# Network Bandwidth's Effect on Virtual World Simulator Performance Optimization


| Sean C. Mondesire | Jonathan Stevens | Douglas B. Maxwell |
|---|---|---|
| U.S. Army Research Laboratory | University of Central Florida | U.S. Army Research Laboratory |
| Orlando, FL | Orlando, FL | Orlando, FL |
| sean@cs.ucf.edu | jonathan.stevens@knights.ucf.edu | douglas.maxwell3.civ@mail.mil |



## ABSTRACT

The *United States Department of Defense* (*DoD*) employs virtual and game-based training simulators to train its servicemembers on both individual and collective cognitive and psychomotor tasks. The current employment of these training simulators is typically conducted in a stand-alone manner, with distributed simulation remaining the exception due to interoperability challenges. The *OpenSimulator* (*OpenSim*) is a popular open-sourced virtual world simulator that currently provides a persistent three-dimensional social community for its users. Under the *Army Research Laboratory*'s (ARL) *Military OpenSimulator Enterprise Strategy* (*MOSES*) research program, OpenSim is being developed to serve as a prototypical distributed military virtual training environment for tactical operations.

Virtual worlds for military training is an emerging domain. As such, detailed analysis of critical architecture parameters is required in order to optimize the performance of both the simulator's servers as well as the multitude of client connections. Unfortunately, due to a lack of extensive virtual world performance analysis, OpenSim server administrators often make arbitrary resource allocations to support their environments and training scenarios. Negative consequences to this approach are that typically too few resources are allocated to an overwhelmed server, resulting in an unresponsive environment, while too many resources are allocated to an underutilized server, when those resources could be more effectively applied elsewhere. In this paper, we analyze network bandwidth's effect on virtual world simulator performance so as to support the future creation of a predictive model that will determine the optimal amount of resources required to support a target number of concurrent users in the virtual world. This analysis, and the future development of our predictive model, will provide the OpenSim developer community with the knowledge required to best allocate resources to support expected server load.


## ABOUT THE AUTHORS

**Dr. Sean C. Mondesire** is a Postdoctoral Research Fellow at the U.S. Army Research Lab (ARL) and holds a doctoral degree in computer science from the University of Central Florida. His research focus is on the expansion of virtual world technologies to be used in tactical, military training.

**Dr. Jonathan Stevens** is a Research Scientist with the Institute for Simulation and Training (IST) at the University of Central Florida (UCF), specializing in training simulation research. Dr. Stevens is a retired Lieutenant Colonel of the United States Army with over 22 years of military experience as both an Infantry and Acquisition Corps officer. He received his Ph.D. in Modeling and Simulation from UCF**.**

**Douglas B. Maxwell** is a Science and Technology Manager at the U.S. Army Research Laboratory-Human Research and Engineering Directorate, Simulation and Training Technology Center (ARL-HRED STTC) located in Orlando, Florida. Mr. Maxwell is the founder of the Military Users of Virtual Worlds working group and serves as a virtual world technology advisor to the Office Deputy Under Secretary of Defense Readiness. He is responsible for the Military Open Simulator Enterprise Strategy (MOSES) simulator's architecture development and deployment. Mr. Maxwell is currently pursuing a Ph.D. in Modeling and Simulation at UCF.

# Network Bandwidth's Effect on Virtual World Simulator Performance Optimization

**Sean C. Mondesire**
**U.S. Army Research Laboratory**
**Orlando, FL**
**sean@cs.ucf.edu**

**Jonathan Stevens**
**University of Central Florida**
**Orlando, FL**
**jonathan.stevens@knights.ucf.edu**

**Douglas B. Maxwell**
**U.S. Army Research Laboratory**
**Orlando, FL**
**douglas.maxwell3.civ@mail.mil**

## INTRODUCTION

The U.S. Army continues to employ and refine its use of *simulation-based training* (*SBT*) in novel ways (Mishkind, Boyd, Kramer, Ayers, & Miller, 2013), in large part due to its proven effectiveness in training (Sotomayor & Proctor, 2009); (Lisk, Kaplancali, & Riggio, 2011); (Blow, 2012) and the need to decrease the cost of training. While the current employment of SBT is typically conducted in a stand-alone manner, a relatively new approach the Army is adopting is the use of persistent virtual environments for training, where these environments are defined as "persistent immersive simulated environments in which a participant uses an avatar (a digital representation of oneself) to interact with digital agents, artifacts, and contexts" (Dawley & Dede, 2013). One of the unique qualities of this simulation based training approach is the true distributive nature of virtual world training for Soldiers.

The Army has been an early adopter of virtual worlds for training (Schulzke, 2013), even as the technology is still relatively immature and evolving (Lele, 2013). However, recent advancements in simulation technology have enabled the rendering of sufficiently realistic environments in virtual worlds that can support effective training (Wong, Nguyen, & Ogren, 2012). Due to its distributed nature, virtual world training represents a major potential cost reduction to military training (Dutta, 2013); (Stafford & Thornhill II, 2012) and thus has attracted intense interest from the Department of Defense.

In light of the potential cost savings, as well as preliminary research indicating that virtual worlds may be effective for military training (Lackey, Salcedo, Matthews, & Maxwell, 2014), the Army's Learning Model explicitly calls for the increased use of this technology for training. In response to this mandate, the *U.S. Army Research Laboratory* (*ARL*) has established the *Military Open Simulator Enterprise Strategy* (*MOSES*) project to examine the efficacy of virtual worlds for military training (Ortiz & Maxwell, 2015). The project has heavily leveraged the open-source OpenSimulator platform to create a robust prototype framework in which to conduct training effectiveness research. The MOSES project's research goal is to determine whether or not persistent virtual environments are suitable for military training, particularly at the collective echelon and for distributed training activities.

Virtual worlds for military training is an emerging domain, with limited research devoted to the employment of this technology for collective training. Up to the present, virtual worlds have been primarily employed for training either at the individual level, focused on tasks such as public speaking anxiety (Kothgassner, et al., 2012), or for collective learning activities, such as a classroom, conference or virtual event. As an example, the Linden Lab's Second Life virtual environment was used to host approximately 800 conference attendees attending a virtual, distributed conference (Shami, Erickson, & Kellogg, 2011). While the pedagogical approach to using virtual worlds for training continues to evolve, current research is addressing the notable absence of any embedded instructional strategy in virtual worlds (Vogel-Walcutt, Fiorella, & Malone, 2013) and the creation of a robust training effectiveness framework (Landers & Callan, 2012). The MOSES research team is preparing to examine the suitability of virtual worlds for collective military training by conducting the required software development of essential capabilities, such as realistic models (weapons, ammunition and their effects), terrain and artificial intelligence as well as exploring the optimization of the architecture required to support distributed, non-deterministic training.

As we attempt to employ the MOSES virtual world at higher echelons of training, architecture limitations must be addressed. The foremost challenge remains the number of human avatars that may operate simultaneously in the virtual environment. In our prior research (Mondesire, Stevens, & Maxwell, 2015), we discovered that server configuration, specifically the number of cores dedicated to the MOSES server, had the greatest effect on simulation performance than any other variable examined. In this paper, we seek to expand our research to both optimize the performance of the simulator's server(s) as well as the multitude of client connection machines. To do so, we examine the effect that network bandwidth had on virtual world simulation performance. Network bandwidth will serve as an independent variable for the future development of a predictive model that will allow us to allocate the optimal amount of resources to maximize the number of human participants training in the MOSES virtual world environment. We analyze network bandwidth's effect on OpenSimulator resource demands in order to support a targeted number of concurrent users in the virtual world. With this analysis, and a future predictive model, the MOSES team will be able to better allocate resources to support expected server load in support of future virtual world training events.

**BACKGROUND**

This work analyzes the effect of network bandwidth on the three-dimensional, virtual world simulator *OpenSimulator* (*OpenSim*). OpenSim is an open-sourced, persistent simulator developed with the goal of providing its users with an immersive virtual world focused on social activities. Inside the simulator, users can socially interact within content-rich environments, dynamically change their appearance, and modify the environment in real-time (Overte Foundation, 2014). Figure 1 is a screenshot of users interacting inside an OpenSim virtual space.

Figure 1. A Screenshot of users interacting in OpenSim.

The simulator is heavily influenced by the popular and established social virtual environment, Second Life® ( Linden Research, Inc., 2015). Both products share software libraries and external applications to give users compatible and similar experiences in both virtual worlds. To date, hundreds of thousands of users have supported OpenSim by hosting their own simulators, contributing to the open-sourced development, content design, and interaction with the environment.

The simulator's persistence and on-demand environment changes are two features that make this product attractive to the military training community. By being a persistent simulator, OpenSim is designed to allow changes to its environment to persist indefinitely and support long durations of simulator runtime without restarting. In addition, the simulator supports dynamic, real-time changes to the environment. This feature allows users to load content into the virtual world and change the environment without pausing the simulator.

ARL has created the MOSES project to leverage the key features of OpenSim for military-based virtual world training. For military training research, MOSES runs multiple instances of OpenSim and manages these discrete virtual spaces as a connected world, where users can seamlessly navigate from one different environment and scenario to another. Already, the ARL MOSES project has contributed to the creation of a military virtual world through the development of an open-sourced virtual space management system, detailed hardware analyses of the simulator, simulation-based training effectiveness studies, and architectural changes to improve OpenSim's scalability.

The MOSES project strives towards adapting OpenSim into an effective large-scale virtual trainer, where hundreds or thousands of military personnel can interact within a single, simulated space to conduct collective training. Unfortunately, at present, OpenSim does not scale well beyond one hundred concurrent users. Previous MOSES research has explored increased hardware vertical scaling in OpenSim (Mondesire, Stevens, & Maxwell, 2015; Mondesire, Stevens, & Maxwell, 2015). From these research efforts, the virtual world community has learned that dedicating more server CPU cores to an OpenSim server has a greater positive effect on scalability than vertically scaling up by simply upgrading the server's CPU capability to more modern technology.

The work presented in this paper continues this vertical scaling investigation and analyzes the impact of network bandwidth on OpenSim's scalability. With an understanding of the role that network bandwidth has on OpenSim performance, the minimum network requirements for hosting a simulation-based training event, how many users the different network types can support, and the expected latency for each data transfer rate can be determined. More specifically, this network analysis provides insights on the type of military networks MOSES can be deployed on and sets expectations on how many personnel can participate in the virtual world under specific bandwidth limitations.

## METHODOLOGY

To determine the impact of network bandwidth limitations on a virtual world simulator, OpenSim's performance was evaluated on a range of bandwidth restrictions. Specifically, OpenSim's scalability was analyzed when the simulator was limited in *megabits per second* (*Mbps*) data transfer rates of 3G (1.4*Mbps*), 4G (6 *Mbps*), residential broadband (25 *Mbps*), and commercial T1 (55 *Mbps*) network speeds, averaged for both upload and download speeds. This data rate range was selected because it covers the typical spectrum from cellular access point to commercial network speeds common to OpenSim hosting. Additionally, MOSES may be deployed on a gigabit military network that limits OpenSim data rates due to other traffic concurrently residing on the network; this throttling can range from slow and shared cellular data rates to dedicated, commercial speeds.

### Simulator Server Configuration

An OpenSim simulation follows a client-server architecture, where the simulator is hosted by an instance of an OpenSim application (the server) and each user (the client) is realized in the simulation by running a separate viewer application. The server application processes everything to run the simulation, including, but not limited to, physics calculations, location updates, user logins and credential verification, user-to-user communication, and network message transfers. The *viewer* is the interface between the user and the simulator. The viewer renders all graphics the user sees and provides the controls that directs the user's avatar movements and actions. With this architecture, the server and each of its connected viewers constantly exchange network data to provide simulation updates and synchronize the simulation's state. The impact of various network bandwidth limitations are investigated in this paper because data synchronization is known to anecdotally affect OpenSim scalability, performance, and the user's experience. In this presented work, we attempt to quantify these impacts by following the scientific method.

During this experiment, the OpenSim server was physically hosted on a 2015 Intel Prototype server with 120 Intel Xeon E7-4890v2 CPU cores and 1.5 TB of RAM, connected to a 1 gigabit shared commercial network through a 10 gigabit switch. The OpenSim server instance was hosted on a Proxmox virtual machine (*VM*) running Ubuntu Server 14.04 LTS 64-bit, where the VM was allocated 8 CPU cores and 32 GB of RAM. This VM configuration is the typical hardware allocation used to host a MOSES simulation environment (*a region*). During each simulator experiment run, the VM was assigned 3G, 4G, residential broadband, or commercial T1 bandwidth limitations through its virtual ethernet (nic) card.

Each instance of OpenSim hosted a 256 square meter region with over 7000 primitive objects (*prims*) pre-placed throughout the space. The prims made up the content of the region, adding landscapes, buildings, and objects that avatars could collide and interact with. This test region represented the typical MOSES region that is used for virtual conferences, meetings, and training events.

### Simulation Load and Client Configuration

To place realistic and repeatable load on the simulator, simulated human user avatars (*bots*) were used during this experiment. The bots were realized by modifying the open-source OpenSim and Second Life client-application

*Phoenix Firestorm Viewer*. The modification allowed the popular viewer to be scripted to support automated avatar logins, movement, and posturing. By modifying the Phoenix Firestorm Viewer, interactions between the simulator and the clients generated similar network traffic and avatar movement behaviors as human-users.

With the viewer modification, each avatar in the simulation space required one instance of the viewer to be launched and loaded. Once logged in and loaded into the simulator, the modified viewer continuously selected and performed random actions for the remainder of the simulation run. The bot behavior pool included: move forward, backward, turn left, turn right, crouch, stand still, and jump. Each movement sent a command from the viewer to the server, stimulating network traffic and placing processing load onto the simulator.

For this experiment, two AMD Bulldozer servers were used to execute the bot applications and thus simulated the expected load of a distributed training event. Each Bulldozer ran a native install of Ubuntu 14.04 Desktop with Gnome and contained 48 AMD Opteron 6100 CPU cores and 500GB of RAM, hosted on a separate network from the OpenSim server. By hosting the viewers and server on separate networks and developing bots that continuously communicated with the server, simulator loads were similar to those of human-users; consistent and repeatable through bot scripting.

**Experiment Procedure**

The goal of this experiment was to determine the effect that network bandwidth had on the virtual world's performance. To gather statistics for the bandwidth and simulation analysis, automated testing that controlled both the simulator and the bots was performed. The following sequence of instructions was executed 10 independent times for each of the four network bandwidth levels, for a total of 40 simulation runs:

1. A VM was configured with a network bandwidth limitation (3G, 4G, residential, or T1).
2. The VM was started.
3. An instance of OpenSim was launched and initialized.
4. Simulator performance was continuously collected at one second intervals, using an automated script.
5. Every 30 seconds, a new instance of the viewer was launched from one of the two client (Bulldozer) servers to automatically log in a bot into the simulation's test region, with the maximum number of avatars to be launched set to 60.
6. Once all of the bots attempted to log into the server, the simulator waited for 2 additional minutes to allow pending logins to complete the process.
7. The simulator halted statistics gathering and shut down the OpenSim instance.
8. All of the bot viewer instances terminated on the two Bulldozers.
9. The simulator's VM was shut down.

**Data Collection**

For each experiment trial, performance data collection was enabled when the OpenSim server was fully operational and ready to accept logins from avatar accounts. This data collection mechanism allowed the state of the simulator to be captured in a quantifiable manner at a one-second interval. In previous vertical scaling analyses, metrics such as simulator frames per second, physics frames per second, total frame time, and process and host CPU and memory utilization have been used to compare performance speed-ups with hardware configuration changes. In this work, the measurement focus switched to avatar count and the simulation uptime in relationship to fully realized avatars in the virtual space. These two measures were employed due to the data latency bandwidth limitations placed on the login process of avatars and the rate at which simulation update messages were communicated between the clients and server. Furthermore, avatar count is used because each avatar introduces additional network load the simulator processes at any given time, meaning avatar count is used as an indicator of load. Also, simulation uptime is used to measure the duration required to processes each incoming avatar.

**RESULTS**

The purpose of this study was to examine the effect that network bandwidth had on virtual world simulation performance. This experiment employed one independent variable, the network bandwidth allocated for virtual world server and simulation operation. The independent variable had four discrete levels: 3G, 4G, residential broadband and T1 commercial broadband bandwidth allocations. These four levels were chosen in order to accurately replicate the typical bandwidth options a simulation operator would encounter when developing a virtual world training exercise and were discussed in detail in a previous section. The dependent variables examined were the maximum number of avatars introduced into the virtual world, the time required to load the maximum number of avatars and the time required to load the first five avatars into the simulation. The experimental procedure was executed forty times - with ten samples collected for each bandwidth allocation setting. For each bandwidth allocation, dependent variable data was captured in one-second intervals. The experimental procedure was executed identically for each bandwidth allocation.

**Maximum Avatar Count**

The first dependent variable analyzed was the maximum number of avatars introduced into the virtual world. This metric was chosen for examination due to the criticality placed on supporting all unit members composing a collective echelon who must utilize the virtual world for collective training. Our current primary use case is infantry maneuver training at the company echelon and below, with the training objective focused on movement, maneuver and communication. Thus, if an echelon is composed of 18 unit members then all 18 trainees must be able to train in the virtual world; partial collective training is not acceptable to the military Therefore, the higher the avatar count, both the higher the echelon that may train in the virtual world and the higher the probability that the entire echelon may train collectively. Table 1 and Figure 2 depict the virtual world's maximum number of avatars supported for all four bandwidth allocation levels. This maximum measure was calculated by averaging the maximum number of avatars observed across the 10 runs of each tested network bandwidth.

Table 1. Maximum Avatar Count Descriptive Statistics

| | 3G | 4G | Broadband | Commercial |
|---|---|---|---|---|
| Average | 10.1 | 22.8 | 24.1 | 24.9 |
| Standard Deviation | 1.7 | 1.1 | 2.3 | 1.4 |

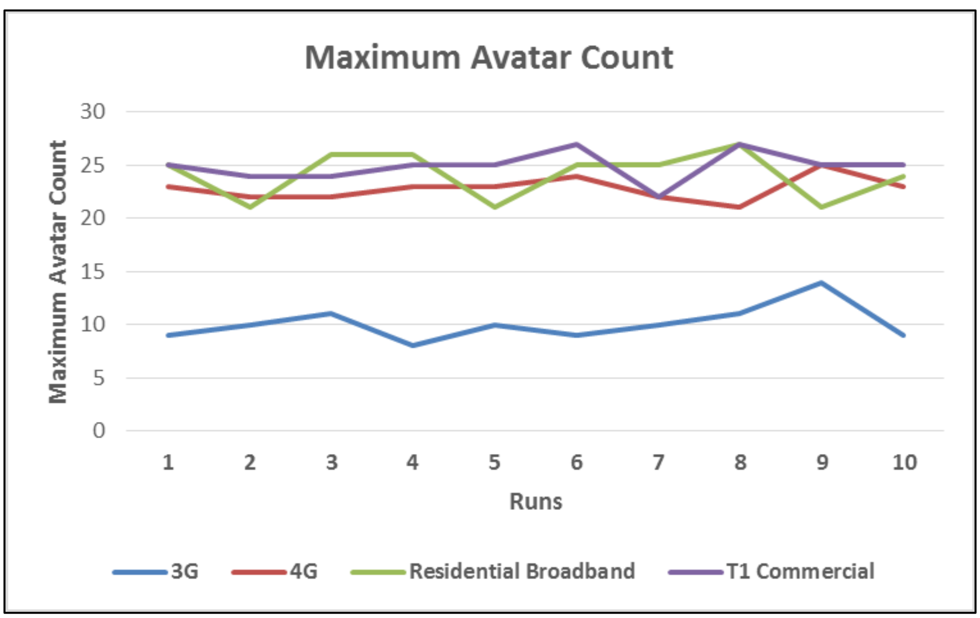


Figure 2. Maximum Avatar Count

ANOVA found a significant main effect of bandwidth allocation on the maximum number of avatars supported F(3, 39) = 170.99, $p < 0.0001$. ANOVA was conducted at $\alpha = 0.05$. Post-hoc pairwise comparisons using the Tukey HSD test indicated that the mean maximum avatar count for the 3G bandwidth allocation was significantly lower than the other three levels, $p < 0.0001$. However, there was no significant difference discovered in the maximum number of avatars supported between the 4G, Residential Broadband or T1 Commercial bandwidth allocations.

### Time to Load the Maximum Number of Avatars

The second dependent variable analyzed was the time required to load the maximum number of avatars into the virtual world. This is another metric of significant interest to the military as the number of exercise participants may exceed 100 at the company echelon. It is not acceptable to delay training so as to systematically inject avatars into the virtual environment in an artificially controlled and manually intensive manner. In short, the virtual world must be able to support both rapid and asynchronous entrance of exercise participants. Thus a lower time required to load the maximum number of avatars was better. Table 2 and Figure 3 depict the virtual world's time required to load the maximum number of avatars for all four bandwidth allocation levels.

Table 2. Time (seconds) to Load the Maximum Number of Avatars Descriptive Statistics

| | 3G | 4G | Broadband | Commercial |
|---|---|---|---|---|
| Average | 462.8 | 519.7 | 496 | 496.1 |
| Standard Deviation | 114.3 | 37.4 | 6.9 | 7.0 |

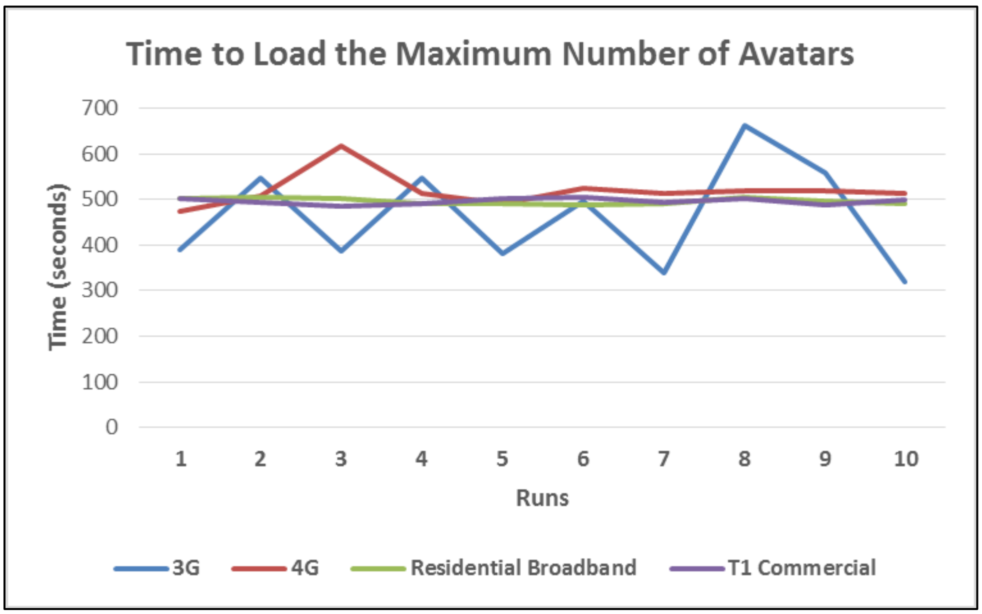


Figure 3. Time to Load the Maximum Number of Avatars

ANOVA did not indicate a significant main effect of bandwidth allocation on the time required to load the maximum number of avatars F(3, 39) = 1.50, $p = 0.23$. ANOVA was conducted at $\alpha = 0.05$. Therefore we conclude there was no significant difference discovered in the time to load the maximum number of avatars supported between the 3G, 4G, Residential Broadband and T1 Commercial bandwidth allocations. This result was not surprising since the maximum number of avatars differed between bandwidth allocation levels, as previously discussed.

### Time to Load the First Five Avatars

The final dependent variable analyzed was the time required to load the first five avatars into the virtual world. This dependent variable was selected for analysis because it is an efficient, yet accurate, method to compare bandwidth's effect on scenario initialization. For a typical virtual world training event, we estimated the average number of exercise support personnel required to be approximately five. Pre-execution, these five exercise design personnel are responsible for the creation, modification and maintenance of the virtual environment in preparation for and support of a high quality virtual experience. During execution, their ability to rapidly enter (and exit) the virtual world is paramount in order for a quality training event to occur, as these personnel typically serve as the simulation's opposing forces (OPFOR) and civilian role-players, depending upon the training scenario selected. We also believe that this load emulates the scenario initialization burden based upon our experience. Thus a lower time required to load the first five avatars was better. Table 3 and Figure 4 depict the virtual world's time required to load the first five avatars for all four bandwidth allocation levels.

Table 3. Time (seconds) to Load the First Five Avatars Descriptive Statistics

| | 3G | 4G | Broadband | Commercial |
|---|---|---|---|---|
| Average | 203.8 | 163.2 | 146.3 | 143.9 |
| Standard Deviation | 23.4 | 21.4 | 14.8 | 19.7 |

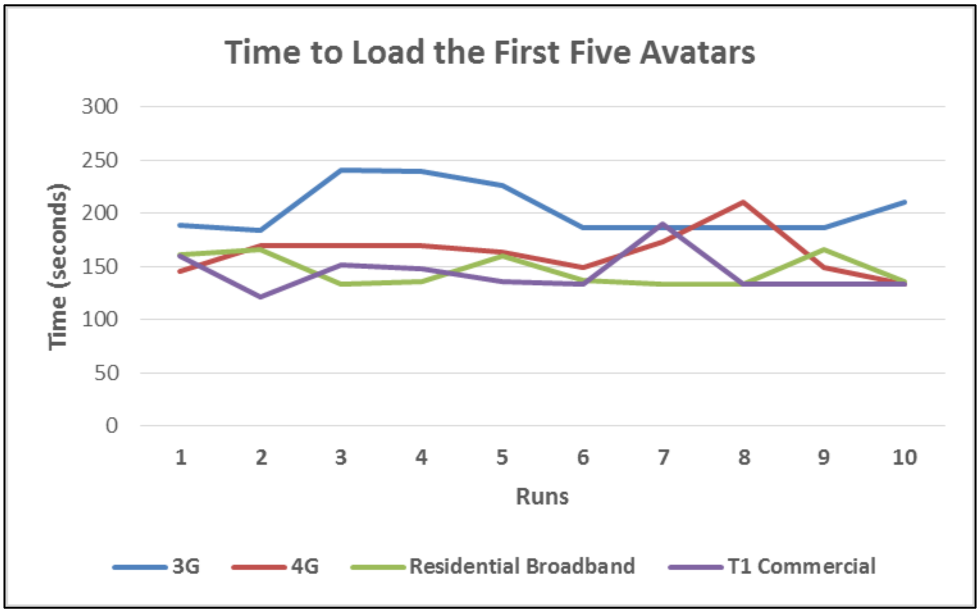


Figure 4. Time to Load the First Five Avatars

ANOVA found a significant main effect of bandwidth allocation on the time required to load the first five avatars supported $F(3, 39) = 19.10$, $p < 0.0001$. ANOVA was conducted at $\alpha = 0.05$. Post-hoc pairwise comparisons using the Tukey HSD test indicated that the mean time required to load the first five avatars for the 3G bandwidth allocation was significantly higher than the other three levels, $p < 0.001$. However, there was no significant difference discovered in the time required to load the first five avatars between the 4G, Residential Broadband or T1 Commercial bandwidth allocations.

## CONCLUSION

The purpose of this research was to examine network bandwidth's effect on virtual world simulator performance so as to maximize the number of trainees that may participate in a virtual world training event. Virtual worlds for training represent a novel employment of simulation-based training. The military has dual requirements to both increase the quality of training while at the same time reduce its cost. In order to achieve the latter objective, we must first determine a method whereby virtual worlds may support a higher number of concurrent users than what is currently possible. Significant simulation architecture challenges must be addressed to meet this goal. This paper addresses one of those challenges, network bandwidth, by examining its effect on the simulation's performance. The results of this study will directly contribute to the future development of a predictive model that encapsulates all of the major variables affecting the number of users that may participate in a distributed virtual world training event.

The results of this research are noteworthy, as we explored the effect that network bandwidth had on virtual world simulation performance. This experiment employed network bandwidth as the independent variable. Bandwidth's effect was analyzed through four discrete levels, corresponding to the most frequently available levels encountered by simulation practitioners: 3G, 4G, residential broadband and T1 commercial broadband. The dependent variable was virtual world simulation performance, further decomposed into the maximum avatar count, time to load the maximum number of avatars and time to load the first five avatars.

We discovered that network bandwidth had a significant effect on the maximum number of avatars that could operate in the virtual world simultaneously. The 3G bandwidth allocation supported a significantly less maximum number of concurrent users than the other three bandwidth allocations. This is an important finding as effective collective training must include all of the members of the collective echelon. In essence, we demonstrated that 3G bandwidth will support a lower echelon of training than the other three bandwidth allocations analyzed. Furthermore, we discovered that network bandwidth had a significant effect on the time required to load the first five avatars supported. Similar to above, the 3G bandwidth allocation performed significantly slower than the other three bandwidth allocations. This is noteworthy as we estimated the average number of exercise support personnel required to be approximately five. Thus this finding highlights that the 3G bandwidth allocation does not support the virtual world exercise overhead burden as efficiently as the other three levels. Finally, we found no difference in the time required to load the maximum number of avatars between different bandwidth allocations. This finding was not surprising as we determined in this paper that the maximum avatar count was found to be a function of the bandwidth allocated.

By determining the impact network bandwidth has on simulator scalability, this paper extended previous research endeavors that explored and analyzed vertical scaling in virtual worlds and produced a number of key observations. First, the scalability cost of deploying a virtual world simulator on a wide range of bandwidths is now known. The results confirmed that the number of active and concurrent trainees interacting in the virtual world does have a direct relationship to the simulator's network bandwidth speed. Second, although simulator scalability significantly increased from a 3G to 4G bandwidth increase, the system experienced diminished returns quickly as residential and commercial broadband speeds did not produce significantly different scalability improvements than 4G. From this conclusion, virtual world administrators now have an example of the feasibility and impact of deploying their simulators on slower, less expensive networks. Similarly, administrators also now have a demonstration of diminished scalability returns when a simulator is given faster, more expensive bandwidths. Finally, this network analysis confirms that there is a relationship between network bandwidth and the scalability of the tested virtual world simulator. Because this relationship exists that links an increased bandwidth to increased trainee count, future predictive modeling will be based on network bandwidth, in addition to server hardware allocation and concurrent users, for higher accuracy. Future work and research focuses on generating this predictive model to provide the community with a method of estimating the maximum amount of concurrent users the virtual world can support while maintaining a positive user-experience.

## REFERENCES


Linden Research, Inc. (2015). *Second Life Official Site*. Retrieved 3 13, 2015, from http://secondlife.com/

Blow, C. A. (2012). Flight School in the Virtual Environment: Capabilities and Risks of Executing a Simulations-Based Flight Training Program . *ARMY COMMAND AND GENERAL STAFF COLLEGE*(SCHOOL OF ADVANCED MILITARY STUDIES).

Dawley, L., & Dede, C. (2013). Situated learning in virtual worlds and immersive simulations. In *Handbook of research on educational communications and technology* (pp. 723-734). New York: Springer.

Dutta, D. (2013). Simulation in Military Training: Recent Developments. *Defence Science Journal*, 49(3), 275-285.

Kothgassner, O. D., Felnhofer, A., Beutl, L., Hlavacs, H., Lehenbauer, M., & Stetina, B. (2012). A virtual training tool for giving talks. *Entertainment Computing-ICEC 2012*, 53-66.

Lackey, S., Salcedo, J., Matthews, G., & Maxwell, D. (2014). Virtual World Room Clearing: A Study in Training Effectiveness. *Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2014.* Orlando, FL.

Landers, R. N., & Callan, R. C. (2012). Training evaluation in virtual worlds: Development of a model. *Journal For Virtual Worlds Research*, 5(3).

Lele, A. (2013). Virtual reality and its military utility. *Journal of Ambient Intelligence and Humanized Computing*, 4(1), 17-26.

Lisk, T. C., Kaplancali, U. T., & Riggio, R. E. (2011). in multiplayer online gaming environments. *Simulation & Gaming*, 1046878110391975.

Mishkind, M. C., Boyd, A., Kramer, G. M., Ayers, T., & Miller, P. A. (2013). Evaluating the Benefits of a Live, Simulation-Based Telebehavioral Health Training for a Deploying Army Reserve Unit. *Military Medicine*, (12), 1322-1327.

Mondesire, S. C., Stevens, J., & Maxwell, D. B. (2015). An Analysis of Increased Vertical Scaling in Three-Dimensional Virtual World Simulation (Publication Pending). *SimuTools.* Athens, Greece.

Mondesire, S., Stevens, J., & Maxwell, D. (2015). Vertical Scalability Benchmarking in Three-Dimensional Virtual World Simulation (Pending Publication). *SummerSim Conference.* Chicago, IL.

Overte Foundation. (2014, 11 10). *Main Page*. Retrieved 3 13, 2015, from OpenSimulator: http://opensimulator.org/

Schulzke, M. (2013). Rethinking Military Gaming America's Army and Its Critics. *Games and Culture*, 8(2), 59-76.

Shami, N. S., Erickson, T., & Kellogg, W. A. (2011). Common Ground and small group interaction in large virtual world gatherings. *ECSCW 2011: Proceedings of the 12th European Conference on Computer Supported Cooperative Work.* Aarhus Denmark .

Sotomayor, T. M., & Proctor, M. D. (2009). Assessing Combat Medic Knowledge and Transfer Effects Resulting From Alternative Training Treatments. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*.

Stafford, R. J., & Thornhill II, M. W. (2012). The Army Learning Model: Changing the Way Sustainers Train. *Army Sustainment*, 44(2), 28.

Vogel-Walcutt, J. J., Fiorella, L., & Malone, N. (2013). Instructional strategies framework for military training systems. *Computers in Human Behavior*, 29(4), 1490-1498.

Wong, J. H., Nguyen, A. B., & Ogren, L. (2012). *Serious Game and Virtual World Training: Instrumentation and Assessment (No. NUWC-NPT-TD-12-118).* NEWPORT, RI: NAVAL UNDERSEA WARFARE CENTER DIV .