White Paper

for

Virtual Control Room

Proof of Concept Phase

October 2015

Project Members: William Little, Benjamin Tully-Hanson

Table of Contents

A. Executive Summary

The Virtual Control Room (VCR) Proof of Concept (PoC) project is the result of an award given by the Fourth Annual NASA T&I Labs Challenge Project Call. This paper will outline the work done over the award period to build and enhance the capabilities of the Augmented/Virtual Reality (AVR) Lab at NASA's Kennedy Space Center (KSC) to create the VCR.

A-1 Overview

A number of new computer technologies reaching the commercial market today are making the fields of Augmented Reality (AR), Virtual Reality (VR), and Natural User Interface (NUI) accessible to the public. These technologies represent the next phase of the computer revolution by fundamentally changing the nature of human-computer interaction (HCI) away from the traditional keyboard-video-mouse (KVM) interface to one employing 3D head mounted displays and motion tracking systems, to mention a few. Computer game oriented devices such as the Microsoft Kinect, the Leap Motion, and the Oculus Rift carry the potential for significantly altering the way humans interact with both the virtual world and the physical world. These new technologies are not yet mature; significant study needs to be done to understand them, to quantify their strengths and weaknesses, and to research solutions to improve the accuracy and reliability of the data they generate. Prototype systems incorporating these devices need to be built and tested to determine how best to use the technology in control rooms, firing rooms, and research labs of the future.

A-2 Problem Statement

Given recent advances in computer input/output (I/O) devices, the very nature of how humans and computers communicate with one another is radically changing. Evidence of this change can be found in the definition of the acronym HCI. The acronym once meant human-computer interface, implying a boundary that the human needed to learn how to cross in order to understand how the machine works. It is evolving into human-computer interaction, implying a deeper level of communication in which the responsibility for understanding how each side works is shared between the human and the computer. As the computer industry continues to develop and improve on AR/VR/NUI devices, reliance on traditional keyboard/video/mouse (KVM) technologies will diminish. As seen with a number of legacy control systems at KSC, aging architectures will become increasingly difficult to maintain and repair. The "adapt or die" reality of rapidly changing technologies will continue to assert itself. It is essential that we continue researching and refining new AR/VR/NUI technologies and their application to our mission in order to develop innovative solutions to future problems.

A-3 Background

At the inaugural 2012 KSC Innovation Day Kickstarter competition, an award was given to investigate the use of the Microsoft Kinect as an I/O device for use in controlling an external device, in this case a Lego Mindstorms robot. The result of the project was the Virtual Control Panel (VCP). Using the Kinect to monitor user movements, a Sony HMZ-T1 head mounted display (HMD), and software written to integrate the Kinect data with a virtual display, it was shown that an operator could control the movements of a robotic device using gestures alone, without the use of KVM devices (Figure 1).

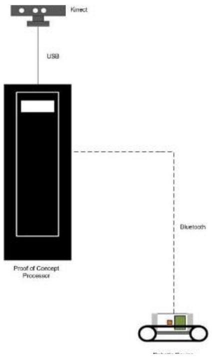

Figure 1: Initial configuration

One immediate finding of the VCP project dealt with the accuracy and reliability of the data generated by the Kinect. The Kinect tracks human motions using an infrared (IR) emitter, an IR sensor, and onboard software that produces a set of points that represent skeletal joints on the human body in 3D. Thus, the location and motions of the operator can be tracked in real time. Three major problems became quickly apparent: joint occlusion, bone length, and jitter. Joint occlusion occurs when the user moves into a position in which a joint is hidden from the Kinect's view. Onboard software extrapolates the joint's location from its last known position, guessing at where the joint currently is. It invariably guesses wrong, generating inaccurate data, which causes a distortion of the virtual skeleton. In the physical world, bone length, the distance between two adjacent joints, is constant. The Kinect's software has a problem keeping the distance between joints consistent, so as the user moves an arm or a leg, the length of a bone may visibly change in the generated skeleton. Jitter occurs when any two nonadjacent joints are brought into close proximity to one another, causing rapid fluctuations in the generated data for those joints. For example, when hands are kept apart, jitter is at a minimum. When the hands are brought close together, the system has difficulty differentiating which joint is which, and jitter occurs.

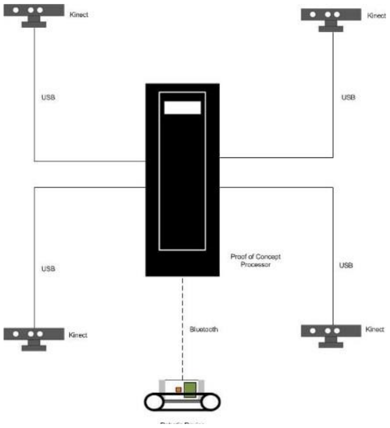

Figure 2: Multiple Kinect configuration

Resolving these issues began with a survey of academic and industry white papers on the subject. This review revealed a consensus: The accuracy/reliability issue of Kinect skeleton data cannot be resolved via use of a single Kinect. Multiple Kinects generating multiple sets of skeleton data are required (Figure 2). Software to resolve the data streams into one integrated skeleton is also required. While the multiple Kinect approach seems to be the correct way to handle these problems, the solution brings with it a new issue. The Kinect is a USB device, and due to data bandwidth limitations, it is inadvisable to connect multiple Kinects to a computer on one USB bus. Each Kinect requires its own unique USB bus. Limitations on the number of PCIe card slots on a computer motherboard limits the number of Kinects that can be hooked up. If multiple Kinects are inputting data to a single processor, processor capabilities to handle the amount of incoming data are taxed, and latency becomes a serious issue. The resolution to this issue is to expand the system architecture to include dedicated processors to input each Kinect's data and ship it out to a host processor via Ethernet, where the host handles the task of using the data to generate the final skeleton, as shown in Figure 3.

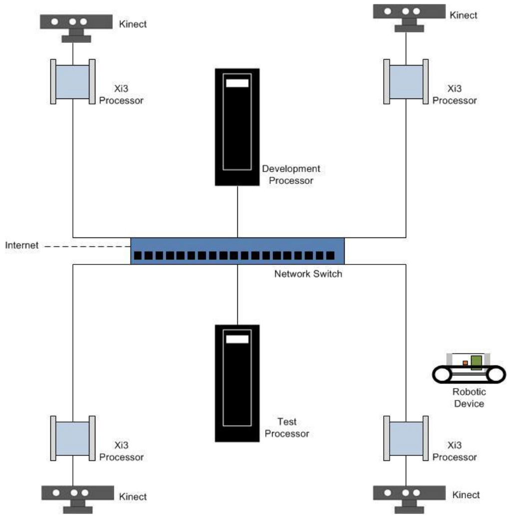

Figure 3: Networked configuration

Another finding concerns the level of detail of the skeleton points generated by the Kinect. The Version 1 model generates a total of 20 points. Each hand is represented by one point, thus making the hand a club at the end of the arm. No tracking of individual finger movements is possible. Therefore, it became necessary to add the Leap Motion device to track the user's finger motions. The Leap Motion uses the same IR emitter/sensor technology as the Kinect to track hand and finger movements; it can therefore better determine the exact locations and motions of individual fingers. It does however suffer from some of the same joint occlusion and jitter problems as the Kinect, to a lesser degree. Resolving this issue will be the subject of future investigations.

Lastly, the choice of the Sony HMZ-T1 was not the best possible for the task. This HMD is actually little more than an extension of a video monitor, offering no rotational/positional tracking of the operator's head, and no true support for 3D visualization. Also, the device is poorly balanced and will incur neck strain on a user who attempts to wear it for an extended period of time. These problems were mitigated by the introduction of the Oculus Rift Development Kit 2 (DK2), which solves all the problems mentioned, and more. In addition, the Oculus offers a mounting bracket for the Leap Motion, which allows the user to track hand and finger motions from the user's viewpoint.

A-4 Conclusion

The progress made on resolving the joint occlusion/bone length/jitter issue was a major step ahead in this project. A full resolution is within our grasp. Upgrading the local processors dedicated to handling Kinect data streams allowed us to replace the Kinect V.1's with Kinect V.2's. The inclusion of the Unity software tool vastly improved our ability to create 3D VR environments. The switch to the Oculus Rift was a huge improvement over the Sony HMZ-T1. It is a true VR HMD, it integrates seamlessly with

Unity, and it incorporates head position and motion tracking features. Adding the Leap Motion to the system was another significant improvement. The device takes up where the Kinect leaves off, allowing tracking of hand positions and motions down to the individual finger level. Lastly, upgrading from the Lego Mindstorm robot to a Makeblock robot proved beneficial, as the Makeblock is more responsive to commands sent via Bluetooth.

Figure 4: The AVR Lab

B. Use Cases

Each of the technologies outlined in this paper have been used primarily in the computer gaming field. Microsoft originally designed the Kinect to go with its Xbox gaming console. It was after the user community “hacked” into the device’s software that Microsoft issued the software development kit needed for programmers to start working with the device for purposes other than gaming. The Oculus Rift is, as of this writing, not yet available to the public; however, the DK2 has been made available to the developer community for software development and testing. The predominant use of the device has so far been in the gaming realm.

NASA can use these technologies for a number of purposes, from collaborative engineering design to virtual training to command and control of external robotic devices and systems. Consider a spacecraft in its design phase. Scientists and engineers could work together, using 3D models and simulations to determine system design and function by viewing and manipulating the models so that optimal size and placement of components could be worked out, interfaces between the craft and its delivery systems could be designed, and form, fit, and function tests could be performed long before actual manufacturing begins.

Using a VCR to perform command and control of robotic devices and systems could prove helpful in reducing risk to humans during hazardous operations. For example, fueling a spacecraft can be extremely hazardous, particularly when hypergolic fuels are involved. Hypergol loading is a Self-Contained Atmosphere Protection Ensemble (SCAPE) operation, requiring evacuation of all nonessential personnel from the fueling area. Personnel performing the operation are required to wear protective gear in the event of a fuel spill. Replacing humans in this scenario with robotic devices equipped with manipulator arms, cameras, sensors, and instrumentation, a remote operator in a VCR could perform the fueling operation by controlling and monitoring the robot as it does the physical work. The operator would be able to see what the robot sees, monitor the progress of the operation, and be able to respond in real time to off-nominal conditions like a leak or spill of toxic materials. The hazardous nature of the operation does not change – the area must still be cleared of humans. However, the risk to human life is greatly reduced by removing the human operators from the hazardous area; at the same time, the human remains “in the loop” via the VCR.

A similar scenario could be envisioned for the International Space Station (ISS) or a human-tended Mars mission. An astronaut could use a VCR to control and monitor a free-flyer device designed to perform maintenance and repair tasks outside the spacecraft, thus reducing the need for astronauts to perform risky extravehicular activities (EVA).

C. Risk Identification

There are no risks associated with this project.

D. Setbacks

Procurement: There were difficulties purchasing IT hardware due to government restrictions on devices manufactured in China. Reliance on the NASA Assessed and Cleared List website for current information on approved devices was problematic, as the site does not seem to be updated in a timely manner. Time spent researching the country of origin for needed hardware and filing RFI's was time spent away from the lab, which negatively impacted progress.

E. Successes

The Kinect joint occlusion/bone length/jitter problems were significantly resolved by the tireless efforts of team member Tully-Hanson. The bone length and joint occlusion issues are close to 100% complete; jitter is still a factor, but further software development and testing will mitigate the problem. Beyond those issues, Tully-Hanson identified the problem of how the system can know, in a multi-Kinect environment, which direction the user is facing. (The Kinect software assumes at startup that the user is facing the device.) The initial solution, putting a marker on the user's chest, while violating the concept of a markerless tracking system, resolved the problem flawlessly, and suggested a possible future means for identifying the direction in which the user is facing without markers.

Integrating the Unity development environment with C# command and control software, the Oculus Rift, and the Leap Motion promise a true VR environment that will allow the user to control and monitor external physical devices. While attaining this goal is still a work in progress, confidence is high that a demonstrable capability is a certainty in the very near future (months, if not weeks).

F. Recommendation

The progress made during the past months have continued to reinforce the conclusion that the technologies investigated have a significant place in future NASA projects and programs. We are increasingly convinced that the work accomplished aligns quite strongly with Technical Area 4 (TA4) of NASA's Robotics, Tele-Robotics, and Autonomous Systems Roadmap and TA11 Modeling, Simulation, Information Technology, and Processing Roadmap. As stated in Section E above, a number of the problems associated with the Kinect have been resolved or are near resolution. Making this technology more accurate and reliable has been a high priority during this effort, and we believe we are well on the way to making it a viable tool for NASA's use in future command and control environments. Therefore, we recommend continuing work on designing and building a prototype VCR to demonstrate the use of AR/VR/NUI technologies.

F-1 Feasibility

The solution to building a VCR is both technically and financially feasible. Our PoC work has shown that many of the hardware and software elements of the solution are in place; the effort to integrate, test, and demonstrate a prototype system will prove that the feasibility of the VCR concept is well within our grasp.

F-2 Value

As we see it, this line of research has value to the agency on two major accounts. First, as we have noted, the use of AR/VR/NUI technology is being found increasingly in Commercial-off-the-Shelf (CoTS) products. As the industry moves ahead with development and refinement of these technologies, is in NASA's best interest to remain current with, if not be a leader in, the future of HCI. Second, as these technologies begin to permeate industry and academia, users of devices making use of AR/VR/NUI will become comfortable with nontraditional forms of HCI; future scientists, technologists, and engineers will come to understand and rely on these means of communicating with computers and robots. By investing in innovative research into these technologies now, NASA will be poised to provide its future workforce with the tools they will need to continue and advance our mission.

There are no security concerns at this level of application.

As the cost of AR/VR/NUI devices continues to come down, return on investment is assured.

G. Appendix

G-1 Lessons Learned

In the 2014 White Paper for Investigating Alternative Human-to-Computer Input Technologies for NASA Applications (Delgado, Little, Thompson, McGuire, Chang), we postulated that the Kinect V2 would go a long way to “fix” the joint occlusion/bone length/jitter problems. It did not. While significant improvements were made in the device (higher infrared and RGB camera resolution, increased depth of field, increased skeleton points), these problems remain beyond the scope of a single Kinect, and may never be successfully resolved using a single device. A multiple Kinect environment is still required for the solution. In fact, the solution lies outside the device itself, in the hardware and software receiving the data being streamed from it. Believing the difficult problems will be resolved in “the next release” is, more often than not, wishful thinking.

Another conclusion drawn in the above-mentioned white paper was that the “… next hurdle is to completely remove the need to ‘touch’ the display in order to control the interface.” Until such time as the technology to create a direct human brain-to-computer interface is truly feasible, this conclusion is somewhat flawed. Touch is every bit as important in developing AR/VR/NUI technologies as sight and sound. What Virtual Control Panel ultimately found is that without the sense of touch, the user is more often left guessing as to whether or not a gesture had its intended effect in a VR environment. A solution to this issue is currently in the very early PoC stages of development.

G-2 Benchmarking

The processes and procedures for benchmarking in this project are similar to those found in industry. That is, the technologies are still immature; applying standard development methodologies is difficult to do in so rapidly evolving a field. In fact, most, if not all, of the commercial developers of devices such as the Kinect, Oculus Rift, and Leap Motion, have gone out of their way to offer their development tools, particularly software development kits, to software developers and tinkerers to see what is possible. As the technologies mature and their capabilities become more fully realized, best practices are sure to evolve into industrywide standards.