Integration of Low-Cost HMD Devices in Existing Simulation Infrastructure

Tomer J. Michael	Yaniv Minkov	Rami Rockah
IDF Ground Forces Command	IDF Ground Forces Command	IDF Ground Forces Command
Tel-Aviv, Israel	Tel-Aviv, Israel	Tel-Aviv, Israel
tomerjmichael@gmail.com	yanivminkov@gmail.com	rami.rockah@gmail.com

ABSTRACT

Recent years have seen a sharp increase in the availability of low cost Head Mounted Display Devices or, as they are colloquially referred to, "VR Goggles". These devices pair the live tracking of the orientation of their user's head, with a full stereoscopic view of a 3d environment. Thus, providing users with the illusion that they been transported to a virtual world where they are free look around in a realistic manner. It is this functionality that brought the HMD to the attention of the research simulation world, in particular because of the field's vested interest in providing its test subjects with the most realistic experience possible within a virtual environment.

However, the task of integrating the HMD presented a set of unique challenges. From the logistical, such as the lack of visibility between the human subject and input devices, to the physiological, such as the potential and prominent increase in so called "simulation sickness" (a subset of motion sickness), sometimes associated with even short encounters with the device. These phenomena raise questions in regards to the HMD's usefulness in research environments, where unintended side effects directly clash with the realism of the virtual environment and, by extension, with the validity of a given experiment's results.

This paper describes an attempt made between 2013 and 2014 to integrate Oculus VR's "Rift" HMD with the IDF Ground Forces Command Battle-Lab's existing simulation infrastructure. It discusses solutions and lessons learned for the integration of the device, technical hurdles encountered in making an HMD work with simulators built on existing frameworks like Vega Prime and Virtual Battlespace, and the application of different methodologies - explored for setting up or converting different simulators for use with an HMD - and their respective effectiveness with human participants.

ABOUT THE AUTHORS

TOMER J. MICHAEL Is a software developer currently serving with the IDF Ground Forces Command Battle-Lab, and is the developer in charge of integrating VR HMD suites with the lab's evolving infrastructure and systems. As a member of the Battle-Lab, Tomer specializes in and is involved with the integration of general purpose software API's and generic software libraries. He is an amateur game designer and theorist, and a member of the Israeli video game developer association - GameIS.

YANIV MINKOV is Human-Factors Engineer and military systems expert. Yaniv Holds B.Sc and M.Sc degrees from Ben-Gurion University of the Negev. Currently, Yaniv serves as head of the IDF Ground Forces Command Battle-Lab that is dedicated to the research of human-machine interactions in complex, multi-operator, military systems.

RAMI ROCKAH is an analyst and a military operations research expert. Rami Holds a B.Sc from Bar-Ilan University in applied mathematics and is working on his master's degree in operation research at the Tel-Aviv University School of Mathematics. Currently, Rami serves as head of section at the IDF-GFC Battle-Lab and leads experiments focused on future systems and concepts for the ground forces.

Integration of Low-Cost HMD Devices in Existing Simulation Infrastructure

Tomer J. Michael
IDF Ground Forces Command
Battle-Lab
Tel-Aviv, Israel
tomerjmichael@gmail.com

Yaniv Minkov
IDF Ground Forces Command
Battle-Lab
Tel-Aviv, Israel
yanivminkov@gmail.com

Rami Rockah
IDF Ground Forces Command
Battle-Lab
Tel-Aviv, Israel
rami.rockah@gmail.com

BACKGROUND

The IDF Ground Forces Command Battle-Laboratory

Israel's Ground Forces Command Battle Laboratory (Battle-Lab) specializes in the analysis of future military technologies through simulated application. Most often when considering the implications of a new piece of military equipment a Serious Game will be constructed from a network of several manned simulator stations, and staged Artificial Intelligence entities. Together the networked elements present a virtual facsimile of a typical combat scenario. A simulated form of the new technology will then be introduced into the network (either directly operated by the appropriate soldiers or embedded into one or more simulators), so that a set of scenarios (both with and without the technology) can then be recorded for further analysis.

A proposed technology does not necessarily have to be physically manifested before it is simulated at the Battle-Lab. Insight is gained from exploring recorded simulations of its use to help steer its early development, where making changes and adjustments to its design (or even dropping it wholesale) is relatively inexpensive.

Human participants stand at the heart of the Battle-Lab's Serious Games, as their interactions with each other and with the simulated technology create the content from which a recorded scenario is formed. The closer these interactions are to the way they would be in reality, the closer the recorded scenario should - in theory - correlate with reality, directly influencing the validity of any conclusions derived from the game.

The goal of most of the Battle-Lab's Serious Games is to provide tactically relevant data for analysis, rather than to perfectly simulate the physical "performance" of combat for a scenario participant. Therefore, a series of abstractions is required to facilitate a participant's interaction with the virtual world of the scenario. The participant does not really move his/her feet to go forward, instead he/she presses the forward key. Likewise, he/she does not perish when his/her avatar does, and instead his/her screen simply goes black and he/she will be politely asked to exit the room.

A *"Ludic"* Interface is usually defined as either a novel or a playful method for interacting with technology. As opposed to the traditional "mouse and keyboard", or "gamepad" interfaces, it is the very way in which a Ludic Interface is used which emulates action being performed within the simulation. Perhaps one of the most prevalent ludic interfaces in the simulation world today is the steering wheel controller, which allows users to "perform" the action of steering a car, rather than abstracting it through a series of keyboard commands.

The Battle-Lab is no stranger to the integration of Ludic Interfaces in its simulator infrastructure, and has always sought to further its integration with new technologies, both as a temporary elements of a specific experiment, or as permanent fixtures in the Lab's growing repertoire of tools.

THE HEAD MOUNTED DISPLAY (HMD)

In 1957 Morton Heilig filed a patent for his "Stereoscopic-television apparatus for individual use" (U.S. Patent No. 2,955,156, 1960). Mounted over the eyes and fastened to the user's head with a belt, Heilig's device would separately project two simultaneously captured images to each one of the user's respective eyes, effectively forming a fully three dimensional or "stereoscopic" image. As the images shown in the device were to be prerecorded on film, the notion of tracking the user's head and projecting an image appropriate for its current orientation would have been neigh impossible. Practically speaking, Heilig's device would have been an early Head Mounted Display (HMD) device.

The modern HMD, however, has probably more to do with Ivan Sutherland's Head-Mounted Three-Dimensional Displays (Sutherland 1968). Sutherland's HMD was created as part of a research project for Harvard University, starting in 1966 with funding and support from the Department of Defense (DoD), the Naval Research Center (NRC), and Bell Laboratories. Although his device was rather large and cumbersome, Sutherland (1968) (as opposed to Heilig) used dynamic computer generated imagery (basic three dimensional wireframe designs) as his primary image source, which allowed his device to change the angle at which the simulation was rendered in relation to movements of a user's head. These were collected through a series of sensors, creating for the user the illusion that the image was moving along with him.

Perhaps with the subtraction of the large array of machinery and computers necessary to create Sutherland's device, the modern HMD has not changed much. Today's devices consist of a display, worn on the head, which usually provides the user with a stereoscopic image. Often when used in Simulation and Video Games, HMDs will also make use of sensors to track a user's head, so that the application will be able to provide the display with an image that correlates in the virtual environment to where the user's head is located in the real world.

For many years, HMDs either lacked the technological capability to be relevant for high fidelity simulation, or have been too prohibitively expensive or complex for most organizations to adopt and iterate on effectively. Only recently, with the introduction of Oculus VR's "Rift", Sony's expected project Morpheus, Avegant's "Glyph", and several other Augmented and Virtual reality solutions, has the technology begun to shift towards widespread circulation.

Most if not all of these new devices are consumer products, designed for use with video-games, or as multimedia platforms. Few of the new devices, if any, specifically target military systems and simulations. However, if they were to prove beneficial in military environments, they would definitely not be the first consumer technology of its kind to have made the leap towards military applications.

Given its relatively open set of tools, developer focused community and demonstrated capability, the Battle-Lab decided to attempt integrating Oculus VR's "Rift" with its simulation infrastructure. The integration was conducted between late 2013 and early 2014, and resulted in the integration of the Rift both as a toggleable part of the Battle-Lab's own in-house infantry simulator, constructed with Presagis' "Vega Prime 5", and as a plugin for simulator stations running Bohemia Interactive Simulations' "Virtual Battlespace 2" (VBS2).

THE INTEGRATION

Initial Experiences

In order to better inform the Integration process, it was decided to first examine a number of existing Rift compliant applications, and then use them to attempt to establish a baseline of the device's capabilities.

The focus group for this stage of the process consisted of ten members of the Battle-Lab's development team, not otherwise directly involved with the Rift's integration. Information gathered at this stage amounted to a qualitative feedback, gained from exposure to the device over varying stretches of time, using different VR applications.

The most striking phenomena immediately noted was the magnitude and recurrence of simulation sickness among participants (a plight effecting users of the device, described in full in the Rift's own best practices documentation) (Yao et al., 2014). Nearly all users experienced varying levels of nausea and disorientation when first interacting with the device. However, most (but not all) had developed a certain resilience to it by their second or third use of the device.

Secondly, it was noted that while using the Rift's configuration utility had not significantly increased or decreased the feeling of disorientation for most users, the act of simply wearing the Rift first for a total of about five minutes, while going through the configuration process, helped to better prepare users for wearing the device for an extended period of time.

The use of the interpupillary distance measurement suite, in particular, helped users both to accept the Rift as a simple screen first before being "bombarded" by the allure of its motion tracking. It also helped them to get used to exploring the screen space by moving their eyes to read on-screen cues (a skill that would later prove invaluable for reading information from a head's up display). Surprisingly, we discovered that at this stage, most users could not tell to which eye an image was projected without closing one eye after the other to see when it disappeared, a realization which proved invaluable later on.

Additionally, while comparing the "Tuscany" simulation from Oculus VR's own website (a simple first person simulation of a villa in the Tuscan country side with a large, open garden and an ocean view), and another freely available simulation of a small New York apartment ("Jerry's Place" based on the fictional apartment of television's Jerry Seinfeld), it became apparent that for their initial experience, users were somewhat more comfortable and less disoriented in the more confined, less brightly lit stage of the New York apartment. On the other hand, some described the experience as claustrophobic, perhaps due to the combination of the apartment's lack of an outdoor view, the reduced comfort induced by wearing the device, and the Rift’s slightly reduced field of view. The lack of ambient noise in the apartment simulation was also proven to be rather jarring. This would prove to be a recurring complaint with other simulators and applications later on.

Lastly, all demos examined bound movement direction to head direction together, something that some users noted as feeling "unnatural" to them, although others indicated that it helped them stay oriented.

First prototype in VEGA PRIME

Concurrent with the Rift's arrival at the Battle-Lab, several new in-house combat simulators were in development with Presagis' "Vega Prime 5", a simulator visualization toolkit and code library. As it was desired that the Rift be as portable as it could be between the simulators, it was decided to try and create wrappings for the Rift, based on Vega Prime's own framework.

Rendering to the Rift requires displaying the rendered scene once from the perspective of each eye, and then displaying the output side by side on its screen and re-centering the images so that their centers match the location of the user's respective pupils. In order to account for the distortion caused by viewing the image through the Rift's lenses, a "reverse" distortion must then be applied to the image through software,

Figure 1. Distortion Example

The Lab was fortunate enough to receive code from constituents at the IAF who had previously experimented with using the Rift in a Vega Prime environment, and built a "module" that could be installed and used in "Vega Prime 4" to both create the proper "cameras" needed for Rift viewing, and to distort the image appropriately. As the module was not compatible with Vega Prime 5, and the Battle-Lab had made a conscious decision to avoid the development of new "explicit" modules for Vega Prime, development shifted towards recycling its source into a C++ library.

The classes created were curated through a master globally exclusive object, named "vpOVR". vpOVR initialized the Rift's own library driver, and then wrapped its provided data classes with logically appropriate Vega Prime classes. vpOVR would then create a window to be displayed on the screen of the currently attached Rift, which was then fed visual information from a "camera" (or in Vega Prime's own Lingo, an observer) placed within the virtual space on an axis (a Vega Prime "Tranform") oriented according to information gathered from the Rift's sensors. This axis was then re-parented to an object within the three dimensional space of the simulation so that its orientation and position would be determined relatively to its parent, effectively forming both a body (the parent object) and a head (the axis). Finally, using the code salvaged from the Vega 4 module, the application was rendered to the window appropriately set up for use with the Rift via information gathered from its own configuration utilities.

As work had not yet been completed on an either an infantry or vehicle simulator by the time the library was to be tested, the initial prototype was instead tied to a debug observer object, that could either float freely through the simulated world, or move freely at a fixed distance from the terrain below it. This was mostly intended as a tester for gauging how the simulators would look through the Rift's lenses, but following the earlier experience with motion sickness, quickly developed into a test bed for different use practices with the device.

The experience manifested that body movement could be de-coupled effectively from head direction, as long as head direction remained relative to it, and that an indicator of some sort (either a vehicle "body" or, as was later used, an infantry reticle) would consistently point in the direction of motion.

Integration Design

Following both initial experiences, it was finally decided to attempt introducing the Rift into the Lab's infantry simulation system.

A typical Lab designed infantry simulator will behave similarly to a modern "First Person Shooter" video game. A simulated soldier's compass direction, eyepoint direction, and weapon direction will all be tied to the same set of heading and pitch settings and controlled via mouse movement while forward movement, strafing, reverse motion, sprinting, pose control and weapon selection correspond to a selection of configurable keys on the keyboard. Usually, a Lab simulator will not let the participant fire his weapon unless he is currently aiming through its sights. This action is performed by pressing a button, at which point an overlay abstracting a weapon's sights will appear, and the participant will be able to fire by pressing the left mouse button.

One historical problem the Lab would often encounter with this design was the strong binding of the eyepoint, aiming and movement directions. As when bound together, these do not effectively allow a combatant to scan his environment for threats, better orient himself while moving in certain direction, or aim his firearm at targets.

However, with the use of an HMD, the Battle-Lab saw an opportunity to unbind the eyepoint direction from the other components, and match it to the orientation of the current participant's head. This direction would be relative to the simulated soldiers body and movement directions, as had been in the initial Vega Prime prototype.

In the interest of creating as modular a solution as possible, the Rift enabled simulator was be developed to be as similarly controlled to the vanilla simulator as possible, maintaining all keyboard and mouse bindings with the above exception notwithstanding.

One of the biggest challenges involved with integrating the Rift with the Battle-Lab's systems had been (and still is) the question of responsibility. As it was integrated, the Rift HMD was a first generation development kit, and a common concern was whether or not particular negative aspects of using the device were the result of its own technical shortcomings. On the other hand, conversely had the software developed at the Lab been at fault for overlooking certain integration tasks? And even if the device's technical shortcomings had truly been at fault, should complex attempts be made to try and Jury Rig a solution? Further weighing in on the latter question was the awareness that this being a development model, the device's capabilities were likely to change, and work made towards mitigating certain issues may end up being unnecessary in the long run.

For example, the first generation Rift suffered from a relatively low pixel density, which increased aliasing significantly, while its touted second generation replacement has reportedly double the resolution, bringing into question the importance of anti-aliasing measures for the first generation device.

Additionally there was the issue of latency. The time period between when motion is registered by the Rift's sensors, and when an image is displayed on screen correlating to that motion is described as the device's latency. To overcome this issue the Rift can use a prediction mechanism, which requires input based on the device's predicted latency. At the time of our examination there existed a separate tool (created by Oculus VR) designed to test latency, however the second generation kit was reported to have this sensor built in, and it was decided at the lab to table the issue until the newer devices arrival.

Finally, were comfort issues with the device itself, its lack of counterbalance for its screen, and the discomfort some users express with having one so close to their faces.

At this point it is important to note that it is always a possibility for the Rift to eventually be found inadequate for our, and another HMD solution could possibly be chosen for implementation at the lab. With that in mind, it was important for code written to be as lightly coupled with the specific device as possible, to make room for the possibility of replacing it later on, and making Rift specific "hotfixes" a dicey proposition.

The decision was made to try and fix any glaring problems as directly as possible. However, as the current implementation was designed for evaluation, areas which were expected to be resolved in later versions of the Rift's development kit, or in other HMD solutions, were given appropriately simple "hotfixes", instead of going through the process of creating entire replacement systems. Additionally, a significant lack of automation was noted to exist with some of the Rift's earlier unofficial utilities, particularly in the field of the sharing configuration information. In these cases manual reconfiguration was preferred, as a large scale multi-user implementation was not yet planned.

Integration in Virtual Battlespace 2

While development of Rift integration was under-way with Vega Prime, the opportunity presented itself to try and make the device work with *Bohemia Interactive Simulations*' *Virtual Battlespace 2* (VBS2). VBS2 was a second simulation platform also occasionally used in addition to, and sometimes in conjuncture with, Vega Prime at the Battle-Lab. Once an architectural design was conceptualized for its integration with VBS2, the process proved relatively straightforward and reliable. This architecture consisted of three main software Elements.

Firstly, a plugin dll written in C++ was developed for use through with VBS through its integral "Application Scripting Interface" which connected with the Oculus Software Development kit (LibOVR) and passed relevant information (including current orientation) about the currently connected HMD to the VBS application via a globally referenced array accessible within VBS via script.

Secondly, a script was designed to work within VBS2 to interface with the dll. First the script would disable mouse controlled head motion by calling "player setDisableGunnerInput 2", then it would initialize the aforementioned dll, and finally it would loop setting the orientation of the player's head according to information gathered from the Rift through the dll.

Due to a desire to conform with VBS2's existing interface for the reorientation of the player unit's head, the Head orientation was set relative to an axis consistent with the body’s compass direction alone, and perpendicular to the player unit's body, whilst the players weapon could be oriented vertically independently of this axis.

Lastly, while the idea of using VBS2's premium *Fusion API* to gain access to its *Direct 3d* calls and distort the image appropriately for use with the Rift screen was initially explored, it was soon abandoned. Not only was the study of the API beyond the scope of the Rift integration project, but it was also discovered that trying to distort the image naturally with callbacks provided by the Fusion API, whilst maintaining VBS2's two dimensional interface (such as its heads up display) was both cumbersome, inefficient, and would have been difficult to maintain for anyone without some sort of understanding of Direct 3d

Instead it was decided to use the Open-Source "VIREIO Perception" direct 3d proxy, an application developed for forcing backwards compatibility for video games which were not originally VR compatible, with embedded support for the Rift Dev-Kit 1 device (Meant to be Seen, n.d.). VIREIO masks the application's natural calls to direct 3d with modified ones, which already take care of the side by side and distortion rendering necessary to create stereoscopic Rift compatible images.

In addition to relieving the Battle-Lab integration of the need to develop the distortion tool for VBS2, and being externally supported by its own community, being that VIREO links to an application mostly through its calls to direct 3d, it did not require an additional fusion license to use with VBS2.

However, VIREIO was not designed with VBS2 specifically in mind, and is still very much a piece of beta software. Associated with it were a number of odd phenomenon, most notable of which was its incompatibility with mouse input in VBS2's two dimensional menus (overcome through the applications ability to load directly into an in-game session via command line). Additionally, VIREIO's use would also sometimes result an odd assortment of rare but noticeable visual artifacts (such as the occasional white flicker of the screen every few minutes). It is also technically limited in that in its popularly compiled form it only works with 32bit applications, and that at the time only supported direct 3d 9 applications, and not the newer direct 3d 10, or direct 3d 11 distributions.

As future versions of VBS2 are expected to support the Rift natively, and because VIREO has a number of commercial counterparts which do currently support these later distributions, it was decided that for the time being that developing a distortion handling plugin for VBS2 was simply not a necessity.

Additionally, as the Rift might not be chosen for use in the lab, the choice of an entirely external application had the advantage of being modular enough to eventually be switched with implementations for other HMDs.

At this point VBS2's UI was reconfigured to appear mostly around the center of the user's screen, instead of at its edges where it is obscured by distortion, and its field of view settings were adjusted to match the Rift's capabilities.

While developing the VBS2 implementation it was discovered that a closer approximation of aiming a weapon through its sights could be achieved than originally expected. By allowing the user a limited area around the middle of the screen where mouse movements would only be bound to weapon motion, the user could move the weapon's sights so that they aligned not with the center of his/her head or body direction, but rather with *one* of his/her pupils. This process allows him/her to aim through one eye and close the other, as he/she would in real life. Within VBS2 this style of loosely controlled weapon and body control is referred to as "floating zone" controls, and can be toggled from within an application's settings.

The end result was a simulator which behaved similarly to that which is described above, perhaps except the aforementioned loosening of the ties between weapon and body direction. Additionally, due to a quirk of VBS2's head orientation set command, head direction remained relative only to a user's heading but not to his body's (or his weapon's) pitch.

Final Integration with the Battle-Lab Infantry Simulator

Informed by the successful experience with VBS2, an attempt was finally made to integrate the Rift with the Battle-Lab's own in-house Infantry simulation.

First, the old vpOVR class was initiated as it had been in the past, and initialized both the Rift's sensors and their Vega wrappers, except that now the old "render to separate window function", and Rift distortion rendering functions carried over from the Vega Prime 4 code, had been deprecated, and were no longer in use.

Instead, a simple side by side view of the world was created. First the application would iterate through the list of all viewports and overlays (Vega Prime's "channels") attached to the application window. Then vpOVR would resize them to the confines of its right half of the window, and create a new viewport to fill the now empty left side. The new viewport was then configured with information from what had previously been the viewport representing the simulated soldier's own view, which was now one of those covering the window's right half.

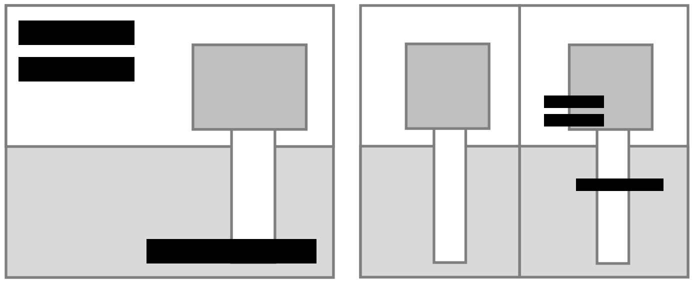

Figure 2. On the left, caricature of unaltered Vega view, on the right, adjusted Battle-Lab HMD view

Both the new viewport, and the original one (herein referred to as the left and right channels respectively) were then assigned to display the view generated by a shared camera object. Through Vega Prime's offset functionality for channels, they were set to display an image, sourced on opposite sides of the camera, which was exactly half as far away from the camera as the Rift reported was the distance between both of the user's pupils (his interpupillary distance, or IPD).

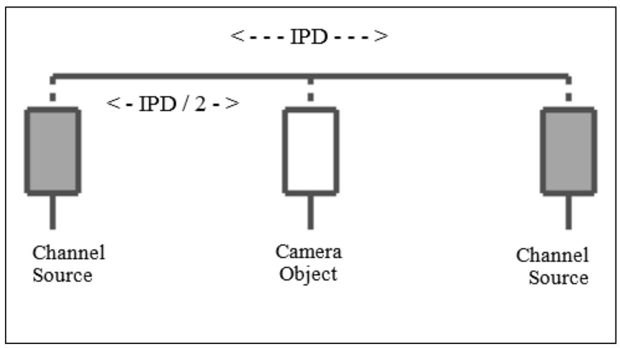

Figure 3. Vega Channel offsets

Otherwise, the natural freelook camera behaved almost exactly like it's counterpart in the earlier Vega Prime prototype, with the exception that its axis was now situated on the abstract representation of a three dimensional soldier, rather than on a debug observer.

As the existing simulator had not made the soldier's body visible to himself, and as he could only see his weapon when aiming through sights, it was decided to communicate current body / movement direction to the user via a two dimensional overlay reticle. To achieve this the overlay was set to track an invisible object parented to the user's virtual body so that it would always be located 20 meters ahead in its current facing direction.

An attempt was also made to make vpOVR's use of the Rift's particular SDK as opaque as possible, to better facilitate the potential integration of other HMDs into the system.

Among the hurdles encountered in integrating with the Vega Prime simulator was the issue of weapon rendering. At the time of integration the non-"Rift Enabled" simulator would use two dimensional sprites overlaid as HUD elements to represent a participant's currently held armament. While this solution had been admissible when viewed through a 2D monitor, the Rift's stereoscopic rendering managed to break the illusion. In 3D ethe 2D sprite would have to be placed somewhere within the space of the scene so that it could be appropriately resolved by the eyes and used to aim. Hypothetically, were we to follow this path, the weapon would appear to the player as a "cardboard cutout" floating at a distance, breaking his immersion.

A compromise was achieved through a quirk of the simulator, specifically that it was not designed to handle fire or show the weapon sprite while not aiming through a weapon's sights. Given the assumption that one aims through a weapon's sights with one eye closed, and that when aiming both head and weapon motion would be tied together, the decision was made to exclusively render to the eye used for aiming and to restore simulator input to the state where both head orientation, weapon direction, and body direction are bound together to mouse movement. In effect when aiming the simulator behaves exactly as it did before the Rift was enabled, except that everything is rendered exclusively to the half of the screen used by the aiming eye.

Two alternatives were considered for the sake of distortion rendering in Vega Prime. The first involved using the two eye channels as targets from which a texture would be derived, then distorted with a modified CG shader through Vega Prime's shader classes, and then rendered itself to another channel, atop which the 2D interface channels were placed. The second approach considered called for the channels to be rendered side by side undistorted, and for a second application to capture the simulator's output and itself handle distorting the image and displaying it on the Rift's screen. Ultimately the latter was chosen, as the former required several assumptions to be made about the specific simulator configuration to be hard coded, whereas simply creating two "side-by-side" channels (as described above) could easily be handled via simpler coding interfaces.

With the exception of the aforementioned changes to weapon handling, the "Rift enabled" version of the Lab's infantry simulator remained mostly true to the initially discussed integration design.

OUTCOME

Live Evaluation with VBS2

Over the course of a week in mid-2014 an evaluation was held at the Battle-Lab of the Rift-Integrated VBS2 simulator. The evaluation was intended to try and quantify whether or not the Rift's use would have a detrimental effect on a soldier's performance in a virtual scenario, and to qualitatively make observations of its effectiveness. For the sake of the evaluation, a pair of soldiers seated next to each other would be requested to cooperate in one of three predefined 5-15 minute combat scenarios. During these scenarios one soldier was to wear the Rift, and the other participated through a mostly unmodified three screen VBS2 infantry simulator. Upon the scenario's completion the soldiers would then be requested to switch, and a different scenario would be loaded for them to engage in.

At the end of each scenario both soldiers were requested to answer a series of questions about their well-being and their opinions of how well certain aspects of the scenario went. These were mostly about how well they felt they had communicated during the exercise, how disoriented (if at all) they felt during the exercise, and how much (if at all) nauseous or uncomfortable had they felt while wearing the Rift.

A conscious choice was made not to use the "SSQ" Simulator Sickness Questionnaire (Kennedy, Lane, Berbaum & Lilienthal, 1993), both due to the relatively long overhead time associated with its completion, and the fact that while *the presence* of motion sickness was to be examined as part of the evaluation, its content was beyond its parameters.

Additionally, other than by playing with lens choice to better accommodate users with glasses, no other personal configuration was made to the device during the evaluations course, and a default interpupillary distance of 6.4 cm was used. This was done with the intention of replicating a worst use case scenario, where personal device configuration become impossible due to a high volume of rotating participants.

Overall, a preliminary examination of the evaluation showed that on average, participants only felt a minimal amount of nausea or discomfort while wearing the device for the first time, with 7 of the 13 participants reporting no nausea or discomfort at all on their first use of the device. Only three participants reported feeling significant discomfort. When they were asked to rate the discomfort on a scale of 1-10 (with 1 representative of "not at all nauseous" and 10 representing "extremely nauseous"), the responses were 6.5, 5 and 4.5 respectively. Many participants reported an increase in their drive to complete the scenario successfully, as well as in their sense of immersion in the virtual environment.

Of the 20 soldiers who had participated both in the monitored evaluation, and in the play-testing of its scenarios, 85% reported that they felt some improvement to their understanding of the virtual space.

However, just prior to the evaluation, while play testing scenarios with a soldier who was not part of the final evaluation, it was found the he could not at all bear wearing the Rift, and rejected it immediately upon putting it on. It should be noted that none of the participants in the evaluation itself manifested this sort of response, and none of them removed the Rift prematurely before the scenario's completion, so it would seem that this may be a rarer, but still possible response to the device.

Early play-testing with Vega Prime

At the time of writing, extensive play-tests with the Rift enabled version of the Battle-Lab's infantry simulator have not yet been performed. On the other hand, early play testing has begun, and (similarly to our initial experiences with the device) consisted mostly of various sessions with Lab personnel.

One odd feature of the Battle-Lab's infantry simulator is that due to the nature of its collision detection system, its frame rate while in motion drops from a smooth 50hz to just under 20hz. This was initially theorized to be detrimental to the use of the Rift, as even its own documentation reports a high frame rate to help prevent simulation sickness (Yao et al., 2014). However in early testing people tended to move their heads the most while motionless, making the drop in frame-rate less noticeable. Additionally, due to the fact that data was collected from the Rift's sensors only after the collision detection system finished its function, the overall data reported and used to render the viewport corresponded relatively well to the orientation of the device at the point in time at which it appeared on the screen. This left only large gaps of frozen, yet somewhat accurate images, instead of consistently highly latent ones.

The users also noted that the experience had been slightly less disorienting then the previously constructed prototype. This is most likely due to its nature as an actual simulator, a representation of something real, rather than the unnatural "floating unsupported in midair" of the early application.

CONCLUSIONS AND LESSONS LEARNED

As a Ludic Interface

Perhaps the greatest parameter by which we can measure an HMD’s success is whether or not it succeeds as a Ludic Interface, whether or not it succeeds in emulating an action better than a potential previously used abstraction.

From a technical standpoint the Rift (when faced with an unhampered user) does seem to significantly increase a user's immersion in the world of the Serious Game he/she is participating in, as well as to make the world around him/her more accessible for visual analysis, particularly while in motion.

Perceptually the avatar becomes more of a virtual skin rather than a far removed object in a computer screen. And as such there appears to be a certain emergent sense of personal investment in the scenarios users participate in.

This can present us with another set of problems, as the Rift significantly shifts the balance of a simulation's narrative, from what is explicitly told to the participant about what he/she is doing, to what it is implied he/she is doing through the simulated environment's behavior. This implicit story, or "story told by play" in the Serious Game environment is often referred to its Ludonarrative, and when it clashes with the explicit narrative a user is told he is supposed to be conforming to, the effect can be described as "Ludonarrative Dissonance" (Clint Hocking, 2007).

This sort of dissonance often results in users being forced to explicitly choose a narrative with which to frame their experience. Often they would either revert to experiencing the simulation as a video-game, or get stuck trying to understand the relation between the abstractions of the simulated environment and reality in extraordinary detail.

While it appeared from the Lab's experience that the Rift helped enforce the implicit story of play, it did not enforce the explicit one, suggesting that further steps may need to be taken to help experiment conductors communicate with participants through the Rift's own interface. The use of on screen textual communication, or an integrated headphone PA system come to mind, but have not yet been tested at the Battle-Lab.

Simulation Sickness

The prevalence of Simulation Sickness associated with the Rift's use, remains perhaps the Lab's top concern when considering it and other Head Mounted Displays for integration.

The Battle-Lab environment requires a participant to be seated at a simulator station for multiple sessions ranging from between 60 – 100 minutes in length, over the course of several days. Participants are expected to play their actual roles within a fighting force, and the loss of even one soldier to Simulation Sickness, particularly an officer or commander, could potentially "unbalance" the force.

A trend of slow adaptation to the device was apparent throughout the Battle-Lab's development team, and it is believed that perhaps the best way to get over simulation sickness may be repeated or long term exposure. What remain then are the outlier cases, those people who cannot bear the device for even the shortest periods of time. Further examination is necessary to see if further exposure to the device, or any other means, might be able to help, or if perhaps it is the finer tuning of the device which is necessary.

Future Work

The Battle-Lab's initial experience with the Rift has been overall educational. While the device's use may not have yet been proven as either entirely positive or negative, the initial evaluation has shown that the device is at the very least practical for use, and functions effectively.

What remains to be seen if how effective a group is when functioning together whilst wearing the device, as well as for extended periods of time. The simulators themselves, although functional, are still limited by a cumbersome initialization sequence, which would have to be streamlined in order to make the Rift practical for large groups of people. One could further imagine a situation where a large group is used to identify the aforementioned outliers, and their presence could be used to measure the effectiveness of various techniques designed to reduce the potency of simulation sickness.

Additional play testing will also have to be performed to explore the effects of HMD use with vehicle broadcasts over a network. Network communications are not always perfect and are effected by jumps and lags. This could potentially be jarring to a user "within" the vehicle already compromised by wearing an HMD.

It is also important to note that although the Rift may have been seen as representative of a new wave of HMDs, it is not without competition, and the examination of its competitors may be an important step in solving the earlier mentioned questions of responsibility, and to ascertain which issues might be exclusive to the Rift device or which are shared by the range of HMD devices.

Finally, embedding the Rift highlighted a number of key points related to the integration process itself. The main points being the value of testing the device as early and as often as possible and the need to identify and solve the causes of simulator sickness. The use of other applications designed for use with the Rift had proven very useful in forming a baseline understanding of its abilities, and the creation of an early test bed had allowed us to quickly iterate through implementation ideas.

ACKNOWLEDGEMENTS

The authors of this paper would like to thank all members of the IDF Ground Forces Command Battle-Lab for their support in the writing of this paper (and in particular Arthur Bendersky), and for their logistical as well as technical assistance. We would also like to thank Avshalom Gil-Ad, and Alberto Saul of the IDF IAF HFE Lab for providing us assistance as well as the early code around which we developed the early Vega Prime Oculus application.

REFERENCES

Heilig, M.L. (1960). *U.S. Patent No. 2,955,156*. Washington, DC: U.S. Patent and Trademark Office.

Hocking, C. (2007). Ludonarrative Dissonance in Bioshock. *clicknothing*. Retrieved from http://clicknothing.typepad.com/click_nothing/2007/10/ludonarrative-d.html

Hocking Kennedy, R. S., Lane, N. E., Berbaum, K. S., & Lilienthal, M. G. (1993). Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The International Journal of Aviation Psychology*, 3(3), 203-220.

Meant to be Seen. (n.d.). Vireio Perception. Retrieved 8 June 2014, from http://ww.vireio.com

Sutherland, I.E. (1968). A Head-Mounted Three-Dimensional Display. *Proceedings of AFIPS 68*, 757-764.

Yao, R., Heath, T., Davies, A., Forsyth, T., Mitchell, N., & Hoberman, P. (2014). Oculus VR Best Practices Guide (v0.008). Retrieved from http://static.oculusvr.com/sdk-downloads/documents/OculusBestPractices.pdf