The Joint Training Experimentation Program Approach to Distributed After Action Review

Reginald Ford, Victoria Lamar
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
650-859-4375, 650-859-5670
reginald.ford@sri.com, victoria.lamar@sri.com

Richard Giuli, Scott Oberg
SRI International
4111 Broad Street; Suite A-7
San Luis Obispo, CA 93401
805-542-9330
richard.giuli@sri.com, scott.oberg@sri.com,

Keywords: Distributed After Action Review; LVC Training Systems

ABSTRACT: *The Joint Training Experimentation Program (JTEP) is a multiphase, multiyear effort to develop a distributed training capability for the California National Guard (CNG) that includes live, virtual, and constructive training simulation to support multi-echelon training. The second JTEP demonstration was a battalion-level exercise conducted in December 2003. This demonstration linked the Joint Combat and Tactical Simulation (JCATS), a constructive simulation; the Close Combat Tactical Trainer (CCTT), a virtual simulation; the Deployable Force-on-Force Instrumented Range System (DFIRST™), a live instrumented training system; and observers at the Office of the Adjutant General (OTAG).*

The first JTEP demonstration, in May 2003, introduced a JTEP After Action Review (AAR) capability. A key goal of the second demonstration was a distributed AAR (DAAR) that would allow commanders and observers at the three sites to view exercise playback on the JTEP high-resolution 2-D and 3-D map displays, and to interact via a standard video teleconference (VTC). This paper describes the DAAR conceptual and technical approaches, and the lessons learned in conducting the first DAAR.

The master site was located in the CNG Distance Learning Classroom (DLC) at Camp San Luis Obispo (SLO). Three additional participating sites were located at Camp SLO, Camp Roberts, and OTAG. Each site was a node for DLC-hosted VTC capabilities. One option was for participants at remote sites to view 2-D and 3-D map displays via the VTC; however, VTC resolution and update rate limitations would have resulted in an unacceptably degraded view. Instead, we opted to run synchronized local instances of the 2-D and 3-D displays at every site. Each of the four sites hosted a full complement of JTEP AAR software and hardware.

The primary technical challenge was to provide lossless remote viewing while staying within the limited 700 kbits/s bandwidth allocated to JTEP on the California Army National Guard (CA-ARNG) network. The allocated bandwidth was adequate for sending out Distributed Interactive Simulation (DIS) data and voice protocol data units (PDUs) to all sites during live operations; however, the bandwidth was not sufficient to carry this load plus a VTC and the greater number of PDUs generated by fast-rate playback. The approach taken was to run local instances of the playback process at all sites, each playing local copies of the data (recorded locally during the live exercise, or distributed in compressed format after the exercise), and send only control information over the network.

1. Introduction

1.1 JTEP overview

The Joint Training Experimentation Program (JTEP) is a National Guard Bureau program managed by the California National Guard (CNG). The Guard currently uses advanced live, virtual, and constructive (LVC) systems[1] to support training, but each system is stand-alone. JTEP was conceived to bring to the Guard the benefits of integrating existing or readily available training environments, and to enable LVC interaction over non-dedicated wide-area networks (WANs).

JTEP is an experimentation program that will leverage the integration successes of other programs whenever possible, but will also advance the state of the art in system and simulation interoperability as needed to meet Guard training needs. JTEP started with an initial study to determine which candidate systems and integration mechanisms will achieve the greatest training impact. After the initial study, the first demonstration linking live and constructive training systems was conducted in May 2003 [1]. The second demonstration, conducted in December 2003, built on the successes of the first and the results of the initial systems analysis study.

The second demonstration provided a battalion-level training capability for the California National Guard by linking existing live, virtual, and constructive training systems [2]. In particular, JTEP linked two live training systems, the Deployable Force-on-Force Instrumented Range System (DFIRST™), which provides instrumentation and engagement simulation for ground vehicles, and the Integrated Global Positioning System (GPS) Radio System (IGRS), which provides tracking for dismounts and interface to the Multiple Integrated Laser Engagement System (MILES) 2000 for engagement simulation. Live entities were able to engage other live entities and constructive Joint Combat and Tactical Simulation (JCATS) entities. Additionally, the demonstration included the Close Combat Tactical Trainer (CCTT) virtual simulation, which was capable of engaging JCATS entities, and virtual-constructive Unmanned Aerial Vehicles (UAVs), which had a common view of the battlespace and all LVC entities [3,4]. The systems were integrated according to Distributed Interactive Simulation (DIS) protocols, and the participants at each simulation type (L, V, and C) communicated via DIS radios. In accordance with JTEP program goals, each demonstration is designed to establish an integrated LVC training capability that can serve as the basis of a leave-behind capability suitable for routine usage in training. Demo 2 provided the basis for a battalion-level LVC training capability [5]. For further information on the JTEP program see www.jtepforguard.com.

1.2 DAAR Goals

The first JTEP demonstration in May 2003 introduced a JTEP After Action Review (AAR) capability that included:

- Synchronized playback of recorded tracking and engagement data and tactical voice nets.
- A 2-D display that shows maneuvers and engagements against the background of a tactical map and maneuver planning graphics.
- A 3-D display showing the same data and graphics on geo-specific Camp Roberts terrain.

The second demonstration added a distributed After Action Review (DAAR) capability. Following the local unit-level AARs at each site, the Battalion Commander at Camp San Luis Obispo (SLO) led the DAAR. Command elements and observers at the main site and three remote sites were able to view exercise playback on the JTEP high-resolution 2-D and 3-D map displays, and to view and interact with the DAAR presenters via a standard video teleconference (VTC). Each site had two screens, one showing the VTC, and the other switching between 2-D and 3-D displays.

1.3 Scope and organization of this paper

This paper provides a technical overview of the DAAR capability that was developed for the second JTEP demonstration in December 2003. Section 2 describes DAAR functional and performance requirements and explains the JTEP system context in which the DAAR is embedded. Section 3 discusses the technical details of the DAAR design, and describes projected enhancements. Section 4 summarizes results.

---

[1] A live "simulation" comprises real people, real vehicles, real environment, and simulated weapons. A virtual simulation comprises real people, simulated vehicles, simulated environment, and simulated weapons. A constructive simulation comprises some real people, some simulated people, simulated vehicles, simulated environment, and simulated weapons.

2. DAAR Requirements

Figure 1. JTEP LVC Demonstration Components Mapped to Scenario Entities

The DAAR objective was to review the exercise conduct, including all constituent JTEP systems and components. Except for the VTC, the DAAR components were designed to operate using the machine and network assets of the live system. The JTEP Demo 2 components and network are shown in Figure 1. The major JTEP components in Demo 2 were as follows:

- JCATS, CCTT, DFIRST, and IGRS broadcast and received DIS entity state, weapons fire, and detonation protocol data units (PDUs).
- Voice communications were broadcast among systems as DIS signal PDUs. In the Tactical Operations Center (TOC) and at the JCATS stations, standard tactical handsets were connected to DIS radio interface boxes. CCTT had organic DIS radio equipment. In DFIRST, tactical SINCGARS radios were connected to the DIS radio interface boxes.
- The DFIRST and IGRS player instrumentation tracked participants in their true Camp Roberts coordinates. Player positions were translated to the coordinates of the stitched Grafenfels-Camp Roberts (GrafenBob) terrain in the DFIRST Base Station.
- The 2-D SRIDisplay and 3-D MetaVR Virtual Reality Scene Generator (VRSG) received DIS entity state, weapons fire, and detonation PDUs, and displayed entities and engagement effects on 2-D and 3-D representations of the stitched GrafenBob terrain.
- One MetaVR VRSG display was adapted for use as a virtual UAV. JCATS simulated the UAV flight path. The VRSG operator controlled the UAV's camera viewpoint.
- DFIRST software mediated indirect and direct fire engagements between DFIRST and JCATS. Since pop-up targets were modeled as JCATS entities, target up and down commands were associated with data sent from JCATS to DFIRST; i.e., a target emerging from defilade was raised, and a killed target was dropped. Commands were issued to personnel controlling the target via voice radio.
- In the CCTT mobile stations, JCATS computer-generated forces were used along with native CCTT semiautomated forces (SAF).
- The computers at Camp SLO, Camp Roberts, and OTAG were connected by a virtual local-area network (LAN) over the California Army National Guard (CA-ARNG) WAN. DataRouter software managed communications among sites.
- Because an Ethernet connection to the CA-ARNG network was not available at the DFIRST Base Station location, a wireless LAN was used to send and receive JTEP network traffic. A PIX501 Firewall was used to secure the CA-ARNG network against possible access intrusion from a wireless LAN eavesdropper.

To the above, the DAAR added a VTC component. The master DAAR site in Camp SLO and the remote site at OTAG were both nodes on the dedicated CA-ARNG Distance Learning Classroom (DLC) network. The networking center in Sacramento temporarily bridged the JTEP DLC nodes to the JTEP nodes on the CA-ARNG WAN to allow the non-DLC nodes to participate in the DAAR.

2.2 Functional requirements

Demo 2 DAAR requirements that were added to or different from Demo 1 local AAR requirements include the following:

- The master site is located in the California National Guard DLC at Camp SLO. Three additional participating sites are located at Camp SLO, Camp Roberts, and OTAG.
- Planned duration for DAAR is 20 minutes.
- The DAAR system is responsible for AAR data, while the organic VTC system is

- responsible for the human interaction portion of the DAAR.

- Each site will have two screens: (1) map/replay, and (2) VTC/whiteboard.

- Map display quality at remote sites must preserve the clarity and responsiveness it has at the originating site.

- The map/replay screen must be controlled from one location. Remote sites must see the same view at the same instant.

- The map/replay screen must be capable of displaying all LVC entities from all sites at up to 60× real time replay rate.

- Live and post-exercise annotations should be collected from multiple sources and sites. The presenter should be able to create a culled/ augmented set of annotations.

- The map/replay screen must be switchable quickly between 2-D and 3-D views (but the 2-D view is expected to dominate).

- Recorded voice from at least seven nets must be operator selectable for synchronized replay with display data.

- The DAAR may be played back through heterogeneous AAR/display systems.

2.3 Performance requirements

A primary technical challenge was to provide lossless remote viewing while staying within the limited bandwidth allocated to JTEP on the CA-ARNG network. The allocated bandwidth was adequate for transmitting DIS data and voice PDUs to all sites during live operations; however, the bandwidth was not sufficient to carry this load plus a VTC and the greater number of PDUs generated by fast-rate playback.

3. DAAR Technical Approach

3.1 Component architecture

In developing the DAAR architecture, we considered two options. One option was for participants at remote sites to view 2-D and 3-D map displays via the VTC. However, VTC resolution and update rate limitations would have resulted in an unacceptably degraded view. Instead, we opted to run synchronized local instances of the 2-D and 3-D displays at each site. Each of the four sites hosted a full complement of JTEP AAR software and hardware.

The DAAR component architecture is shown in Figure 2. The software architecture supports various machine configurations. It is not required that every site have the same number of machines. The SRI DataRouter facilitates assigning multiple functions to a single machine, for example, allowing multiple processes on the same machine to listen to the same port. DIS signal PDUs originating from numerous (e.g., seven) radio nets are redirected to two speakers.

The DAAR architecture included a remote viewing location inside the DFIRST Base Station/AAR Trailer located on the range at Camp Roberts. The same wireless LAN used to connect DFIRST to the rest of the JTEP network was used to provide the DAAR. This is significant because it demonstrates the ability for soldiers in the field to participate in a DAAR without going to a dedicated site on post. More time can be spent in the field training rather than commuting to and from such a site.

Implementing the JTEP DAAR architecture required solutions to several technical challenges, particularly in the areas of distributed logging and playback. Although the technical approach was tailored for the particular JTEP requirements and context, it may contribute a partial solution to general DAAR issues addressed recently in [6].

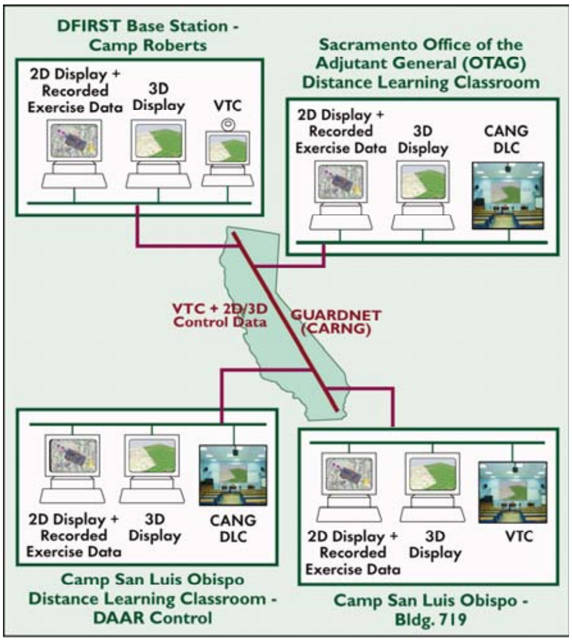

Figure 2. Distributed AAR Architecture

3.2 Data logging

3.2.1 Distributed logging

During the exercises, we logged the DIS PDU data and voice network traffic at each location (see also [7]).

Each DIS-enabled application broadcasts PDUs on the LAN. We used the SRI DataRouter to distribute the PDUs to remote locations. Because we configured the DataRouter to use the reliable TCP (Transmission Control Protocol) for long-haul WAN communications, all sites received identical data. A client DataRouter at each remote site rebroadcast the PDUs on its LAN. Because the best-effort UDP (User Datagram Protocol) broadcast protocol occasionally loses a packet, there were small differences among the distributed log files (see also [8]).

For this demo, it was feasible and appropriate to distribute all data to all sites for real-time consumption and logging. However, the JTEP distributed logging capability is designed for the more general case in which particular systems may be interested in only a portion of the network traffic, e.g., only entities with which it is interacting, and distribution of all data is not possible due to network constraints (see also [9]). In future demos, we anticipate applying site-specific interest criteria to filter voice and data traffic as needed.

The mirror log files at each site could have been used for distributed playback, but we opted to send a compressed master log of the entire exercise to each site at the end of each exercise. Our raw log file of data and voice PDUs compressed rather well using a standard pkzip compression to reduce its size by 75%, which is much better than compression as applied to individual messages because of the advantage of long run-time encoding. Since the exercises were approximately 1 hr 30 min in duration, and the interval between endex and DAAR was about an hour, this compression ratio made it feasible to transmit full replacement files. In future JTEP demos, we will create separate log files for data that are not distributed in real time to some of the sites. These files will be compressed and transmitted to the appropriate remote sites as fill-in data supplementing their local log files.

3.2.2 Logger design

The logger design was driven by the need to capture portions of the full data flow at different locations, compose them into a complete data set for DAAR playback, support centralized playback control of distributed files, and support direct access by time stamp. To meet these requirements, we leveraged services provided by the SRI DataRouter and built a generic logging and playback system, based on the data logging, indexing, and playback techniques employed in SRI-developed live training systems.

The primary job of the DataRouter is to distribute data among networks according to rules defined in the sender's XML configuration file, and to handle the data according to rules defined in the receiver's configuration file. It occurred to us that a log file can be regarded as just another routing destination, and so we designed the logger as a plug-in to the DataRouter. The DataRouter's XML configuration mechanism gave us a flexible method for defining rules at run time that direct different data to different files, which enables a fine-grained approach to distributing DAAR data logs. The rule sets supported for Demo 2 included source IP address and port, which we planned to use to separate voice and data traffic (see also [10]). However, because CCTT required that voice and data PDUs use the same port address, we changed the configuration to record all PDUs to the same file. In future demos, we will add flexibility to the rule sets, including filtering on selected message contents (e.g., message type).

The DataRouter approach also facilitates using a single machine for both logging and network routing when the host operating system does not allow two different applications to listen on the same port. The DataRouter was the single listener on the ports used for DIS PDU broadcast, serving both functions.

To support direct access to time-series data during playback, the logger captures the raw data in one file and stores indexing information in another. The index references messages in the log files to the system time on the logger machine at the time of capture. Because loggers at the remote sites do not receive PDUs at the same instant as the master logger, there are slight variances in the indexed times, but these variances are too small to make any perceptible difference in a DAAR.

The preferred method of time synchronization among the loggers is to synchronize to a standard internet time service. However, we were not able to use such a service in the demonstration because time service access to the internet was blocked by the CA-ARNG firewall. A secondary method assigns the local time server function to a machine that gets its time from a GPS receiver. This method was feasible for JTEP because the DFIRST live system provides GPS time, and we established a local virtual private network (VPN) among all sites that permitted use of the time service port. If no truth source is available, accuracy sufficient for DAAR playback coordination can be obtained by hand-setting the clock on the time server machine, or by simultaneous hand-setting of all logger machines if network time services are blocked.

3.3 Playback

3.3.1 Playback files and players

To minimize bandwidth usage during the DAAR, data files and players were located at each of the participating sites. In the JTEP playback architecture, only operator control and player synchronization are centralized, as explained in Sections 3.3.2 and 3.3.3.

At each site, a separate player instance reads and plays each file, under the control of a single play control process located at the master site. We have used this method of separating file players and play control in DFIRST, and found that it scales well and promotes software simplicity. It also supports playback composition in which players from heterogeneous systems are used together in synchrony. JTEP currently combines DFIRST and DIS playbacks, but the method could be extended to include any system that exposes a mechanism for external control of its playback. By using this method, a playback more faithfully replicates the conditions of the live exercise. For example, we have played back combined JCATS and DFIRST exercises side-by-side, with DFIRST using its own recorded data for its native entities and DIS-recorded data for external JCATS entities, and JCATS using only DIS-recorded data. We would have preferred JCATS to display its native entities from data derived from its own recorded file, but the JCATS playback does not have a mechanism for external control and therefore does not support playback composition.

DFIRST-in-JTEP provides an example of the potential advantages of composing playbacks from heterogeneous native recorded files. DFIRST-native data are preferred for DFIRST entities because some information is lost in the DIS PDUs published for use by external systems. For example, DIS entity state PDUs do not report data quality, i.e., all data are assumed to be reliable, whereas data for live instrumented entities are occasionally unreliable, e.g., loss of GPS-derived gun pointing angle due to a GPS receiver hardware failure.

The DIS and DFIRST recorded files have overlapping data for DFIRST entities. To avoid displaying duplicate data from the non-preferred source during playback, an argument is passed to SRIDisplay instructing it to suppress DIS PDUs originating from DFIRST. In the future, we may use a refined log routing control mechanism, as explained in Section 3.2.2, to route potentially duplicative data to separate files so that playback composition may be effected simply by selecting which file player instances are started.

It would be possible to support a DAAR with different player configurations at different sites; for example, the audience at a CCTT locale might see the DIS recording of DFIRST data, while the DFIRST locale sees the native DFIRST recording. However, such heterogeneity is probably not desirable because it could lead to situations in which the presenter refers to something on his display that differs from what audiences at remote locations are seeing.

3.3.2 Playback operational control

The operator controls for the DAAR playback are integrated with the SRIDisplay toolbar and menus, but they could be extracted from this context and used as a stand-alone player. The playback time controls include direct selection using sliders, jump forward and backward by 1 or 5 minutes (configurable), and jump-to times associated with annotated events. We plan to add 2-D and 3-D map view specification as an option for annotations in the future. Annotations can be entered during live operations or playback. Voice channels may be selected from a menu that includes information identifying the contents of each channel. The playback rate may be set anywhere between 1× and 60×.

For Demo 2, we used startup scripts to give playback controls uniquely to the instance of SRIDisplay that was running in the Camp SLO DLC. In the future we will add a mechanism for passing a token among multiple playback-enabled instances so that playback control can be shared among sites.

For smooth DAAR conduct, it is essential that control segues be rapid and transparent to the presenter. During the preparation period between endex and DAAR, the system operator usually plays back the exercise for the presenter at a fast rate. Annotations are entered, and associated 2-D and 3-D viewpoints are specified.

3.3.3 Playback synchronization

When the operator selects a time, the JTEP playback control process synchronizes the players by sending an initial time and then a sequence of time pulse messages. Players interpret each time pulse as a “play-to” directive. They play all data in the interval between the last pulse received (or the initial time) and the current pulse. Playback pauses or stops when the flow of time sync messages stops. The granularity of the intervals is configurable—for example, once per second if the finest granularity of indexed time stamps among the log files is 1 s, or 10 times per second if the granularity is finer than 1 s.

It is not necessary to send rate information to players because fast playback is effected simply by scaling the play-to times in the pulse messages. For example, at 10× rate, the interval between the play-to time in two successive time pulse messages sent by the JTEP controller is 10× as large as in normal speed playback. An alternative method would be to send the same messages as with 1× rate, but reduce the spacing between the messages by a factor of 10.

The control messages sent to JTEP DIS log players are formatted as DIS SetData PDUs. Control messages sent to DFIRST players are in a native DFIRST message format. In the future, JTEP will add support for other interoperability mechanisms, e.g., the Test and Training Enabling Architecture (TENA).

The initial start time should be sent reliably, but time pulses can be sent best effort because it does not matter if a message is occasionally lost in transmission. Because DIS uses best effort UDP rather than reliable TCP, and messages are occasionally lost, we include a redundant copy of the current initial time or "jump-to" time in all time pulse messages.

During the Demo 2 DAAR, only control information was sent over the network, including the messages described above, and the 2-D and 3-D view control messages described in Section 3.4. Control PDUs consumed very little bandwidth, leaving most of it available for the VTC. This one controller/many players approach succeeded in seamlessly synchronizing the operation of the distributed playback and display components.

3.4 Map viewpoint control

The JTEP 2-D and 3-D displays at remote sites were slaved to masters located at the Camp SLO DLC (see also [11]). We adapted the 2-D SRIDisplay to allow a designated master to send messages reporting map center and zoom factor. Slave displays recenter and resize their views accordingly. In the future, we will add other controls, such as filtering specifications that control what entities are shown or hidden on the display.

We configured the MetaVR VRSG 3-D display at the Camp SLO DLC as a master display that exports its "camera" viewpoint, which the operator manipulates using a controller device. The eyepoint is reported in DIS entity state PDUs as position and orientation. The remote displays are configured to set their eyepoint to mimic the fictitious entity that represents the master's eyepoint.

Because the 3-D eyepoint is sent as a standard entity state PDU, it will be displayed on the 2-D map display unless specifically hidden. Although it can sometimes be useful to operators to see the location of the 3-D eyepoint represented on the 2-D display, the 3-D operator's control actions may cause it to move rapidly across the map; we assigned the entity a special DIS "enumeration" that we could categorically "hide" on the display so as not to distract the audience.

The fixed master-slave relationship between the Camp SLO DLC and remote sites met the DAAR requirement for Demo 2, but in the future we would like to have a mechanism for passing control between sites. For the 2-D SRIDisplay we can adopt a "token" passing mechanism similar to the one we mentioned in Section 3.3.3. Master-slave relationships are currently set by the operator in MetaVR VRSG 3-D display, and so vendor modifications will be needed if we choose to implement control switching in the future.

In the Demo 2 DAAR, the 2-D and 3-D displays shared a single screen. A KVM (keyboard-video-mouse) switch was used to control which display was seen by the audience. At each remote site, an operator was cued by a voice signal to throw the switch. We relied on the VTC to deliver the signals. An improvement planned for the future is single-action switching effected by a message sent from the controlling site to each remote site.

During the Demo 2 DAAR, the presenter used a laser pointing device to draw attention to particular areas and events of interest on the 2-D display. However, the pointing was visible to remote audiences only through the VTC. Future improvement plans include replacing the laser pointer with an electronic device such as a wireless mouse controller. The master display's mouse pointer location could be sent to remote sites and echoed on the slave displays. Remote audiences would be able to follow along as the presenter moves the mouse pointer.

3.5 VTC

The DAAR used organic DLC VTC appliances and associated networking services and protocols. However, only two of the four remote viewing locations were DLC classrooms, and therefore the other two sites were not nodes on the dedicated DLC network. Guard networking personnel in Sacramento solved this problem by bridging from the DLC network to CA-ARNG network to bring in non-DLC sites. Our testing was limited to a few prearranged intervals, which were dominated by troubleshooting an apparent incompatibility between the VTC protocol used by the DLC VTC appliances and the CISCO PIX routers that were used to encrypt data over the microwave link that connected the DFIRST Base Station to the Camp Roberts CA-ARNG network infrastructure.

Insufficient rehearsal with the VTC appliances resulted in a presentation that was less effective than planned. Most importantly, the microphone setup used on the day of the demonstration made it difficult for remote audiences to hear the main presenter in the Camp SLO DLC, and it was overbalanced in favor of the Camp Roberts site. In addition, we failed to detect a resolution mismatch between heterogeneous equipment, resulting in a reduced size image at one of the remote sites. Also, a standard practice used in most VTCs was not ideal for our two-screen approach. The view angle of the VTC camera at the master site took in both the presenter and the 2-D/3-D map display. Some observers found it disconcerting to see the high-resolution local 2-D/3-D display on one screen side-by-side with a second screen showing a low-resolution view of the master displays behind the presenter.

Future JTEP DAARs are likely to include participating sites that are not DLC classrooms. Since VTC is a commodity, it may be easier both operationally and technically to provide a lightweight organic JTEP VTC capability.

3.6 Takehome package

The JTEP Takehome package contains all the DIS and DFIRST log files recorded during the exercise for playback on a user's PC. It is delivered on a CD-ROM and installed using the InstallShield™ setupwin32.exe program. In addition to the log data, the installation includes the JTEP Takehome software, 2-D maps for the exercise area, and the Java runtime if a compatible version is not detected on the user's computer.

The package includes a launcher application that allows a user to select an exercise. Once a user starts the exercise, a 2-D SRIDisplay map appears. All playback controls used by the DAAR operator are included on the map display's toolbar.

The voice playback software in the Takehome package is not the same as the DAAR software. Due to integration compatibility issues with CCTT, all DIS radios in Demo 2 used CVSD (continuously variable slope delta modulation) encoding. At present, we are unable to distribute a license-free CVSD compatible DIS radio player. Therefore, we have included a µlaw player and converted the recorded voice data to work with this player.

3-D viewing of the Demo 2 exercise is possible if the user has a license for a compatible viewer, such as MetaVR VRSG or MaK Stealth Observer. This capability would also require a separate distribution of 3-D terrain data. The terrain data files are too large to be included on the Takehome CD, and need to be customized for the capabilities of the particular viewers.

4. Summary and conclusions

The second JTEP demonstration, an integrated LVC exercise, was conducted on December 11, 2003 at Camp Roberts, Camp SLO, and OTAG in Sacramento, California. This demonstration concluded with a DAAR comprising a VTC and synchronized playback of the recorded tracking, engagement, and voice net data. Maneuvers and engagements by constructive JCATS, virtual CCTT, live DFIRST, and live IGRS players were displayed on high-resolution 2-D and 3-D displays at each of the four participating DAAR sites.

The key technical challenge was to provide a lossless view of the exercise at all sites while obeying the constraints of a limited inter-site bandwidth allocation on the CA-ARNG network. This paper describes the technical design of the DAAR, including distributed logging, synchronized playback of distributed logs, and viewpoint control of the distributed 2-D and 3-D displays. The DAAR worked as planned, except for the VTC, which was insufficiently tested and rehearsed.

This paper also describes enhancements that are planned in each of the technical areas. Some of these are designed to improve JTEP support for exercises similar to Demo 2, and others will be needed to meet the next year's goals in JTEP's two current major areas of focus: (1) Combat JTEP, which supports combined arms training exercises and (2) JTEP Military Support for Civilian Authority (MSCA), which enhances training for the Guard's mission to support civilian authority in natural disasters and homeland defense. It is anticipated that both areas will generate additional requirements for the DAAR.

5. References

[1] Reginald Ford, John Shockley, Michael Beebe, Mark Faust, Gerald Lucha, Mark Johnson, and John Bernatz., "The Joint Training Experimentation Program: Lessons Learned from the First Demonstration," 03F-SIW-058, Simulation Interoperability Workshop, September 2003.

[2] John Shockley, Reginald Ford, Michael Beebe, Gerald Lucha, Mark Johnson, and John Bernatz, "The Joint Training Experimentation Program: Hotwash from the Second Demonstration," 04S-SIW-132, Simulation Interoperability Workshop, April 2004.

[3] Mark Johnson, Richard Giuli, Scott Oberg, Michael Beebe, Reginald Ford, and John Shockley, "Integration of CCTT and JCATS in an LVC Exercise," 04E-SIW-066, Simulation Interoperability Workshop, June 2004.

[4] Michael Beebe, Mark Faust, Reginald Ford, John Shockley, Mark Johnson, "The Creation of GrafenBob for the JTEP LVC Demo," 04E-SIW-069, Simulation Interoperability Workshop, June 2004.

[5] John Bernatz and John Shockley, "A Funny Thing Happened on the Way to LVC Integration: Great Training," 04E-SIW-065, Simulation Interoperability Workshop, June 2004.

[6] David Fisher, James Bryan, and Terry McDermott, "Distributed Simulation Synchronization Replay Tool (DSSRT)," 04S-SIW-033, Simulation Interoperability Workshop, April 2004.

[7] Daniel Van Hook and James Calvin, "Execution Logging and Replay: Issues and Approaches," 97F-SIW-118, Simulation Interoperability Workshop, 1997.

[8] Gerry Magee, Graham Shanks, "Lessons Learned from an Implementation of a Fully Distributed Data Collection Tool," 99S-SIW-084, 1999.

[9] J. Mark Pullen and Vincent Laviano, "Adding Congestion Control to the Selectively Reliable Transmission Protocol for Large-Scale Distributed Simulation," 97F-SIW-018, Simulation Interoperability Workshop, 1997.

[10] Robert Kerr and Christopher Dobosz, "Reduction of PDU Filtering Via Multiple UDP Ports," 13th DIS Workshop, 13-95-05, 10 March 2002.

[11] Kevin Mullally and Andy Cox, "The National Counterdrug Center – Leveraging Military Simulation Technology for Law Enforcement Agency Interoperability Training," 01F-SIW-072, Simulation Interoperability Workshop, 2001.

Author Biographies

REGINALD FORD, Software Development Manager at SRI International, has 24 years of experience in test and training range instrumentation systems for the Army, Navy, Air Force, and Marine Corps. He helped establish SRI's Software Engineering and Development Program and is a member of the Instrumentation and Simulation Program. He manages DFIRST software development and the integration of JTEP software systems and components.

RICHARD GIULI, Software Engineer at SRI International, is a member of the Software Engineering and Development Program and has 6 years experience in software engineering. On JTEP he helped develop the distributed networking, VPN and mapping display.

VICTORIA LAMAR, Senior Research Engineer at SRI International, is a member of the Instrumentation and Simulation Program, and has 17 years of experience in test and training range instrumentation systems. She oversees the JTEP MSCA effort. She developed requirements for the combat JTEP DAAR and coordinated the VTC.

SCOTT OBERG, Software Engineer at SRI International, is a member of the Software Engineering and Development Program and has experience in building and integrating heterogeneous distributed systems. For JTEP he focuses on the logging of simulation data, DAAR, routing and compression of simulation data across the WAN, and the network architecture for the interoperation of JTEP systems.