LESSONS LEARNED FROM THE SPECIAL OPERATIONS FORCES STOW-A HLA EXERCISE

Gilbert González
Science Applications International Corporation
Orlando, FL

Ivan Carbia, Gene Lowe, Glenn Valentine, John Bray
Science Applications International Corporation
Orlando, FL

Víctor Colón
Naval Air Warfare Center Training Systems Division
Orlando, FL

Abstract

The U. S. Special Operations Forces (SOF) conducted their second Synthetic Theater of War-Architecture (STOW-A) exercise during the week of October 25th, 1999. The first STOW-A exercise utilized the Distributed Interactive Simulation (DIS) protocol; the second exercise was designed and executed using the High Level Architecture (HLA). The exercise consisted of two missions, which were classic infiltrate and assault using air and ground assets, requiring joint coordination between the Air Force and Army pilots, Rangers and the other SOF units.

This exercise had two primary objectives. The first objective was to conduct a realistic Computer Aided Exercise (CAX) using manned simulators from the 160th Special Operations Aviation Regiment (SOAR) at Ft. Campbell, KY and the 19th Special Operations Squadron (SOS) at Hurlburt Field, FL. The second was to establish a distributed training architecture using HLA that could be used to refine and validate tactics for multi-aircraft, all-weather operations. The exercise achieved the goals and objectives of all the participants.

This paper captures the lessons learned during the integration, testing and execution of the SOF STOW-A training exercise using HLA. We address the technical challenges the federation developers overcame related to Run-Time Infrastructure (RTI) connectivity over a Wide Area Network (WAN), and the use of a DIS filter and gateway to integrate radios and non-HLA simulation applications into the federation execution. We discuss the elements used to coordinate and execute the federation between distributed sites.

Author Biographies

Gilbert González is a Project Director for SAIC and a member of the ADST II HLA Team. He has been involved in simulation development and distributed simulation technologies since 1987. Mr. Gonzalez received a BS degree in Computer Engineering from the University of Central Florida.

Ivan Carbia, John Bray, and Dr. Gene Lowe are employed by SAIC and are members of the ADST II HLA Team. Mr. Carbia is a Senior Software Engineer, and holds a BS degree in Computer Science and a BA degree in Organizational Communication from the University of Central Florida. He has been involved in distributed simulation technologies since 1997. Mr. Bray is a Project Director, and holds a BS degree in Computer Science from Rowan University. He has been working in the simulation industry since 1985. Dr. Lowe is a Project Director, and holds a Ph.D. in Electrical Engineering from Georgia Tech. He has 30 years of experience in systems engineering and simulation and 25 years in project management.

Victor M. Colón is an Electronics Engineer at the Naval Air Warfare Center Training Systems Division in Orlando, Florida. Mr. Colón has a BSEE in Electrical Engineering from the University of Puerto Rico. He has been working in simulation since 1995.

LESSONS LEARNED FROM THE SPECIAL OPERATIONS FORCES STOW-A HLA EXERCISE

Gilbert González
Science Applications International Corporation
Orlando, FL

Ivan Carbia, Gene Lowe, Glenn Valentine, John Bray
Science Applications International Corporation
Orlando, FL

Víctor Colón
Naval Air Warfare Center Training Systems Division
Orlando, FL

INTRODUCTION

Special Operations Forces (SOF) execute Joint Tactical Missions in conditions of restricted visibility and darkness. These missions require a mix of specialized personnel, state-of–the–art equipment, and unique tactics and training. Virtual environments, such as the one created for Synthetic Theater of War-Architecture (STOW-A), provide the commander with a robust training environment for these diverse missions. One aspect of the Synthetic Environment is the ability to network man-in-the-loop training simulators, computer generated forces, exercise tools, and supporting simulation applications to create a complete interactive training environment. To achieve their simulation interoperability requirements, the STOW-A program chose the High Level Architecture (HLA).

This paper presents the lessons learned during the integration, testing and execution of the SOF STOW-A training exercise using HLA. It also addresses the technical challenges the federation developers overcame related to Run-Time Infrastructure (RTI) connectivity over a Wide Area Network (WAN), and the use of a Distributed Interactive Simulation (DIS) filter and gateway to integrate radios and non-HLA simulation applications into the federation execution. Additionally, the elements used to coordinate and execute the federation between distributed sites will be highlighted.

STOW-A Background and Purpose

STOW-A was the first large-scale training exercise supported by HLA. The mission was an infiltrate and assault mission to extract precious cargo from enemy forces using U.S. Special Operations air and ground assets. The actions required joint coordination between the Air Force pilots, Army Special Forces (Aviation), Rangers, and Special Operations units. A goal of the STOW-A program was to assemble an infrastructure that linked existing simulators and simulations to support a mission planning exercise using new and emerging technologies. HLA was one of the new technologies being evaluated.

The 160th Special Operations Aviation Regiment (SOAR) at Ft. Campbell, KY and the 19th Special Operations Squadron (SOS) at Hurlburt Field, FL had migrated their training devices and supporting simulation tools to HLA. HLA had been thoroughly tested at each site as a standalone system, but the SOF sites had not used this capability in a training exercise over a WAN. STOW-A provided the opportunity to test this new technology for the first time, using a real-world training mission, under rigorous operational conditions.

STOW-A Participants and Locations

Participants for the STOW-A exercise were distributed among two sites in different states: Ft. Campbell, KY and Hurlburt Field, FL. These sites were connected over a dedicated T-1 line using KIV-7HS encryption devices.

The architecture revolved around legacy training simulators: the MC-130E Weapon System Trainer (WST) at Hurlburt Field, FL and the MH-60K Blackhawk and MH-47E Chinook Combat Mission Simulators (CMS) at Ft. Campbell, KY. All three devices are full-motion high fidelity virtual training devices outfitted with sophisticated avionics, radar and

Forward Looking Infrared (FLIR) systems, which replicate the full functionality of the particular aircraft.

These three devices are part of the SOF Federation (figure 1) which also includes the AC-130U Battle Management Center (BMC), the MC-130H Mission Rehearsal Device (MRD), and several Commercial and Government Off-The-Shelf (COTS/GOTS) simulation tools such as stealth viewers, data loggers, computer generated forces (CGF), simulated radios, and HLA gateways. The five SOF Federation training devices had recently achieved HLA compliance using RTI 1.3v6 and the SOF Federation Object Model (FOM) 0.5 (based on RPR-FOM 0.5). All five systems had been migrated to HLA using a middleware approach, while the DIS GOTS/COTS simulation tools utilized the STRICOM Gateway. The gateway translated DIS Protocol Data Unit (PDUs) into the corresponding HLA service calls and data, and vice versa.

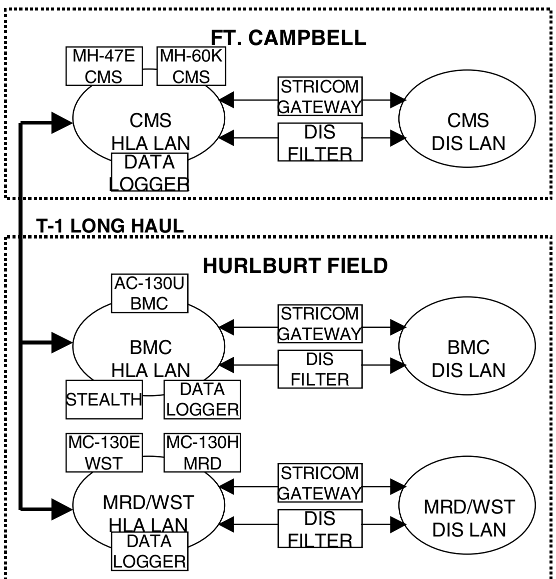

Figure 1. SOF Federation

LESSONS LEARNED

HLA is a robust technology that provides tremendous flexibility for developing distributed simulation network architectures. As a relatively new technology, its use often leads to the discovery of new aspects during implementation and creates new challenges for those who are designing and integrating an HLA long haul exercise for the first time, such as the STOW-A HLA team. During the integration of the STOW-A exercise the HLA team learned a number of useful lessons on how to design, integrate and execute a long haul HLA training exercise. Among the more significant issues and concerns encountered by the team:

- HLA Connectivity over a WAN
- Testing the Federation
- Radio Communication in HLA
- Managing HLA Expectations
- System Performance
- Conducting the Exercise

These topics were dealt with as part of requirements analysis, design, integration, test, and exercise execution.

HLA Connectivity over a WAN

Connecting the SOF federates over a WAN was quite a challenge for the HLA team. With a tight integration and testing schedule for the STOW-A exercise, the HLA team decided to test the long haul connection prior to the first Sub-System Evaluation (SSE). On the first attempt, the HLA team had to deal with T-1 line connection issues, which were not related to HLA. The T-1 line between the two sites had not been tested since the previous STOW-A exercise a year earlier. Following several failed attempts, we determined the problem was a faulty connection at one of the sites. After the connection was fixed we could not get the encryption keys synchronized between the two sites. This problem also took several attempts to fix before security at both sites developed a process that worked. Both of these problems consumed valuable integration time and budget that were planned for subsystem testing between the two sites.

Lesson Learned: If a T-1 link is not used often, the sites should schedule regular connectivity tests. This will help avoid unnecessary loss of time during integration and potential cost impact to the program.

Another challenge the HLA team encountered was related to the RTI implementation of Internet Protocol (IP) communication. Traditional IP communication allows a host to send packets to a single host (unicast transmission) or to all hosts (broadcast transmission). By default the RTI communicates using IP multicast transmission for best-effort messages, which allows a federate to send a packet to a subset of all federates. The subset of federates belong to a member group. Packets delivered to the member group are identified by a single multicast group address. Multicast packets are delivered to a group using best-effort reliability just like IP unicast packets. Initially, this was not a problem because all integration had been done on a Local Area Network (LAN) with all hosts connected to a single hub. IP multicast became an issue when federates had to join over a WAN (see figure 2). Our first attempt to have federates join the federation over a WAN was unsuccessful; federates at Hurlburt Field could not join the federation created at Ft. Campbell. This prompted the HLA team to perform some basic trouble shooting such as verifying that all federates had a common RTI message version and that IPs were set correctly, but no problems were found. A ping test, which is a unicast transmission, worked over the WAN, which meant there was some form of connectivity between the sites. After consulting with other network engineers in our team, we were able to determine that in order to do multicast transmission, routers on both sides had to be capable of IP multicast routing. CISCO confirmed that the routers in question supported IP multicast routing but were not configured properly. After making the configuration changes IP multicast routing still did not work.

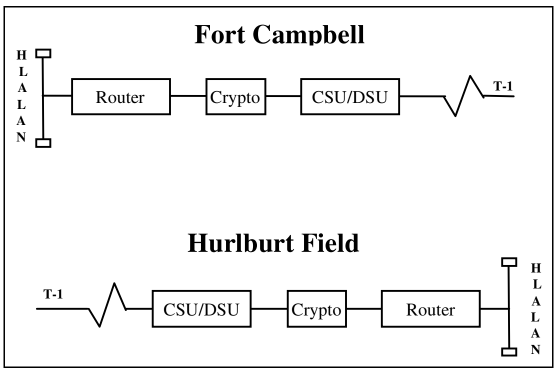

Figure 2. Network Connectiviity Architecture

At this point the Defense Modeling & Simulation Office (DMSO) HLA help desk was contacted for support. They advised us to configure the RTI with a reliable distributor. This configuration creates a point-to-point (unicast transmission) connection between a federate, which is the reliable distributor, and all other federates in the federation. This required the RTI.rid file of all federates, with the exception of the reliable distributor federate, to be modified. It also required that all data in the .fed file be set to reliable. After making these changes we were able to have all federates join a common federation execution over the T-1 line. One disadvantage of this configuration was the increased network traffic. Every message sent went through the reliable distributor and then was sent to each federate individually, plus an additional message was returned acknowledging receipt of the original message. This differs from multicast, which only requires the publishing federate to send a message once for all federates subscribed to receive the message (i.e. best-effort). Initial testing of this configuration confirmed that performance was not a problem. However, our customer required that we make the multicast configuration work to reduce network traffic in case additional bandwidth would be needed during peak traffic. After several experiments and analyses we discovered an error in the router software which prevented IP multicast routing from working. After updating the software, we were able use multicast IP and best-effort reliability. This configuration generated less traffic on the network, which provided more network bandwidth during heavy WAN traffic loads. The WAN performance and reliability during the exercise was exceptional. This would not have been a difficult problem to overcome if we had prior knowledge on some of the details of IP multicast and CISCO routers.

Lesson Learned: If an HLA exercise is distributed over a WAN and the RTI is configured to communicate using a multicast protocol, the routers used at each site will be required to support IP multicast routing. Otherwise, the RTI will have to be configured with a reliable distributor.

Testing the Federation

Testing the HLA functionality and performance became an important part of the STOW-A integration effort. The decision to use HLA as the interoperability solution had not been finalized. The migration of the Ft. Campbell devices to HLA had been completed but was still undergoing performance analysis during the planning stages of the exercise. The exercise would be executed using DIS if HLA could not meet the reliability and performance requirements established by the exercise objectives. A decision to use HLA or DIS would be made at the second SSE. With so much on the line, it was essential to assess the robustness and performance limits of HLA. A large part of this effort had been done during the initial phases of the HLA migration of the SOF training devices, but more rigorous testing involving more entities with the full system configuration was required.

The initial testing focused on verifying that federates were sending and receiving the correct attribute and parameter values. This was followed by reliability testing and performance testing. The reliability testing determined how long the federation could stay up and running before it became unstable. The federation could become unstable if several federates were terminated improperly, leaving orphan objects in the federation. At this point, the RTI would be recycled and the federation restarted. This was an uncommon occurrence and only happened once during the exercise. Verifying performance required stressing the system by increasing the entity and radio traffic. Replicating the actual exercise load on the test network was difficult because not all systems were integrated. For example, while forty radios were required for the exercise, we were limited to two or three radios during testing. We used other techniques to heavily load the system, such as increasing the entity count from 250 entities to 700 entities. With 700 entities distributed over the WAN, no latency or performance degradation was noticed on the simulators, ModSAF, or the ASTi radios.

Lesson Learned: This type of testing validated, with a reasonable degree of confidence, that the HLA network would achieve the performance required for the STOW-A exercise.

Radio Communication in HLA

For the STOW-A exercise, the communication structure consisted of both simulated and live radios. These radios provided the communication between the simulators, ground unit commanders and the live Command and Control (C2) helicopter. The simulated radio used was the Digital Audio and Communication System (DACS) made by ASTi. To bridge the live radio to the simulation radio network a product made by SimPhonics was used. At the time of the exercise both of these systems were DIS compliant and neither had an HLA interface. This created a unique problem for the HLA team. Even though the STRICOM Gateway could translate DIS radio PDUs to HLA and vice versa, previous experiments conducted in the Advanced Distributed Simulation Technology (ADST) II HLA Testbed revealed that reliable communication could not be maintained with a large volume of voice traffic. This performance problem was not an RTI issue but a limitation of the gateway architecture. Therefore, the HLA team decided, early in the STOW-A program, not to network the radios using the gateway. Because of this decision the gateway at each site was configured to ignore radio PDUs. The only alternative was to allow the simulated radios to communicate between the two sites using DIS.

The problem then became how to transfer the radio DIS traffic between the local DIS networks and over the HLA network. The DIS network at each site had a different IP address group that did not match the HLA network IP address group. To achieve this transfer of data packets between networks we used an application called the Protocol Cell Interface Adapter Unit (XCIAU) (also know as the DIS Filter). The XCIAU software, which runs on an SGI system, was developed by STRICOM to filter various types of PDUs between two networks. Figure 3 illustrates how the application works. The XCIAU filter in site A receives PDU X of type "signal" from the DIS network. The application will repackage the PDU and give it an address for the HLA network. The PDU will travel through the HLA network, in our case the T-1 line in a long haul exercise. The XCIAU filter in site B receives the signal PDU from the HLA network. The application will then repackage the PDU and give it an address of the DIS network. DIS applications in site B will now be able to process the signal PDU. This implementation provides the optimal solution to running the DIS radio simulation in an HLA exercise.

Lesson Learned: Most training environments will have GOTS/COTS tools that have not migrated to HLA. Therefore, to achieve a complete system solution, bridging technology that allows using the GOTS/COTS tools in an HLA environment will be required.

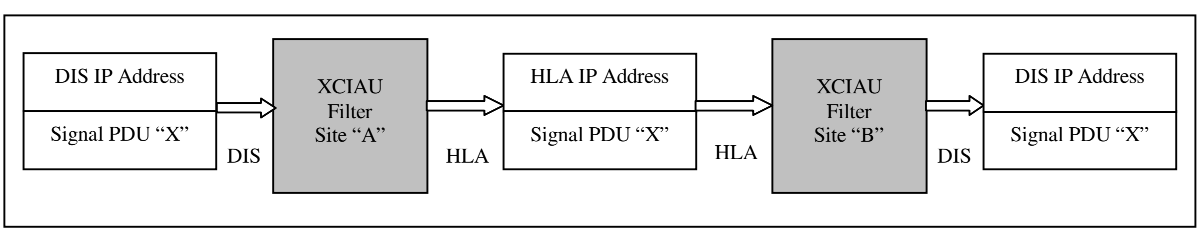

Figure 3. XCIAU Radio Filtering Process

Managing HLA Expectations

HLA is a relatively new technology in the distributed simulation training community. There are still many users who are reluctant to embrace this new technology. In many cases they are satisfied with their current DIS capability and do not see the benefits of HLA. The STOW-A user was no exception. The DIS implementation had recently been completed and replacing it with HLA before a major training exercise was perceived to be a high risk. Their number one priority was to train the soldier for this mission. With all this anxiety, the HLA effort was under the microscope. As simulator upgrades were completed and additional simulation systems were integrated into the training environment during this time, problems would surface. For the most part, whenever problems occurred, there was a tendency to perceive that HLA might be the cause. This forced the HLA team to do additional testing and analysis just to prove that HLA was not the culprit.

For example, during SSE II, a Semi-Automated Forces (SAF) operator unintentionally overloaded his SAF system with several hundred entities. The SAF station began to "gasp". During this "gasp" period, the SAF station stopped sending network PDU data. As a result, the gateway followed the DIS standard for entity time out at twelve seconds. The gateway then issued remove object calls to the RTI. No sooner had all the objects been removed from the HLA network, than the SAF station returned the entities to the network, causing the gateway to rediscover the entities all over again. The SAF station gasped again, causing the problem to return. This process continued for approximately fifteen minutes. Visual latency and entities flashing on the view monitor were the only sign of a problem during the rehearsal. The only way to pinpoint the problem (since the SAF operator was unaware of the problem) was during a network entity check done every 15 minutes. It was then discovered that approximately 600 entities were participating in the exercise. Parser prints on a federate system showed the entity count going up and down like a roller coaster ride. The problem was then narrowed to the SAF station. This process of removal and discovery of objects also caused orphan entities to exist, and a decision was made to recycle the federation. Recycling the federation took no more than two minutes.

Over time the customer and other integration teams became more and more confident with the HLA capability and HLA was seen as one of the most stable components in the system.

Lesson Learned: As an HLA integrator, your team will need to understand the big picture of how the different systems interact with each other and that problems are not automatically HLA problems.

System Performance

Performance measurements of HLA during the STOW-A exercise were limited due to time constraints and scope of effort. The performance requirements in HLA for the STOW-A exercise were twofold. First, the out-the-window view of all training systems must show no visual stepping. Second, the network and federates had to handle somewhere in the range of 250 to 300 entities and forty radios with no significant audio degradation.

The HLA interface units of the MH-47E and MH-60K simulators require a 60 Hz update rate with the simulator host. The HLA interface unit of the MC-130E simulator requires a less strict update rate of 10 Hz with the simulator host. Performance analysis, therefore, focused on the MH-47E and MH-60K simulators. Entity load conditions in excess of those expected for STOW-A were simulated using ModSAF and the STRICOM Gateway. Functional federation operations under these loads were confirmed. Performance analysis of these load conditions did reveal an issue with entity updates sent to the respective image generators of the Ft. Campbell simulators.

The 60 Hz cycle rate proved to be challenging to meet for both the DIS and HLA versions of the interface units. While performance improvements in the HLA interface source code yielded good results, it was determined that too much processor time was being spent on processing network entity data. To further reduce this processing time, several filters were implemented. Only entities within a user-defined range from the federate were processed. An additional entity enumeration filter prevented processing of specified entities in a database. It is important to note that data distribution management proved to be too challenging for the time allotted to prepare for this exercise.

The ability to handle a high entity count is a concern for all network protocols. The HLA migration team identified the STRICOM Gateway as both the key to interoperability as well as the source of possible failure. According to the gateway developers, the gateway operating on an SGI Indigo2 10000 chip requires approximately one to two milliseconds per translation, or five hundred to a thousand translation events per second, assuming a constant distribution of less than 200 bytes of data in both protocol directions. Larger data sizes and dynamic bursts of data distribution will add latency to these measurements. The team tested several configurations involving SAF entities in DIS translated to HLA through the gateway, which then interacted with several federates including an HLA Logger, a MetaVR Stealth, and the training simulators. Countless hours were spent over 14 months testing and debugging the gateway. As a result, tests on the maximum number of entities were very encouraging. The network and simulation systems were able to handle 600 to 700 entities without any noticeable latency.

Lesson Learned: It is important to identify the federation requirements for execution performance in the planning phase of an exercise. This will help address any limitations of participating federates and possible solutions that could be implemented to meet the performance requirements.

Exercise Management

Critical to running an HLA exercise is having a specific set of procedures that covers the steps required to manage the federation execution from beginning to end. The execution procedures, which can be part of the Network Integration Plan (NIP), should also address how communication will be established between sites to coordinate federation management functions. The procedures should identify the following:

- Which site(s) will start the RTI exec and the Fedex
- Which system(s) will run the RTI exec and the Fedex run on
- Which federate will be the reliable distributor (if this configuration is used)
- In what order will the federates join the federation
- How will federates be removed from the federation if they are not responding
- How will the federation be recycled if it becomes unstable

Communication among HLA distributed federates is critical, especially during the join process and when federates are experiencing problems. This is due in part to the centralized nature of the RTI and Fedex applications, which are only observable by the federate manager. Federates lack the ability to monitor and perform any manual intervention necessary during the course of their participation in the federation. For this reason, it is important that good communications and coordination exist between federates and the federation manager who oversees the RTI and FEDEX. For STOW-A, federation management communications were supported in several ways. Standard non-secure phone lines were used, but they had several shortcomings. The phone lines were not always advantageously placed with respect to where the federation workstations were located and phone access was also shared with other personnel. A major disadvantage was the non-secure aspect, which required users to "edit" their conversations in order to avoid disclosure of classified data. This was particularly troublesome during pre-exercise database and entity testing. In addition to phone lines, beepers were also distributed to federate managers and controllers and backup or alternate phone lines were identified. An innovative and simple solution to the phone security and location problem was the use of UNIX talk windows. The UNIX talk windows had several advantages over the less convenient and more restrictive phone communications. Since network connectivity to all sites went through encryption devices, classified information could be transmitted via the talk window. Talk windows were also self-documenting in that prior parts of a conversation could be recalled simply by scrolling back up the screen. The talk windows were opened at the federation manager and federate controller's workstations and were not utilized by other personnel. While talk windows will facilitate federate manager to federate controller "point-to-point" communications, they do not provide a chat type forum in which multiple federates can communicate among themselves.

Lesson Learned: Execution procedures should be defined and agreed upon by all parties. A reliable user friendly and effective communication strategy must also be decided upon and implemented in order to facilitate the communication essential to execute the HLA exercise.

Conclusion

Performing an HLA long haul exercise requires extensive planning, a strong technical team, and a cooperative team approach. Programs that are planning to run an HLA exercise over a WAN need to be aware of the type of IP protocols used by the RTI and how this can affect a network router configuration. They should also understand that they would have to use HLA to DIS bridging technologies to integrate GOTS/COTS systems that have not migrated to HLA.

The STOW-A program was successful in demonstrating that HLA can be used to link, over a WAN, heterogeneous simulations in a simulation based training environment. But like any new and emerging technology there is still a lot to be learned about how to exploit HLA. Probably the only way this technology will become more main stream, in the simulation based training community, is to have more large-scale exercises using HLA as the interoperability solution.

HLA is a step forward in distributed simulation technology and should be taken seriously as a viable solution for designing distributed training environments.