SIMULATION 73:5, 273-280
© 1999, Simulation Councils, Inc.
ISSN 0037-5497/99
Printed in the United States of America

Developing an HLA Virtual Command Post

John W. Shockley
SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025
E-mail: john_shockley@sri.com

Kirk Parsons
Lionhearth Technologies, Inc.
25401 Spanish Ranch Road
Los Gatos, CA 95030
E-mail: kirk_parsons@lionhearth.com

Mark Morgenthaler
Lionhearth Technologies, Inc.
25401 Spanish Ranch Road
Los Gatos, CA 95030
E-mail:
mark_morgenthaler@lionhearth.com

SRI International and Lionhearth Technologies have been developing an HLA-compliant version of the Virtual Command Post (VCP), a distributed simulation in which users interact and share battlefield information. The project involves developing HLA federates for individual functional VCP capabilities. The goal is to develop an HLA-compliant VCP as a stand-alone federation, as well as a set of functional elements that can interact as separate HLA federates in other federations. As a federate, the VCP can participate in federations requiring a command, control, communications, computers, and intelligence federate. The individual VCP federates can act as separate elements in other federations. Among issues examined is run-time infrastructure (RTI) performance sending high-volume audio and video data. We are now examining the H-Anim standard for representing human commanders. We have found that the HLA and the RTI are capable of meeting our functional requirements for a distributed simulation. We also believe we can develop an HLA-compliant simulation that will promote its use in HLA federations. This paper describes the HLA VCP, the experiments and results we have obtained to date, and efforts we have planned for the future.

Keywords: Virtual Command Post, high-volume data, HLA, commercial standards

1. Introduction

Since early 1998, SRI International and Lionhearth Technologies have been developing a High Level Architecture-compliant version of the Virtual Command Post (VCP). The VCP is itself a distributed simulation providing a virtual collaborative environment in which geographically distributed users can interact and share battlefield information in a realistic depiction of a command center. The HLA VCP project involves developing HLA federates for individual functional capabilities of the VCP. The goal of this project is not only to develop a VCP that is compliant with basic HLA requirements as a stand-alone federation, but also to develop a set of VCP functional elements that can interact as separate HLA federates in other federations. As an HLA federate, the VCP can participate in a variety of federations requiring a command, control, communications, computers, and intelligence (C4I) federate. The individual VCP federates—e.g., audio, video, and commander federates—can be used as separate functional elements in other HLA federations. As a result of this goal, SRI and Lionhearth have examined several technical issues. In particular, we have examined run-time infrastructure (RTI) performance sending high-volume audio and video data, and we are currently examining the H-Anim standard for representing human commanders in the virtual world. These results have been incorporated into papers published at the Simulation Interoperability Workshop—i.e., 98F-SIW-214 and 99S-SIW-102. In examining these technical issues, we have found that HLA and the RTI are capable of meeting our functional requirements for a distributed simulation. We have also gained confidence that we can develop an HLA-compliant simulation that will promote its use in HLA federations.

The HLA-compliant VCP distributed simulation is an extension of the original VCP project developed as the result of a Small Business Innovative Research (SBIR) contract. The original VCP concept enables commanders in the field to "meet" in a virtual collaborative environment that possesses the principal capabilities of a command post. Such a command post has the advantage of simultaneously being both everywhere and nowhere. Because the VCP needs to connect real commanders in the field, it must do so via technology that is available in a battlefield environment and via limited bandwidth (9600-baud) radios.

The original VCP was developed by Lionhearth; three nodes were delivered to the U.S. Army Communications and Electronics Command (CECOM) in 1997 [1]. Since then, Lionhearth and SRI have recognized the need to update the VCP and to make it compliant with the DoD HLA [2]. In addition, because Lionhearth's technology focus included commercial product development, the HLA VCP needed to comply with commercial requirements and standards. These two goals led us to develop an architecture that we believe will satisfy both.

To satisfy the HLA requirements, we have examined the performance of the HLA RTI[1] in its ability to handle high-volume audio and video data [3, 4]. To satisfy commercial standards requirements, we are also examining the feasibility of and plan for incorporating commercial standards into HLA-compliant simulations.

In the following sections, we (1) describe the basic architecture and capabilities of the HLA VCP, (2) describe some of our experiments to determine RTI performance in handling high-volume data, (3) describe our current efforts to incorporate H-Anim, an international standard for human figure representation, as an HLA reference, and (4) determine what we see as our next steps.

2. The HLA VCP—Architecture and Capabilities

2.1 HLA VCP Architecture

The VCP is an early experiment in the use of distributed simulation and virtual conferencing technologies to provide a "you are there" collaborative C4I environment. Developed by Lionhearth Technologies for the U.S. Army CECOM under a Phase II SBIR contract, commanders and war fighters in geographically dispersed locations can use the VCP to view battlefield information on portable or man-wearable devices, including laptops, personal digital assistants (PDAs), or head-mounted displays (HMDs). The VCP captures each commander's "presence" in the form of compressed voice, line-of-sight, and body position, and can transmit the commander's voice and animated iconic representation over a 9600-baud radio network, leaving bandwidth available for maintenance of a distributed battlefield management database. This database is updated locally and converted at each node into 3-D battlefield visualization with interactive planning, "what-if," and data drill-down capabilities. In parallel, satellite broadcasts or video-teleconferences can be captured and displayed in the VCP on video walls. Figure 1 depicts an initial system block diagram of the VCP and highlights some of the technical details of the system.

More recent efforts in refining the VCP architecture have resulted in a more streamlined diagram. This is

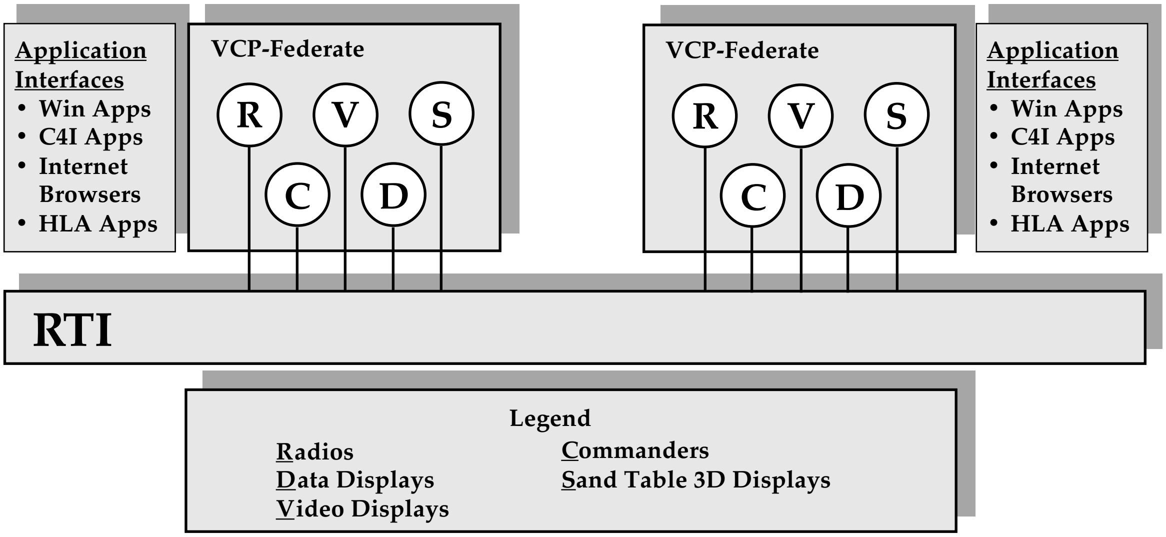

Figure 1. Original HLA VCP architecture

---

[1] RTI versions in these experiments were the DMSO Versions 1.0.2 and 1.3.3 for Windows NT, respectively.

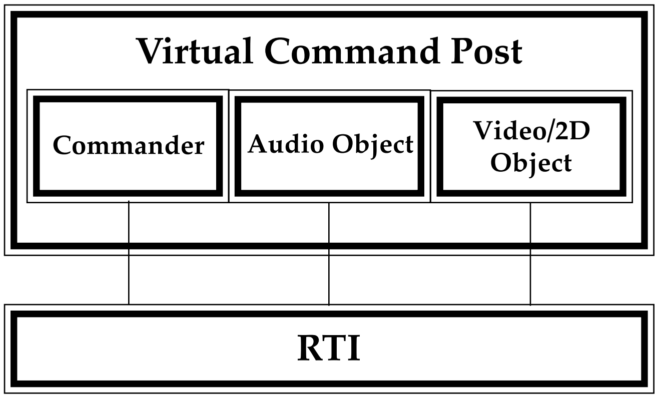

Figure 2. Current HLA VCP architecture

depicted in Figure 2. This current architecture defines the following separate elements as federates: (1) commander or avatar, (2) audio object, (3) video or 2-D object, and (4) the VCP room itself. The first three proposed federates represent technical modules of the VCP, while the fourth (the VCP room) represents the entire VCP as a separate federate. In addition, we are examining the feasibility of including a generic 3-D object or viewer as a separate federate. This division was based on how the original VCP was designed and what we saw as a reasonable division of functional modules that could be separate HLA federates. We envision that each of these could be separate functional entities that are part of an HLA federation other than the VCP.

2.2 VCP Federates

As currently defined, the HLA VCP consists of up to four HLA federates (commander, audio object, video/2-D object, and room, as defined above). The purpose of each of these federates is to provide some flexibility and modularity in the design of the system so that elements of the VCP might be used individually or collectively, as dictated by the application requirements.

Since there are four federates, SRI and Lionhearth needed to develop a simulation object model (SOM) for each. Tables 1 through 4 provide the top-level object class structure tables for commander, audio, video/2-D, and room federates, respectively. Note that the room federate is essentially a representation of the collective capabilities of the entire VCP and, therefore, is a union of the individual SOMs for the other elements.

The object class structure tables are necessarily simple. For the audio and video data, we are only considering the streaming data as an object class for each federate. For the commander federate, we have included a more complex class structure to map into the representation used for H-Anim, the Virtual Reality Model Language (VRML) standard to which we are adhering for avatar representation. We believe this simple SOM representation will most closely map into the capabilities and architecture of the HLA VCP.

2.3 HLA VCP Functional Capabilities

The current version of the VCP is implemented on a per-node basis. Each node consists of a Windows PC, which serves as the primary processor, an HMD, a user mouse interface, a user audio and microphone interface, and user motion capture instrumentation. Except for the motion capture instrumentation, the hardware included in a VCP node is composed of standard, commonly available commercial products. Except for the HMD, these products are also generally included in a typical PC suite.

The motion capture instrumentation is currently based on magnetic detection technology and requires the user to wear a harness tethered to the sensor system. Lionhearth and SRI are currently developing a new motion capture technology to replace this tethered system. The most promising technology is a motion capture system that employs stereo cameras and active target tracking (patent pending). Figure 3 illustrates the current hardware architecture of the VCP.

The functional components of the VCP include the following elements (for each command post represented):

- VCP Room with functional command areas.
- Avatar representation of each networked user (i.e., node) including motion and sound.

Table 1. Object class structure table—commander federate

Object Class Structure Table	
HumanoidRoot (PS)	Joint (PS)
	Effector (PS)
	Displacer (PS)

Table 2. Object class structure table—audio federate

Object Class Structure Table	
AudioSourceRoot (PS)	AudioStream (PS)

Table 3. Object class structure table—video federate

Object Class Structure Table	
VideoRoot (PS)	VideoStream (PS)

Table 4. Object class structure table—VCP federate

Object Class Structure Table		
VCPRoot (PS)	Commander (PS)	Joint (PS)
		Effector (PS)
		Displacer (PS)
	Audio (PS)	AudioStream (PS)
	Video (PS)	VideoStream (PS)

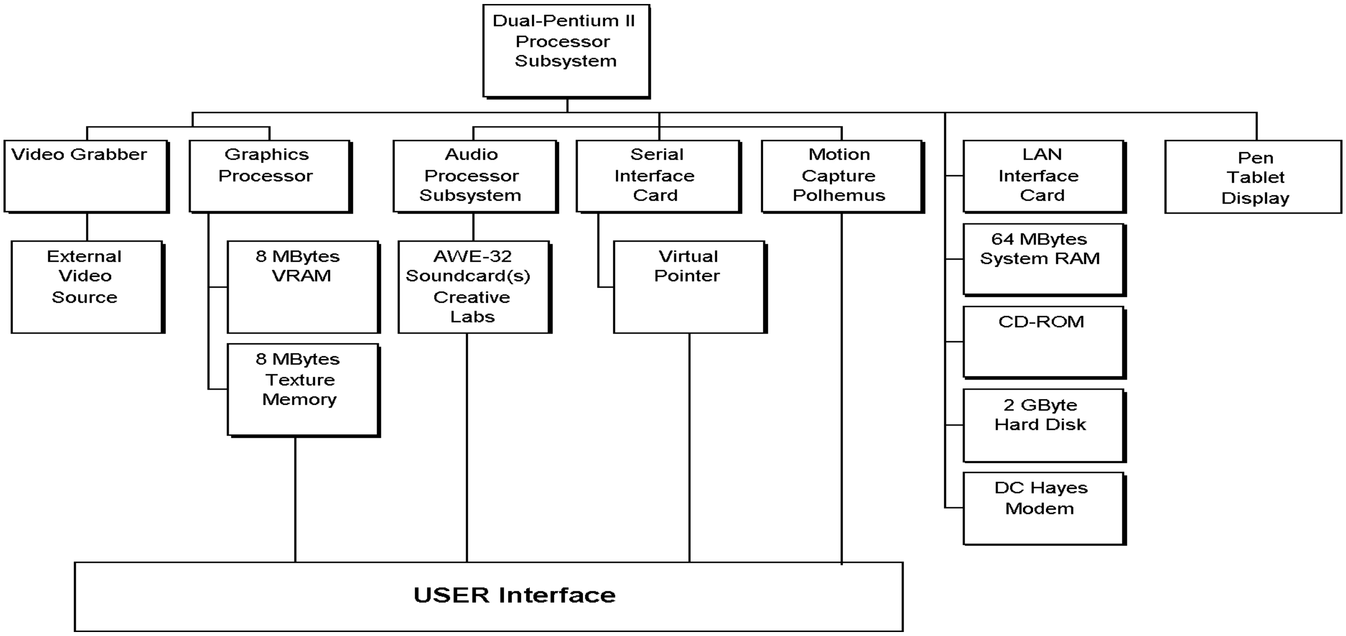

Figure 3. Architecture of a VCP node

- Two-dimensional displays on the walls of the Command Post—each display is capable of displaying video or any Microsoft Windows program.
- Three-dimensional sand table representation of the battle area of interest.

While the other functional capabilities are self-explanatory, the sand table provides a unique capability worthy of further elaboration. The sand table is a 3-D representation of a digital terrain map of the area of interest to the users. It is linked to a 2-D display on the command post wall that provides a plan view of the area. In the plan view, the user is able to create and manipulate iconic representations of other users, routes, phase lines, etc. These representations, in turn, are transformed onto the sand table for display. The unique capability of the VCP sand table (because the command post is virtual) is that users can "teleport" themselves onto the surface of the terrain and actually see the effect of the terrain on movement, line-of-sight, etc.

3. RTI Performance Experiments

Because the VCP concept is based on connecting nodes via a radio net and includes the provision for video teleconferencing, the ability of the HLA VCP to support these kinds of data is critical to its success. For this reason, Lionhearth and SRI have investigated the performance of the RTI in handling high-volume data. This is a performance issue because the data stream is not only large in quantity, but typically must be sent in near-real time to meet the functional requirements of VCP federates. The two experiments we have conducted have examined RTI performance for handling real-time audio and video data.

3.1 RTI Performance in Handling High-Volume Audio Data

In the first experiment [3], we examined the ability of the RTI to transmit real-time audio data. Specifically, we compared the performance of three networked configurations: (1) using the RTI to send real-time audio data directly between two simulated radio federates, (2) using the RTI to send "token" radio data between the federates, while using a dedicated backchannel to send the high-volume data, and (3) using the commercially available application, NetMeeting, to send the audio data. In Option 1, the RTI was both the transport and control layer, while in Option 2, the RTI was only the control layer. Both Options 1 and 2 were designed to be compliant with Rule 3 of the HLA requirements:

> "During a federation execution, all exchange of FOM data among federates shall occur via the RTI."

Table 5. Summary of RTI performance test results*

Test Scenario	Data Latency (ms)
1. RTI as Control and Transport Layer	100-200
2. RTI as Control Layer Only	100-200
3. MS Net Meeting with H.323	300-400

* From, M. Morgenthaler and J. Shockley, "RTI Performance in Sending High Volume Data." #98F-SIW-214, *Simulation Interoperability Workshop*, September 1998, Orlando, FL.

The third configuration was included to compare what might be expected from current commercial capabilities.

These initial experiments suggested that the RTI appeared to be capable of sending high-volume audio data. Table 5 (repeated from Reference 3) illustrates the results of these experiments.

These results were encouraging because the RTI provided results comparable to those of a backchannel and better than those provided by the commercial Microsoft NetMeeting configuration. However, our experiment had some inherent limitations. In particular, the experimental setup only included sending audio data between federates. Eventually we will attempt to transmit more demanding video data.

Also, only two radio federates were in the experimental setup. With the scalability of the RTI a critical concern regarding performance, we recognized the need to add several radio federates to the experiment and to determine the scalability of such a configuration. Nevertheless, because of the encouraging results, we pursued the architecture represented by Option 1 (sending data directly over the RTI) in developing the overall architecture for the HLA VCP.

3.2 RTI Performance in Handling High-Volume Video Data

The second experiment (described in detail in Reference 2) was more of a capability demonstration. In this case, Lionhearth developed a prototype video-teleconferencing (VTC) capability that operated over the RTI. Sending real-time audio and video data, necessary for a useful VTC capability, are the most demanding interactions between the VCP federates. This prototype, therefore, served to demonstrate perhaps the most critical aspect of the HLA VCP.

A system diagram of VTC capability is shown in Figure 4. In this demonstration, three Pentium II 400 PCs were connected over a 10-Mbit Ethernet. Each PC is indicated as a VTC node, where each node has a video camera with a video capture board and a full-duplex sound card. One of the nodes is running Windows NT, while the other two are running Windows 95.[2] In addition, one of the PCs hosted the RTI (Version 1.3.3) executive.

In the VTC federation we developed, each VTC node also constituted an HLA federate capable of transmitting and receiving both audio and video data. Only simple feasibility demonstrations were completed; nevertheless, they provided some enlightenment on the performance of the RTI and its ability to handle high-volume data. Figure 5 is a screen capture of the RTI-based VTC in operation. It shows four video images as well as the *rtiexec.exe*.

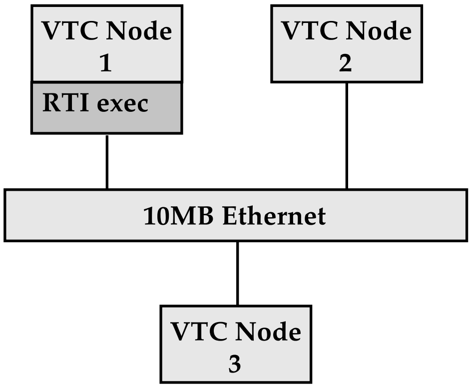

Figure 4. RTI-based VTC capability

We conducted simple feasibility demonstrations using the prototype capability to conduct some normal VTC interactions. Since the tests were informal, there was no test plan and we took no effort to quantify "normal" VTC interactions. Our overall goal was to see if these VTC interactions could be supported via the RTI within an end-to-end latency of 100 ms.

The RTI VTC capability developed demonstrated video teleconferencing between as many as five nodes on a 10-Mbit Ethernet.[3] Some of the key characteristics of these interactions and lessons we learned were:

- Video consisted of 160x120-pixel, gray-only, RLE-packed frames keyed at once per second.
- The audio was unpacked pulse code modulation (PCM) at an 8-kHz rate.
- Overall data rate averaged 80 kB/s for two nodes and 140-180 kB/s for three nodes.
- Latencies were 200 ms for either the two-node or three-node configuration. While this did not meet our 100-ms requirement, we felt that we could eventually meet it.
- We found difficulties sending the multimedia data over the RTI using best-effort UDP; packets built up. We eventually decided to send all data as "reliable."
- We did not concentrate on data packing schemes that we anticipated would improve performance substantially.

In both of the experiments, we were encouraged that our approach to use the RTI to send high-volume data directly was feasible. While we failed to observe our 100-ms latency goal in either case, the performance we observed suggested we would be able to achieve this goal with a more detailed examination and additional effort in the design of the systems.

---

[2] The RTI VTC (VideoPhone) simulation developed for this demonstration runs on both Windows NT and Windows 95.

[3] Although five nodes were demonstrated, only four were capable of showing a video image. This was due to technical issues associated with the Windows NT operating system that operated on that node.

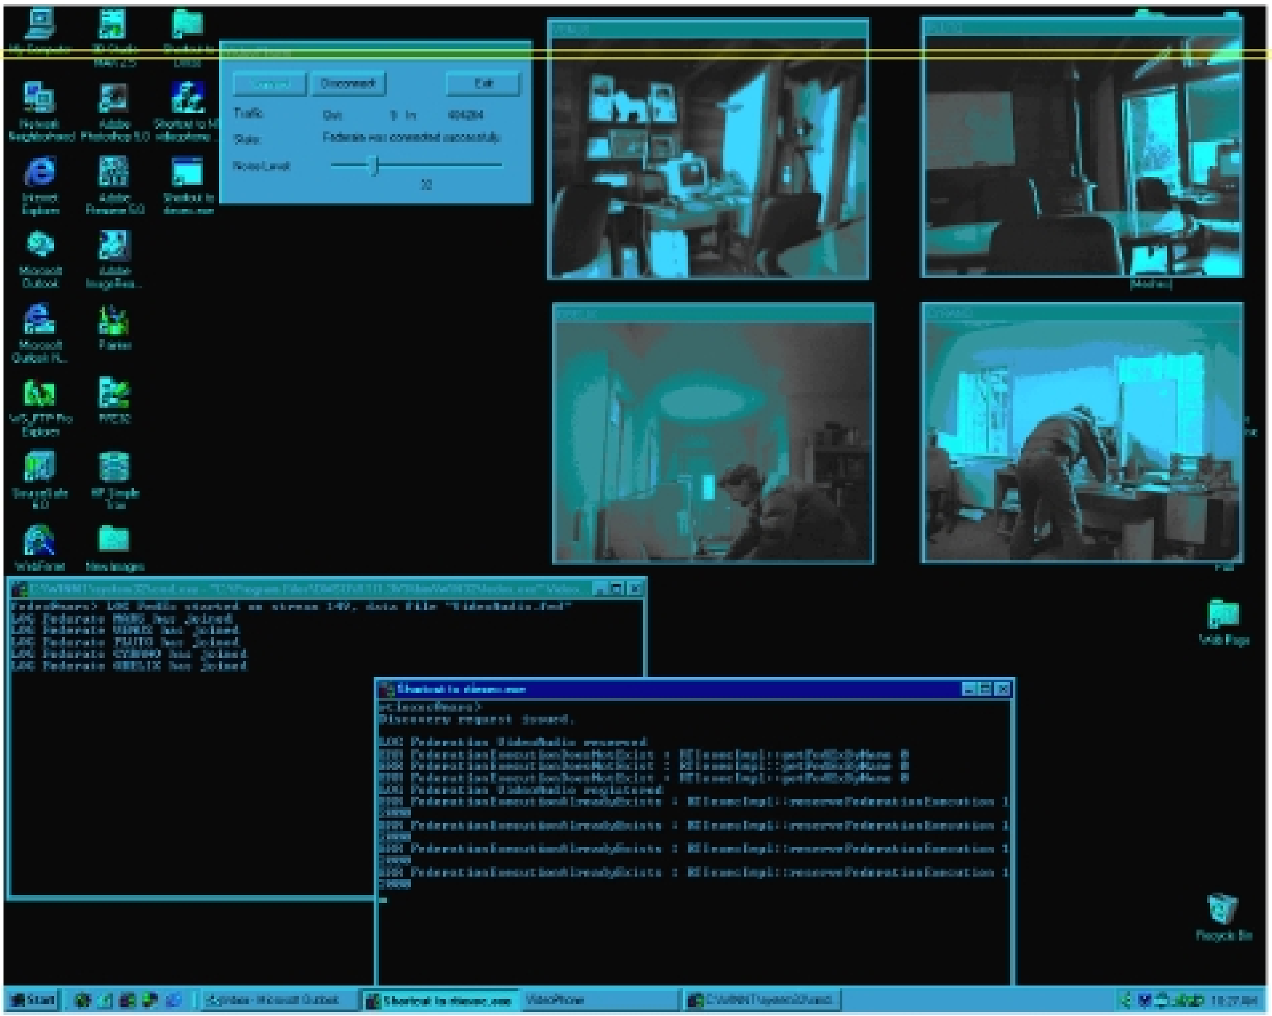

Figure 5. Screen capture of RTI-based VTC in operation

4. Using Commercial and DoD Standards and the Applicability of H-Anim

From the initial development of the VCP, the intent has been to create technologies that can be used in both military and commercial applications. The military applications of the VCP are not only a virtual collaborative environment for field commanders, but also in simulation exercises involving command, control, communications, computers, and intelligence (C4I) assets as a command post simulation. With HLA, we envision additional simulation applications of the individual components as federates—e.g., avatar representations of individuals in a simulation. For commercial applications, we envision a variety of uses of the VCP technology in the creation of virtual collaborative environments for VTC applications, distance learning, and electronic commerce.

Because of the dual focus of VCP technologies, Lionhearth and SRI have needed to focus on satisfying both military and commercial standards simultaneously. The military applications drove us to use HLA (though we also see possible technical advantages to its use, regardless of DoD requirements). The commercial applications are driving us to the adoption of a number of standards. These include:

- Use of the Microsoft Windows operating system.
- Use of the Moving Picture Coding Experts Group (MPEG) standards for representation of high-volume data.
- Use of the Virtual Reality Modeling Language (VRML) standards for avatar representation.

Our adoption of both military and commercial standards in the development of the HLA VCP has presented us the challenge of making both sets of standards work together. The best example of our attempt to accomplish this is our proposal to use the VRML standard H-Anim for representing human figures in HLA-compliant simulations [5]. This effort arose from the definition of the VCP commander as a separate HLA federate.

4.1 H-Anim

The "H-Anim" standard was created by the Humanoid Animation Working Group of the Web3D Consortium to create representations of human figures (i.e., virtual humans, using VRML 97). H-Anim enables avatar designers to create such figures using Web-based avatar tools.

The primary goal of the H-Anim standard is to specify a way of defining interchangeable *humanoids* and *animations* in standard VRML 97 without extensions. This will allow people to author humanoids and animations independently. To achieve this goal, it is necessary to specify a standard set of joint names (and the geometry, or segments, corresponding to those joints) as well as a minimal set of modeling conventions.

The complete set of H-Anim joint names includes 94 joints, which is believed to be sufficient to model the most anatomically correct virtual human model possible. Users of the H-Anim standard are free to use any subset of the 94 standard joints. It is also possible to add joints to the H-Anim hierarchy, provided no non-H-Anim joint is added between two H-Anim joints. Geometry is represented by segments in H-Anim, and there is a segment corresponding to each of the 94 joints in the standard hierarchy.

The advantage of H-Anim is its broad acceptability, not only in the general Web-based community, but also in other commercial environments. It has also been developed as an interchange format for representing virtual humans developed from a variety of 3-D tools and motion capture devices. Furthermore, it is the basis for representation of avatars in the MPEG-4 standard.

H-Anim provides a useful basis for avatars in the VCP primarily through run-time animation reuse—possible through H-Anim's use of a standard set of joints, along with specific modeling conventions. In addition, H-Anim offers an improved "production pipeline"—the process whereby avatars are authored in off-the-shelf authoring tools (e.g., Kinetics' Character Studio™) and then brought into the run-time VCP.

4.2 H-Anim Reference SOM

To evaluate the applicability of H-Anim as an HLA reference, we developed an initial version of a SOM for the VCP avatar federate. The object class structure table was presented in Table 1. Table 6 presents an abbreviated version of the interaction table for such a SOM, and Table 7 presents an abbreviated version of the attribute table. Although these three abbreviated tables obviously do not constitute a complete SOM, we believe they are sufficient to convey how such a SOM is represented.

Our efforts to date indicate that the H-Anim standard maps well into the HLA object model template (OMT). For this reason, we are encouraged that we can develop a commander federate based on H-Anim. Furthermore, the applicability of H-Anim and its acceptance in the international VRML community suggest that it will be a suitable reference SOM for human representation in HLA.

Table 6. Interaction class structure table for H-Anim

Interaction
Translate
Rotate
DisplaceAction

Table 7. Attribute table for H-Anim (abbreviated)

Object	Attribute
Effector	R_hand_tip
	L_hand_tip
	Skull_tip
	R_middistal_tip
	L_middistal_tip
Joint	L_ankle
	L_elbow
	L_hip
	L_kneww
	L-midtarsal
	L-shoulder
	L-wrist
	R_ankle
	R_elbow
	R_hip
	R_knee
	R-midtarsal
	R-shoulder
	R-wrist
	Skullbase
	V15
Displacer	Segment

5. Next Steps

Our efforts to date have been encouraging in all respects. The RTI performance experiments suggest that we will be able to send high-volume data directly over the RTI. H-Anim appears to map well into the HLA OMT and appears to be a suitable reference for a human representation SOM. Overall, our efforts to make the VCP compatible with both commercial and military standards look fruitful. At the same time, however, we have a considerable amount of work remaining.

Our next steps involve completing the efforts addressed in this paper. These include developing a proposal for the H-Anim standard to be an HLA reference. They also include additional performance experiments on the RTI. In particular, we are in the process of developing a latency budget on end-to-end transmission of high-volume data. Through this examination, we hope to be able to identify the major contributors to latency and those on which we can focus additional efforts to reduce.

While we have developed draft versions of the SOMs/FOMs for the HLA VCP, we have yet to develop revised versions that we will use in compliance testing. Once these are complete, we will be seeking HLA certification.

6. References

[1] Lionhearth Technologies, Inc. *The HLA Virtual Command Post,* Response to OSD97-006 Solicitation: Commercialization of Components C4I Interface to Simulation Using HLA, April 23, 1998.

[2] Morganthaler, M. and Steiner, G. *Phase II SBIR Final Report for the Dispersed Virtual Command Post Program*, Lionhearth Technologies, March 23, 1998.

[3] Morgenthaler, M. and Shockley, J. "Performance in Sending High Volume Data." #98F-SIW-214, *Simulation Interoperability Workshop*, Orlando, FL, September 1998.

[4] Morgenthaler, M. and Shockley, J. "Performance in Sending High Volume Data: An Update—The Development of an RTI VTC Capability." #99S-SIW-102, *Simulation Interoperability Workshop*, Orlando, FL, March 1999.

[5] Humanoid Animation Working Group, *Specification for a Standard VRML Humanoid*, Ver. 1.0, December 1997. [Note: Ver. 1.1 was released August 1999]

John Shockley is a Senior Research Engineer at SRI International. Over the past 15 years, he has worked at SRI on a broad variety of projects—primarily in the areas of test and training range instrumentation systems for the U.S. Army, Navy, and Air Force. He began working on modeling and simulation aspects of these systems and has since participated in DIS/HLA standards development activities for the past nine years, concentrating on integrating live and virtual systems.

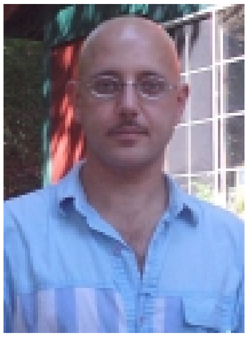

Kirk Parsons is Lionhearth's Avatar Technologist. At Lionhearth, he has been responsible primarily for development of the NetPresenter product, where avatars are used in support of Web-based call centers. Also, he has brought together an "avatar team" of artists and graphics programmers who are creating H-Anim avatars and H-Anim technology. Prior to joining Lionhearth, he co-authored the H-Anim Virtual Human Standard, and is currently a Co-Chair of the working group. He had also founded Attic Graphics, Inc., a VRML avatar development company later acquired by Blaxxun Interactive.

Mark Morgenthaler is the Founder and President of Lionhearth Technologies. At Lionhearth, he has assembled a team of technical, administrative, and marketing professionals to develop virtual environments and sensor technologies to enhance collaboration. Prior to forming Lionhearth, he served as President of Crystal River Engineering. During his tenure there, he revitalized this virtual reality company, resulting in the development of a 3-D sound API—which was subsequently incorporated into Windows 95. In 1993, he joined Trimble Navigation and founded the Avionics Division, responsible for developing differential GPS landing systems. Prior to his work at Trimble, he had 14 years of experience at Hewlett-Packard as an Engineering Manager, working on a variety of graphics, multimedia, and RISC architecture projects.