user's view.
- A set of visualisations which take data from the back-end and are rendered by the front-end. These visualisations should utilise the techniques listed above and provide insight into cyber security issues.
- A back-end, which provides information to the front-end. This information will be the source of the visualisations presented. It is not important that data sources are from the same data set, or that they corroborate, as they are simply used to populate the visualisations.

B. Design Overview

The final 3D scene consists of a collection of *displays* arranged in a cylindrical pattern around the user's position.

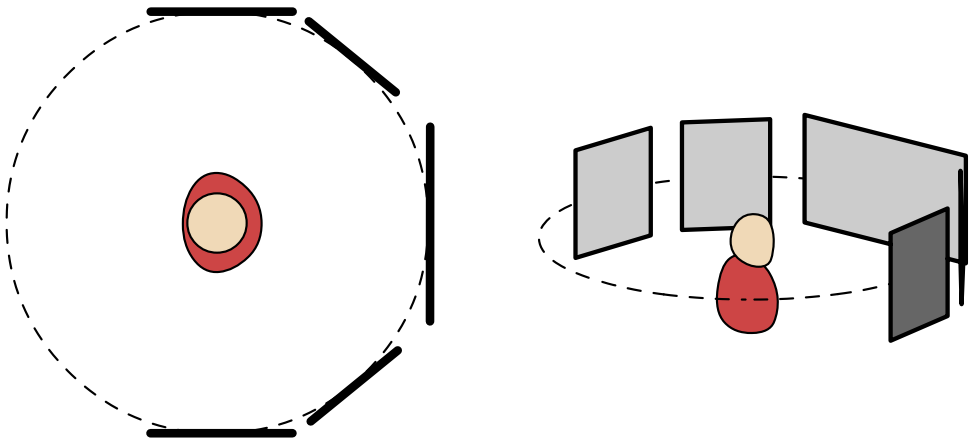

Fig. 6. On the left, a top-down view of the virtual 3D scene, with the user in the middle and virtual displays positioned around them. On the right, a perspective view of the same scene. The dashed line represents the user's camera level

Each display can be instructed to display the output of any of a set of *visualisations* which run in the background. This enables the user to move, replace and duplicate visualisations on the displays around them in any configuration required. The user's head rotation was tracked by the Oculus Rift and used to rotate the camera in the 3D scene. This enables to them to look at displays in a 180° arc around them.

Each display in the system consists of a set of *canvases*, each of which is a 2D vector image. These canvases are arranged in pairs and each pair is placed at a different distance towards the camera. One

canvas of each pair is drawn to the left eye and one is drawn to the right eye. This enables the display of both dichoptic images (by varying which canvas in pair is drawn to) and stereoscopic images (by varying which pair is drawn to).

Canvases that were closer to the camera were scaled accordingly, so that to the user each canvas appeared to be the same size. As an consequence, for any x and y, the points at (x, y) on each canvas were collinear with the user's camera. This means that when using stereoscopic depth in visualisations, the could simply move it from one canvas to a closer one and it would not appear to be offset, larger or out of place.

No interaction was implemented for the system except for a basic terminal system, which allowed the user to interact with the scene to move visualisations between displays, and to send commands directly to specific visualisations. This was not an ideal input method as usage of the keyboard while wearing the Oculus Rift is difficult.

The data used in the visualisations was a combination between the data provided in [18], and synthetic data created purely to populate the visualisations.

The visualisations implemented were as follows:

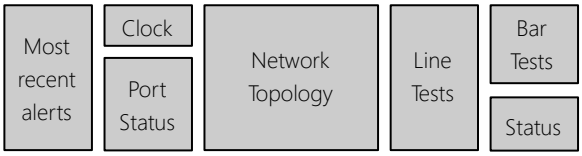

Fig. 7. The layout of the visualisations in the system

1) Text Display: This straightforward visualisation simply displayed text. In this case, it was a list of the most recent alerts.

2) Port Display: This visualisation consisted of two main elements. On the left, a rectangle was displayed, with each pixel corresponding to a port (arranged by port number). If there was a network flow within the last day on a port the pixel would be highlighted in red. If the port corresponded to one of the system ports (i.e. its port number was less than 1024) the port would be dichoptically highlighted, that is, displayed in red to one eye and yellow to the other. The intention of this was to draw the user's attention to these ports and make them stand out clearly. In this example, a simple recolouring of the pixel would not be sufficient to ensure it was differentiable from the other highlighted ports.

On the right hand side of the display, a few key ports were listed. These were drawn in grey when the port was inactive and yellow when active. When active, the ports were also drawn on a plane closer to the viewer to make them stand out further.

3) Clock: This visualisation was a simple digital clock, which served to demonstrate that the scene was capable of being updated in real-time. In a typical use-case, such a display may be necessary as the Oculus Rift prevents the user from being able to see anything other than the 3D display.

4) Network Graph: This visualisation was intended as a high-level display of the network being monitored. The visualisation consisted of a graph, with each vertex corresponding to a network node. Edges were placed to represent communication between nodes on the network. The data used in this visualisation was entirely synthetic in order to make it easier to visualise and to enable a focus on the presentation of the data.

This display made use of two techniques:

- Stereoscopic data presentation: Name labels for each network node were placed closer to the user. This separated them from the network graph and was intended to make them easier to read.

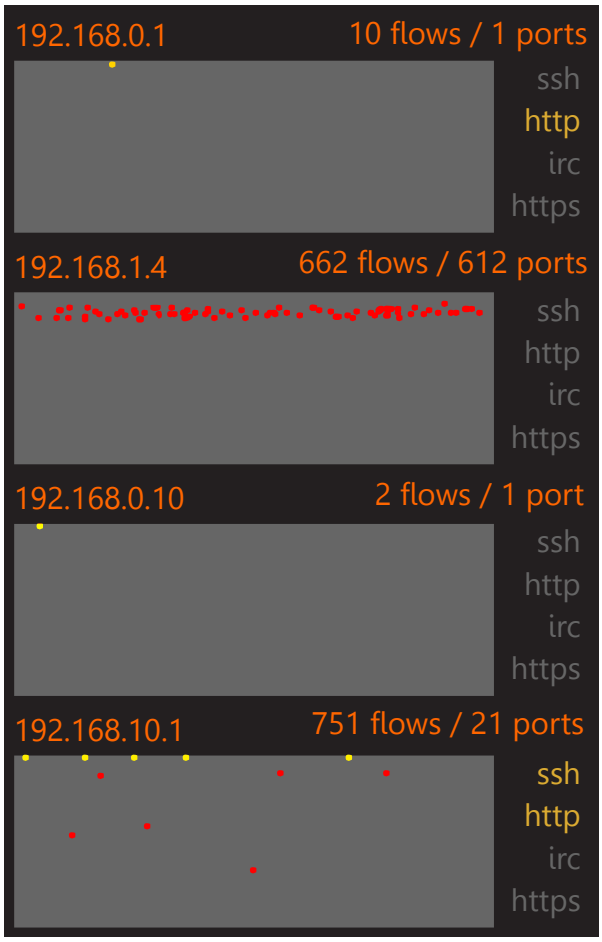

Fig. 8. A 2D mock-up of the port display. Yellow dots in the mock-up were dichoptically highlighted with red and yellow.

- Dichoptic highlighting: Warnings were placed over some of the nodes to indicate (for example) that IDS alerts were being generated connected to that node. One of these alerts was dichoptically highlighted in red and yellow to draw attention. This was seen as a good use-case for dichoptic highlighting as it represented something to which the analyst's attention should be drawn quickly. This warning was also placed closer to the user.

5) Example Graphs: These were a collection of graphs generated from synthetic data which were intended to explore dichoptic effects.

The first graph consists of a line graph with two lines, one to represent historical data (e.g. average server CPU usage throughout the day) and one to represent current data (e.g. current server CPU usage today up until now). The dichoptic data presentation technique is used here, by displaying the average server CPU usage to only one eye. This was intended to make this graph less obtrusive and to keep the user's focus more strongly on the current graph.

The second graph has a single line, and is highlighted with a red block whenever the graph reaches zero. In practice this could correspond to memory usage, and be highlighted whenever free memory became critically low. The highlighting is performed on only one eye. This was done so that it would not be as obstructive of the graph.

The third graph is a bar graph where the tops of the bars are dichoptically highlighted in red. This is intended to draw the user's attention to problematic areas.

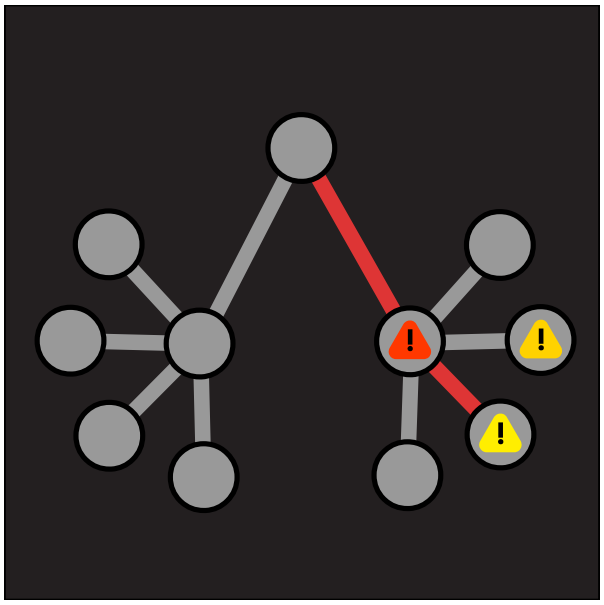

Fig. 9. A 2D mock-up of the graph display. Labels were presented stereoscopically highlighted below each node, but are omitted here for clarity. The warning in red was dichoptically highlighted in red and yellow.

The final graph consists of two entirely separate lines which are dichoptically separated. This is done to enable the user to see them both separately and together.

6) Example Graphs: These graphs were implemented as a test for the techniques and consist of a collection of bar graphs. Some of the bars are stereoscopically highlighted to draw attention and some are dichoptically highlighted. They were not used in the study.

C. Implementation

The front-end was written in JavaScript. This was done to enable access to the relatively easy-to-use WebGL. JavaScript was used to produce Canvas elements with the visualisations. Three.js was then used to create a virtual 3D scene which was rendered for the Oculus Rift.

Each visualisation consisted of a single class that was set up with data and contained a method to output to a given collection of canvases. As a result, conversion of existing JavaScript canvas-based visualisations to this system would be straightforward.

The back-end had two components: First, it exposed a REST API to which the front-end could send requests for data. In this system, these requests were made to a server which served data from a set of data flows [18]. Second, as the source of the data was not as important as the examination of the techniques, some entirely artificial data was created, which was stored locally in JSON format.

Outputting the display to the Oculus Rift was a straightforward process using three.js, and gathering headtracking data was made possible by running the visualisation within a WebVR-enabled version of the Chromium browser. Three.js [19] is a JavaScript library that handles the implementation of WebGL in the browser. It also has support for rendering to the Oculus Rift, which was used in the system.

V. Pilot Study

A. Design

To evaluate the system developed, and to gain feedback on the techniques, a small pilot study of 3 users was carried out. This study enables us to explore the feasibility and usability of the system. Participants

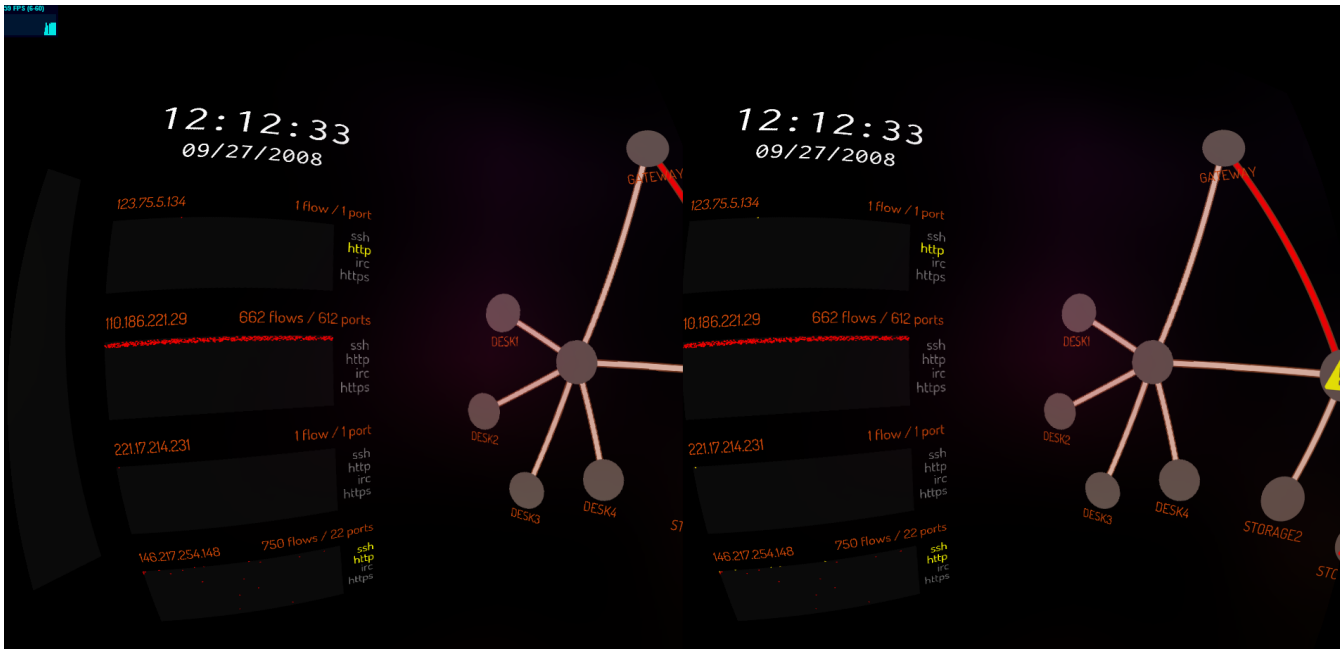

Fig. 10. Example screenshot from the system. The two distinct eye viewpoints can be seen. The image appears slightly distorted because it has been rendered to counteract the effect of the lenses in the Rift. The display on the left is showing the port display, and on the right is the network graph.

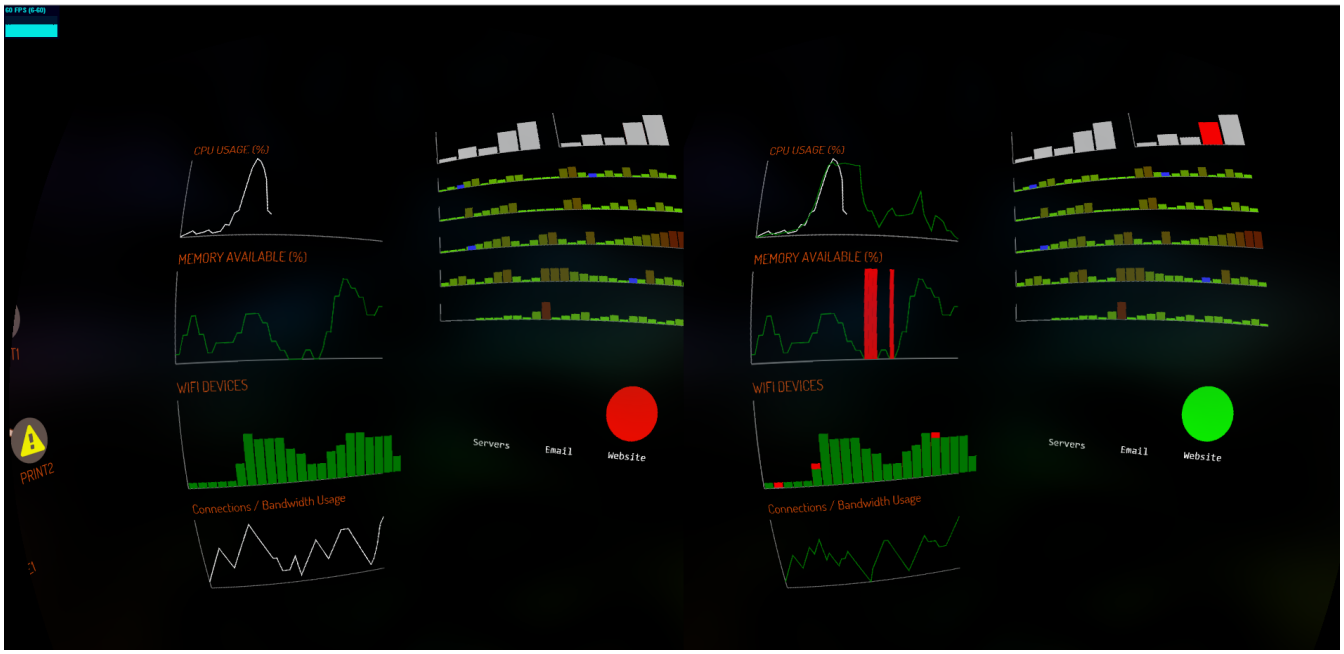

Fig. 11. Example screenshot from the system showing the example graphs. Dichoptic highlighting can be seen in one of the bars in the top right, which appears grey on the left eye and red on the right eye.

were asked to complete a small set of tasks inside the system, and after using the system were asked about their opinions of the system and the potential they felt the techniques have.

Due to the small sample size, this study was not intended to assess quantitatively how well the system performed, either on its own or relative to other systems. Instead, it was intended to examine how other users interacted with the system and whether the techniques are worth future exploration.

The participants were given the Oculus Rift (DK2) and asked to perform a simple set of tasks to ascertain how well they could use the system:

1) Orientation: The first task required participants to check they could see each display, especially those on the extreme right and left of them. This was intended to orient participants and ensure they could use the whole system. It was also an opportunity to assess how comfortable they were when looking at displays far to their side, which required significant rotation of the body or head.

2) Reading Information: The second task required participants to read text from the left-most visualisation. This was done to examine their ability to read text inside the 3D environment. They were asked if the text was clear, if they could read it and if it was comfortable to do so. They were also asked how the readability of the text varied according to its vertical position, as it had been noticed previously that text further from the horizontal was less clear.

3) Finding Information: The third task required participants to use the port visualisation to look up and read off information. Participants were asked to read the number of flows corresponding to a given IP, and were asked to find IPs which had http activity. This was to examine the usability of the system in a more natural use-case, where analysts might need to find information from the system.

4) Example Graphs: The fourth task involved the set of example graphs. Participants were asked some simple questions regarding each graph in order to encourage them to interact with them, so that they could use them to inform the discussions afterwards.

B. Results

No participants were colour-blind, and all participants were asked if they had any history of motion sickness, epilepsy, or previous problems with VR technology. Only one participant had ever used VR technology before and none reported any of the listed problems.

One participant wore glasses; the Oculus Rift was worn over the top of the glasses and this did not seem to cause any problems. Participants reported some initial discomfort and clarity issues when putting on the headset. Once the tasks began they were comfortable with it and did not report any problems when asked. Throughout the tasks, the most common issue brought up by participants was regarding the clarity of the display.

Because this was a pilot study, participants were given the tasks informally and were not recorded during the process. Participants were not explicitly told when techniques were being used until they had attempted the task.

1) Orientation Task: All participants were able to see displays in a 180° arc clearly and comfortably.

2) Reading Information Task:

- All participants were able to read all the text from the display with relative ease.
- Participants did not report discomfort while reading, but some initial difficulty while getting used to the headset.
- As anticipated, text further from the horizontal was reported as being less clear, although still readable.
- Participants reported significant chromatic aberration with the title text, which was both a different colour (white compared to the body of text being orange) and more displaced from the horizontal.

3) Finding Information Task:

- All participants were able to readily find the information corresponding to a specific given IP, and were able to correctly read off the number of flows that corresponded to this IP.
- One participant incorrectly reported the number of IPs with http activity.

- Regarding the dichoptic highlighting of some ports, one participant did not see the highlight as being clear and in particular reported that they saw the pixel and it's highlighted counterpart as two separate pixels, one of each colour. Only when informed of how the highlighting was being performed were they able to interpret it as intended.
- Other participants noticed the highlighting and said that it made the pixels stand out.

4) Example Graphs Task: The intention of this task was to expose the participants to the techniques in order for them to inform their answers later.

- Participants were asked if they could determine the shape of both line graphs in the first example. All participants correctly pointed out that the graphs overlapped at first but diverged later.
- Participants were not aware that the green line was being displayed to only one eye until this was explicitly pointed out to them.
- In the second example, participants were simply asked to confirm that they could see two areas of red highlighting on the graph, which all participants did.

5) Network Graph: Before removing the headset, participants were asked about the network graph display. In particular, they were asked about the warning symbols that were present on some of the vertices.

- Participants were able to identify the warning symbols and that one was dichoptically highlighted.
- One participant was very positive of this highlighting, saying that it strongly drew their attention. They observed the colours superimposing as intended and said this stood out more than a single colour would.
- One participant expressed concerns about the colours used in the graph and said this made it hard to see the highlighting.
- The final participant confused the highlighting with aberration and thought it was a mistake.
- Participants were asked if they could see that some labels were placed closer to them. All participants could, but were mixed about whether this made them clearer or not.
- One participant said it was an interesting feature and that it made the text slightly easier to read.
- Both other participants raised issues with the display, reporting that the text was harder to read because it was not curved around the viewer but a flat plane.

6) Post-session interview: One participant felt that the technology was not mature enough, and that the tool was a complicated way to display information. They particularly mentioned chromatic aberration as a main problem. One participant enjoyed the experience and felt that there were lots of applications. They expressed some concern that it took some getting used to. The final participant was also positive of their experience but felt they may struggle to use it for extended periods.

Advantages the participants felt the tool had over traditional visualisation methods included:

- Stereoscopic highlighting was beneficial.
- You can focus on the data and avoid background interference from other things around you.
- You can use many visualisations at once.
- You don't need many separate monitors to hold all the information.
- It made it easier to handle a range of information as there was no need to have it all in one place.
- The device let you have a lot more information in a lot less area.

Problems or disadvantages reported included:

- Information could be harder to read.
- You can't use other things while using the headset.
- You have to use things specifically designed for the Rift.
- There's no interaction method.

7) Final questions: Participants were given a series of statements and asked how they felt about that statement, with answers as one of strongly disagree (SD), disagree (D), neutral (N), agree (A), strongly agree (SA). P1, P2 and P3 are the answers given by the numbered participant

Question	P1	P2	P3
The tool was useful for looking at the big picture of data	A	A	SA
The tool was useful for looking at detail in data	A	D	SA
The tool was easy to use	N	A	SA
The tool was caused discomfort or was otherwise unpleasant to use	D	N	D
The tool used stereoscopic techniques to increase my ability to interpret or understand data	SA	D	D
Stereoscopic techniques like those used in the tool have the potential to increase my ability to interpret or understand data	SA	SA	A
I felt like I was in a different world while using the tool	A	SA	D

Of particular interest, all participants felt the techniques had potential, with two of three strongly agreeing. Only one participant felt the techniques were useful as applied in the study. This participant was also the participant who was able to see both instances of dichoptic highlighting as intended.

Interestingly, one participant felt the tool was useful for looking at both the big picture and detail of data, but disagreed that stereoscopic techniques were helpful as used. This may indicate that they felt the nature of the headset was a useful tool for visualisation, with features such as headtracking and the larger display area.

C. Analysis

With one exception, all tasks were completed correctly, with participants reading text accurately and finding the required IP. The only mistake was made when identifying the number of highlighted http ports where one user misidentified an http highlighting as a different port. This suggests that overall the ability of users to read and understand information from a VR headset is good.

The most common issue reported was clarity. This seemed to stem from two main problems:

- Firstly, the while the Oculus Rift's resolution has improved greatly between the two available development kits (1280x800 in the first version and 1920x1080 in the second, later version) it still appears relatively poor, especially when contrasted with a normal screen. This is an issue with the hardware of the system and should become less of a problem as the technology behind the headsets is improved - the first release version of the Oculus Rift, scheduled for Q1 2016, is reported to have a resolution of 2160x1200.
- Secondly, the nature of the display inside the Rift means that when viewed through the lenses, a significant amount of chromatic aberration can occur. This issue can be improved through software-based shaders and careful calibration of the device. During the study, all participants reported problems with chromatic aberration, suggesting that there were problems with the setup. While chromatic aberration shaders were utilised by the system, it seems they were not sufficient to deal with the problem entirely.

Overall, the issues regarding clarity should be lessened when newer models of VR headsets become available, and with calibration and design of the system being used. There were no issues reported that suggested any fundamental clarity issues with the use of VR headsets in general.

The third task examined the value of dichoptic highlighting of single data pixels. The results were partly unsuccessful, with one participant unclear that highlighting was taking place. This is likely due to the small area that was being highlighted and due to issues of chromatic aberration. Two versions of the highlighted pixel were reported to be appearing next to each other instead of overlapping (with one on

each eye) as intended. This is very similar to the effect of chromatic aberration, where the image appears blurred to the side in a different colour. This suggests that if dichoptic highlighting is performed, it should be done carefully to ensure it is very clear to users. This issue may be compounded by the fact the pixels were relatively isolated, which gave few reference points for the eyes to aid with binocular fusion.

The participants who did see this as intended felt that it made the pixels stand out more clearly, suggesting that this technique has potential provided it is used correctly.

Similarly, when examining the dichoptic highlighting present in the network graph, participants were mixed on the effect this had. Two participants described it as having no benefit, one explicitly pointing out that they had originally thought it was an issue of chromatic aberration before being informed how it worked. This problem may have been accentuated by the fact that his warning was also placed closer to the user.

The third participant saw the highlighted warning symbol and felt that the highlight did contribute to making it more visible - more so, they felt, than individually brightly coloured warnings could.

VI. Discussion

The results of the study suggest that there is potential in these techniques, but that there were problems in their implementation.

These problems can be improved from the feedback of participants:

- Participants had trouble observing dichoptic highlighting: including visual clues, such as a clear background image, could aid binocular fusion and provide a baseline for the highlighting to contrast with.
- Participants reported problems with chromatic aberration: careful control of chromatic aberration shaders and ensuring that all the settings of the display are personalised for each user would minimise this effect. This is less likely to be an issue in real applications as there would be more time to adjust the settings over time to personal preference.
- Single pixel dichoptic highlighting could be unclear: this issue is closely related to problems of chromatic aberration, as it is likely that users were confusing the highlighting with chromatic aberration.
- Participants had trouble seeing the effect when dichoptic highlighting and stereoscopic highlighting were combined: it may be inappropriate to combine these two effects as together they are difficult to see clearly.
- Content that was above or below the viewpoint was harder to read: this issue could be solved by implementing a sphere-based layout for the displays instead of a cylindrical layout.
- In the layout used (see Figure 6) the displays are in tangential planes to a cylinder around the user. This means that, especially for large displays like the network graph, content towards the edges of the display appeared to be further away and at an angle.

In general, headset users felt they were able to read and understand all the data they were being presented with. This suggests that there is potential for future applications of headset-based visualisation. The ability of the headset to show many displays is a benefit [2] which one study participant picked out specifically. By combining the displays into a single controllable 3D environment, the user gains flexibility and control over exactly how their information is displayed to them.

Study users reported some discomfort and lack of clarity but these seemed to be issues with not being used to VR headsets. In particular, one participant reported problems with clarity as being their main problem throughout the study. Once the final questions were complete, the participant tried the headset on again and discovered that in their second use everything was much clearer and easier to read. This may have been due to being more accustomed to the headset or due to having adjusted the headset when putting it on the second time. This reinforces other observations that suggest VR technology requires getting used to. In a SOC environment this should not present a problem, as analysts should have ample time to acquaint themselves with any system they use.

One possibility provided by the headset not explored in this study was the ability for some displays to be "attached" to the user's perspective, so that they are always in front of the user regardless of where they are looking. In this way, small amounts of important information could be persistently in front of the user throughout their use of the system. In a cyber context, applications could include critical alerts appearing at the top of the user's vision, or summary statistics being displayed at the side. This means that such information is ever-present regardless of which display the user is looking at. Coupled with techniques such as dichoptic highlighting, this would provide a quick way to reliably alert analysts to problems the moment they arise, without requiring them to be looking at a particular location.

Based on the results of the study and from the experience of developing the visualisations, the following observations were made about the techniques proposed:

1) Stereoscopic Highlighting: This technique provides a convenient way to mark elements of the visualisation and make them stand out from others in an unobtrusive way. As a result, it can be readily used in conjunction with other highlighting methods to enhance visual clarity. It does so without compromising the rest of the image. In a cyber context, potential applications could be the highlighting of key events on a timeline.

While highlighted elements were scaled and translated to line up correctly with their surroundings, it is impossible to avoid some lack of clarity in their positioning. As a result, this technique may be unsuitable for visualisations where the precise location of each element is important

2) Stereoscopic Data Presentation: This technique can aid the clarity and readability of information, particularly with regards to text - the detail of which lent itself well to stereoscopy.

3) Dichoptic Highlighting: Problems with the study meant that users were not, in some cases, able to observe dichoptic highlighting clearly. In future applications, it is suggested that dichoptically highlighted elements are embedded within backgrounds that aid binocular fusion.

When observed correctly, however, participants felt that dichoptic highlighting helped draw their attention strongly.

4) Dichoptic Data Separation: Dichoptic data separation in the example took the form of multiple datasets being presented together. In the study, these presentations came across as being either unclear or unnoticeable, and participants struggled to see the benefit of such presentation. This may have been caused by the significantly contrasting colours used.

An alternative method of implementing dichoptic data separation uses multiple discrete levels of binocular lustre to differentiate between data values [17]. This may be a more practical way to use this technique.

A. Future Work

Further research into this area would benefit from a study into the techniques in a general context, looking at both quantitative and qualitative measurements of how participants interact with stereoscopic and dichoptic visualisation techniques. Such a study could firmly determine what, if any, benefits can be gained from using such techniques.

In a cyber context, an attempt to construct a working visualisation that could be provided to security analysts could be useful, particularly if coupled with a study of analysts using the tool. This would be very valuable to help understand how beneficial these techniques would be in practice. Such work would benefit greatly from the general study described above.

More specifically, it has been shown that stimuli moving towards the viewer have increased visual search priority [20]. This, and other animated variants of the techniques proposed, could provide further tools for visualisation designers.

VII. Conclusion

This project identified four techniques that utilise stereoscopic and dichoptic techniques and have potential for application in visualisations. These were:

- Stereoscopic highlighting
- Stereoscopic data presentation
- Dichoptic highlighting
- Dichoptic data presentation

An dashboard proof-of-concept was built to explore the techniques listed in context. Finally, a short pilot study was conducted to explore what potential the techniques could have in the future.

Overall, the techniques have a variety of purposes, many of which are especially useful in a cyber context in particular, dichoptic highlighting makes things stand out clearly and draws attention to them, which has a lot of applications within a cyber dashboard.

However, the results of the study highlighted that it is important to use these techniques carefully to avoid confusing users. This is particularly crucial for users who are new to VR technology.

REFERENCES

[1] Robert Ball, Glenn A. Fink, and Chris North. Home-centric visualization of network traffic for security administration. In *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, VizSEC/DMSEC '04, pages 55–64, New York, NY, USA, 2004. ACM.

[2] Glenn Fink, Christopher L North, Alex Endert, and Stuart Rose. Visualizing cyber security: Usable workspaces. In *Proc. of Intl Workshop on Visualizing Cyber Security (VizSec 2009)*, pages 45–56. IEEE, 2009.

[3] Daniel M. Best, Alex Endert, and Daniel Kidwell. 7 key challenges for visualization in cyber network defense. In *Proceedings of the Eleventh Workshop on Visualization for Cyber Security*, VizSec '14, pages 33–40, New York, NY, USA, 2014. ACM.

[4] Diana Kelley and Ron Moritz. Best practices for building a security operations center. *Information Systems Security*, 14(6):27–32, 2006.

[5] J.R. Goodall. User requirements and design of a visualization for intrusion detection analysis. In *Information Assurance Workshop, 2005. IAW '05. Proceedings from the Sixth Annual IEEE SMC*, pages 394–401, June 2005.

[6] Klaus Julisch and Marc Dacier. Mining intrusion detection alarms for actionable knowledge. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, pages 366–375, New York, NY, USA, 2002. ACM.

[7] J.R. Goodall. Visualization is better! a comparative evaluation. In *Visualization for Cyber Security, 2009. VizSec 2009. 6th International Workshop on*, pages 57–68, Oct 2009.

[8] William Yurcik, James Barlow, Kiran Lakkaraju, and Mike Haberman. Two visual computer network security monitoring tools incorporating operator interface requirements. In *In ACM CHI Workshop on Human-Computer Interaction and Security Systems (HCISEC)*, 2003.

[9] Charles Wheatstone. Contributions to the physiology of vision. part the first. on some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London*, 128:371–394, 1838.

[10] Samuel Gibbs. Oculus vr founder on the key to virtual reality's success: 'people are narcissists', January 2015. [www.theguardian.com; posted 07-January-2015].

[11] Basak Alper, Tobias Hollerer, JoAnn Kuchera-Morin, and Angus Forbes. Stereoscopic highlighting: 2d graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2325–2333, December 2011.

[12] S. Peterson, M. Axholt, and S.R. Ellis. Managing visual clutter: A generalized technique for label segregation using stereoscopic disparity. In *Virtual Reality Conference, 2008. VR '08. IEEE*, pages 169–176, March 2008.

[13] Harry J. Whitlow, Minqin Ren, Jeroen A. van Kan, Frank Watt, and Dan White. Exploratory nuclear microprobe data visualisation using 3- and 4-dimensional biological volume rendering tools. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 260(1):28 – 33, 2007. Nuclear Microprobe Technology and Applications (ICNMTA2006) and Proton Beam Writing (PBW II) 10th International Conference on Nuclear Microprobe Technology and Applications and 2nd International Workshop on Proton Beam Writing.

[14] C. Paloc, E. Carrasco, I. Macia, R. Gomez, I. Barandiaran, J.M. Jimenez, O. Rueda, J. Ortiz de Urbina, A. Valdivieso, and G. Sakas. Computer-aided surgery based on auto-stereoscopic augmented reality. In *Information Visualisation, 2004. IV 2004. Proceedings. Eighth International Conference on*, pages 189–193, July 2004.

[15] U. Obeysekare, C. Williams, J. Durbin, Larry Rosenblum, R. Rosenberg, F. Grinstein, R. Ramamurti, A. Landsberg, and W. Sandberg. Virtual workbench-a non-immersive virtual environment for visualizing and interacting with 3d objects for scientific visualization. In *Visualization '96. Proceedings.*, pages 345–349, Oct 1996.

[16] Haimo Zhang, Xiang Cao, and Shengdong Zhao. Beyond stereo: An exploration of unconventional. In *CHI 2012 Conference on Human Factors in Information Systems*. ACM, 2012.

[17] Soon Hau Chua, Haimo Zhang, Muhammad Hammad, Shengdong Zhao, Sahil Goyal, and Karan Singh. Colorbless: Augmenting visual information for colorblind people with binocular luster effect. *ACM Trans. Comput.-Hum. Interact.*, 21(6):32:1–32:20, January 2015.

[18] A. Sperotto, R. Sadre, D. F. van Vliet, and A. Pras. A labeled data set for flow-based intrusion detection. In *Proceedings of the 9th IEEE International Workshop on IP Operations and Management, IPOM 2009, Venice, Italy*, volume 5843 of *Lecture Notes in Computer Science*, pages 39–50. Springer Verlag, October 2009.

[19] Ricardo Cabello. Three.js. [http://threejs.org/].

[20] Steven L. Franconeri and Daniel J. Simons. Moving and looming stimuli capture attention. *Perception & Psychophysics*, 65(7):999–1010, 2003.

VIII. Appendix - Study Questions

A. Pre-session interview

How often to you play computer or video games?
Do you have prior experience with VR technology?
Do you have any history of motion sickness or epilepsy?
Are you colour-blind?
Do you wear glasses or contact lenses?

B. Tasks

(Before each task, direct participants to the relevant panel)
1) Task 1: Can you see the "Most Recent Alerts" panel on your far left?
Can you see the bar graphs and status circles on your far right?
2) Task 2: Can you read the text on the leftmost panel, "most recent alerts"?
Is it comfortable to read text at that angle? Is the text clear?
Does the readability of the text vary according to how high it is?
3) Task 3: How many flows have there been to the IP starting 146.x?
How many of the monitored IPs have "http" highlighted to signify activity on port 80?
4) Task 4: Each pixel in the port graphs represents activity on a different port. Are any of the pixels highlighted?
Do these pixels stand out significantly?
5) Task 5: Can you see two line graphs, one in white and one in green?
Can you see the patterns of both lines?
6) Task 6: Can you see the areas highlighted in red?
Does it obscure the rest of the graph, or make it hard to interpret?
7) Network graph discussion: The front-most display contains a network graph, where some nodes have warning indicators on them. One of the warnings is stereoscopically highlighted so that it appeared yellow and red.
Is this clear?
Does this draw attention to the warning and make it hard to ignore?
What is your opinion on this warning?
On this display, the labels and warning symbols are stereoscopically closer to you.
Is this effect noticeable to you?
Does it affect the display? Does it make it clearer?

C. Post-session interview

What do you think about the VR technology based on your experience today?
What was your experience with the tool like?
From your experience with it today, can you name the advantages and disadvantages you think the tool has over traditional visualisations?
What other features do you think the tool could include?
What could be improved in the tool?
Was the experience comfortable? If no, why not?
Was there anything in this setup which you found problematic, or would like to highlight?

D. Specific Questions

(Answered as one of strongly disagree, disagree, neutral, agree, strongly agree)

The tool was useful for looking at the big picture of data.

The tool was useful for looking at detail in data.

The tool was easy to use.

The tool was caused discomfort or was otherwise unpleasant to use.

The tool used stereoscopic techniques to increase my ability to interpret or understand data.

Stereoscopic techniques like those used in the tool have the potential to increase my ability to interpret or understand data.

I felt like I was in a different world while using the tool.