AEGIS' EXPANDING TRAINING:
ENABLING TEAM, FLEET, AND JOINT SIMULATIONS

Neil K. Lauver and Warren L. Jones
Computer Sciences Corporation
Dahlgren, VA

1. INTRODUCTION

AEGIS recently expanded its training for Team, Fleet and Joint simulations. This was accomplished through an open architecture reengineering of the training system, the use of the Distributed Interactive Simulation (DIS) protocol, and an alliance with Battle Force Tactical Training (BFTT). This paper explores each of these contributions (Section 2) and presents the corresponding benefits to training (Section 3). First, however, the paper provides a brief historical perspective (Paragraph 1.1).

1.1 Background

For 15 years, AEGIS has been the U.S. Navy's premier surface combatant - an integrated complex of sensors, weapons, command assistance, tactical dissemination, communications, and a hierarchical manning control structure. As a readiness and support component providing training at all levels, AEGIS has included an embedded Combat Training System (ACTS). This system is an assemblage of on board hardware and software - a central controlling suite interacting with the rest of the Combat System where sensor and weapons-related training support processing is embedded. The hardware of the central suite comprises multiple CPUs, multiple communications channels, and mass storage, hardcopy, and operator control peripherals. The software comprises four computer programs or Computer Software Configuration Items (CSCIs): a real-time, "executive" CSCI; a real-time, "tactical team coordination training" CSCI; a non-real-time, "individual operator training" CSCI; and a non-real-time, "post-mission/event-reconstruction/training-debrief" CSCI.

In 1992, the Navy began to discuss expanding AEGIS' training from its historical role of intraship preparedness into one supporting theater preparedness via coordinated intership/multi-platform simulations. These discussions resulted in the conclusion that, in order to achieve all the goals, the legacy Mil-Spec training system needed to be modernized through a comprehensive reengineering effort - one that would continue to support the full, multi-level training for each ship's Combat Information Center (CIC) while also supporting the new, distributed simulations.

The actual training expansion occurred between February 1994 and June 1997. This paper documents that work.

2. DISCUSSION

The expansion of AEGIS' training occurred in two stages. First, the training system was reengineered from a clean slate; only new technologies were employed; and the focus was on standards-compliant, growth-enabling open architecture commodities Commercially-Off-The-Shelf (COTS). Second, leveraging the reengineered system, AEGIS implemented the DIS communications protocol and integrated with BFTT thereby enabling participation in Fleet and Joint simulations. The balance of this section discusses these stages.

2.1 Open Architecture Reengineering

2.1.1 Rationale

There were several rationales for reengineering the training system. The legacy Mil-Spec interfaces were limited to Naval Tactical Data Systems (NTDS) standards and could not support modern LAN connectivity.

The legacy Mil-Spec computing resources were depleted, preventing system performance from being satisfied. Three of the 16 functions in the team training CSCI required 800, 500, and 100 ms, respectively, of the theoretical 1000 available. Further, due to the closed nature of the Mil-Spec suite, it was impossible to add processing on demand or to upgrade to more capable CPUs to overcome the depletion.

Complicating matters, the training requirements were in their fifteenth year of accretion, having gradually moved away from the original system engineering. For instance, the new throughput requirements for the multi-platform simulation traffic, when finally tallied, represented more than an order of magnitude increase over all previous training I/O combined; these requirements, and the accompanying (new) application-level requirements, simply could not be accommodated in the legacy Mil-Spec suite.

Finally, AEGIS' training did not support a high stress, synthetic theater wargaming environment. The track capacities needed to be increased to support the larger volumes inherent in theater operations. Thus, the full power of training with AEGIS was not being realized.

2.1.2 Strategy

The rationales bespoke the need for a more flexible growth path - hence the decision to follow an open architecture approach. An open architecture approach, as used here, is a particular manner of envisioning and constructing a system such that the configuration of hardware and software is flexible, can be used freely, and readily accommodates changes with minimal impact. The approach employs modern, widely-supported, standards-based technologies, and focuses on portability at all levels - the hardware, the operating systems, the communications protocols, the languages, and the applications. As a corollary, the approach avoids solutions that are proprietary, constraining, restrictive, closed, or nearing the end of their life cycle.

To complete the strategy, the training system also needed the following:

a. The complete replacement of the legacy hardware and software, as none of it abided the open principles.

b. Up-to-date performance and throughput calculations and projections for proper budgeting and sizing.

c. Specific assessment criteria for selecting the products.

d. Standards compliance throughout to maximize interoperability and portability.

e. COTS commodities for cost efficiency and schedule availability.

f. To be conceptualized as a set of upgradeable and expandable building blocks.

2.1.3 Selection Process

At the beginning of the selection process, it was clear that the topology of the new system (i.e., the hardware footings) would influence the form of the other products that had to be selected as well as the space allocation requirements in the ships' computer rooms where real estate is at a premium. Therefore, extra care and attention was given at this juncture.

The topologies considered included chassis-enclosed shareable backplanes [FutureBus+, Versa-Module Eurocard (VME), and VME64]; chassis-enclosed shareable subbackplanes [VME Subsystem Bus (VSB), Raceway, SkyChannel, Autobahn, and others]; high-end workstations [from the Navy's then Tactical Advanced Computer (TAC) contract]; hybrid workstation/backplane interconnects; and an Open Systems Module (OSM) suite-extension provided by the developers of the legacy Mil-Spec system. A crucial issue was the need to support six carry-over Combat System communications channels while also supporting at least two new networks - one providing intership/interplatform time coherency and one connecting to the DIS WAN. The topology that best satisfied this issue, the other system requirements for processing and peripheral interfacing, the ships' computer room requirements, and the open architecture engineering strategy (with the most openness) proved to be a single, rack-mountable, ruggedized VME chassis. This decision was vendor-independent.

Having resolved the topology, the selection process for the remainder of the system proved relatively easy for three reasons: one, there was an abundance of COTS offerings; two, the vendors and their sales arms were usually quite accommodating; and three, our up-front investment that pre-established the assessment criteria enabled us to focus on the critical information, separating the wheat from the chaff.

Still, the process proved time consuming due to the array of items needed (an entire subcomplex was being replaced) and the large volume of information that addressed the assessment criteria. In fact, the entire selection process spanned four months. During that time, the assessment criteria were used to build lists containing the candidates within each required category. (The categories were chassis/backplanes, subbackplanes, CPUs/Single

Board Computers, interfaces and devices, global memories, operating systems, languages, and development environments.) It would be fair to characterize this portion of the selection process as exhaustive, as the VME International Trade Association (VITA) was consulted for all available products. In all, more than 150 were assessed. Ultimately, each list was pared down to a "short list" of finalists. From there, candidates were rescrutinized, and the desired vendors were contacted one last time in an effort to perfect our findings.

The recommended system configuration items were then documented and presented to the Navy which, after careful review (including review by AEGIS' Computing Resources Committee), adopted them in their entirety. The Navy then initiated the procurements.

2.1.4 Selections

The following selections formed the basis of the training system's open architecture. In all instances, the items were COTS commodities and standards compliant.

The open architecture VMEbus (ANSI/IEEE 1014 and ANSI/VITA 1-1994) and VSB bus (ANSI/IEEE 1096) solved the training system's computational bounding problems and the cost of the accreting requirements, providing I/O bandwidth to the limits of the networks. A 21-slot VME card cage was selected to account for all processors, global memory, interfaces, and devices - with room to grow. A 6-slot VSB backplane was selected as an extra throughput lane simply for growth and flexibility as new requirements emerged - the idea being to balance the total system throughput so that no one bus would be heavily loaded (eliminating up-front any possibility of a single-bus bottleneck) and so that bus bandwidth reserves could be maintained.

VME64-compliant Single Board Computers (SBCs) provided an increase in processing potential (by a factor of 5.1), local bus bandwidth potential (by a factor of 5), memory capacity (by a factor of 32), and memory speed (by a factor of 7.5). They also provided standardized floating point support (IEEE 754) and standardized interfaces for ethernet (IEEE 802.3), SCSI (ANSI X3.131), VSB, and RS-232. Single-slot, 6U form factors were selected.

Multiple high-speed networks and network interfaces provided modern LAN communications. A Fiber Distributed Data Interface (FDDI) (ANSI X3T9.5/X3.139/X3.148) was selected to accommodate direct connection to the BFTT-provided, shipboard DIS LAN - known as the Synthetic Theater of War (STOW) LAN. A 10Base2 ethernet network (IEEE 802.3) was selected to communicate with the training system's operator console. A separate 10Base2 ethernet network was selected to interchange tactical information between the training processors (SBCs) and to serve as a third throughput lane. An Internet Protocol (IP) compliant backplane network was selected as a TCP/IP boot LAN and as an alternative interprocessor UDP/IP LAN. A Global Positioning System (GPS)/InterRange Instrumentation Group (IRIG) time distribution interface was selected to provide time coherency locally (intraship) and remotely (intership) thereby ensuring interoperation with BFTT.

Modern peripherals replaced the training system's Mil-Spec support infrastructure. SCSI-compliant, read/write optical drives were selected, providing formatted storage capacities and performance in excess of the training requirements. And, an ethernet-compliant laser printer satisfied the shipboard hardcopy requirement.

A POSIX-compliant (IEEE 1003) real-time operating system replaced the non-portable AEGIS executive, providing the required determinism, low latency, interoperability, portability, scaleability, instrumentation, and symbolic debug capability.

Where carry-over interfaces were required (to interact with the rest of the Combat System still in the Mil-Spec equipment), NTDS-compliant VME boards were selected.

2.1.5 Architecturalization

With the configuration items selected and the training requirements in hand, the next step was to complete the architecturalization. This involved systematically arranging the hardware items, apportioning the training requirements to the processors, and performing some modest VME configuration.

The hardware items were partitioned into two prioritized clusters: an I/O cluster and an application cluster. The I/O cluster was arranged further based upon the criticality, frequency, and volume of the specific data on each interface as well as the knowledge acquired about the behavior of each interface during the selection phase. The application cluster did not require further arrangement. As a sanity check, the new hardware arrangement was juxtaposed with the legacy arrangement; where applicable, the relative priorities mapped one-to-one.

Apportioning the training requirements leveraged knowledge from the legacy system regarding CPU utilization. It also took advantage of the relative priorities of the requirements - that is, the I/O requirements were the most important, followed by team training, individual operator training, and, finally, post-mission/training-debrief.

The I/O requirements were apportioned to their own processor to unburden - from the constancy of interrupts, pollings, and buffer management - the processors that would be executing the training programs. (ACTS would be interrupting between 300-400 times per second.) It also was done to isolate and make independent this CSCI.

The team training requirements were distributed across two processors: one for the new, DIS Protocol Data Unit (PDU) filtering and entity management requirements and one for the traditional, Combat System track injection/extrapolation/engagement requirements. As the PDU requirements represented the beginning of the DIS I/O thread, they were added to the I/O control processor.

The individual operator training requirements were apportioned to a third processor, as they were known to consume 600 ms in the legacy suite and since the program could be run in parallel with team training. The post-mission requirements were apportioned to the same processor. This was done for three reasons. One, the post-mission program did not need to execute in parallel with the others; two, the processor in mind was the control point for the system and could leverage the operator ethernet interface; and three, it resulted in a savings to the Navy by not having to buy an extra configuration item for 58 platforms/sites. This engineering decision saved the Navy $470,000.

Configuring the VME included the typical bus ownership, bus tenureship, and interrupt settings. It also included two portable performance optimizations: 64-bit block transfers for the high-bandwidth FDDI interface and 32-bit burst transfers for the carry-over interfaces, where multi-word flows existed. Other optimizations were possible (such as write-postings, mailboxes, location monitors, early DTACKs, and pre-assigned IRQ lines). However, they were avoided for portability reasons.

Figure 1 illustrates the results of the architecturalization. The application cluster is on the left and the I/O cluster is on the right. Within each are the corresponding subdivisions and apportionments. For completeness, the interfaces are identified.

2.1.6 Software Development

The final step in the reengineering stage was the development of the new training software.

Drivers were developed in Ada for the six carry-over interfaces, the GPS/IRIG interface, and the operator control interface. These drivers employed standard UDP/IP sockets for interprocessor communications, providing a growth option for redistributing the training data outside the VME should that prove desirable in the future. The GPS/IRIG driver was designed to provide time accurate to within 100 µs. A pseudo-driver was developed for the optical devices using the SCSI control features of the operating system. This pseudo-driver also employed standard Unix sockets for interprocessor communications, and was implemented in Ada. The FDDI interface came with its own network driver.

The team training requirements were redesigned to overcome the effects of the patched-work accretion, and re-sourced in Ada to facilitate future porting. Significant performance gains and growth features resulted. For instance, the cost in the number of CPU instructions to process a typical AEGIS training track was reduced from 776 to 431. A fundamental change in a periodic processing algorithm reduced the previously-mentioned 800 ms function to approximately 135 ms. And, the format of AEGIS' training scenarios was changed from fixed to variable, enabling larger scripting commands and limitless scripting extensions.

The individual operator training development leveraged an earlier Ada version (this version did not execute in the legacy Mil-Spec suite). This version was extended to support six trainees simultaneously. It previously supported just one.

The post-mission development reused event reconstruction code from another group within AEGIS that has event reconstruction as its charter. This resulted in improvements to the tabular and graphic debrief reports used by the ship's crew, and prevented duplication of effort.

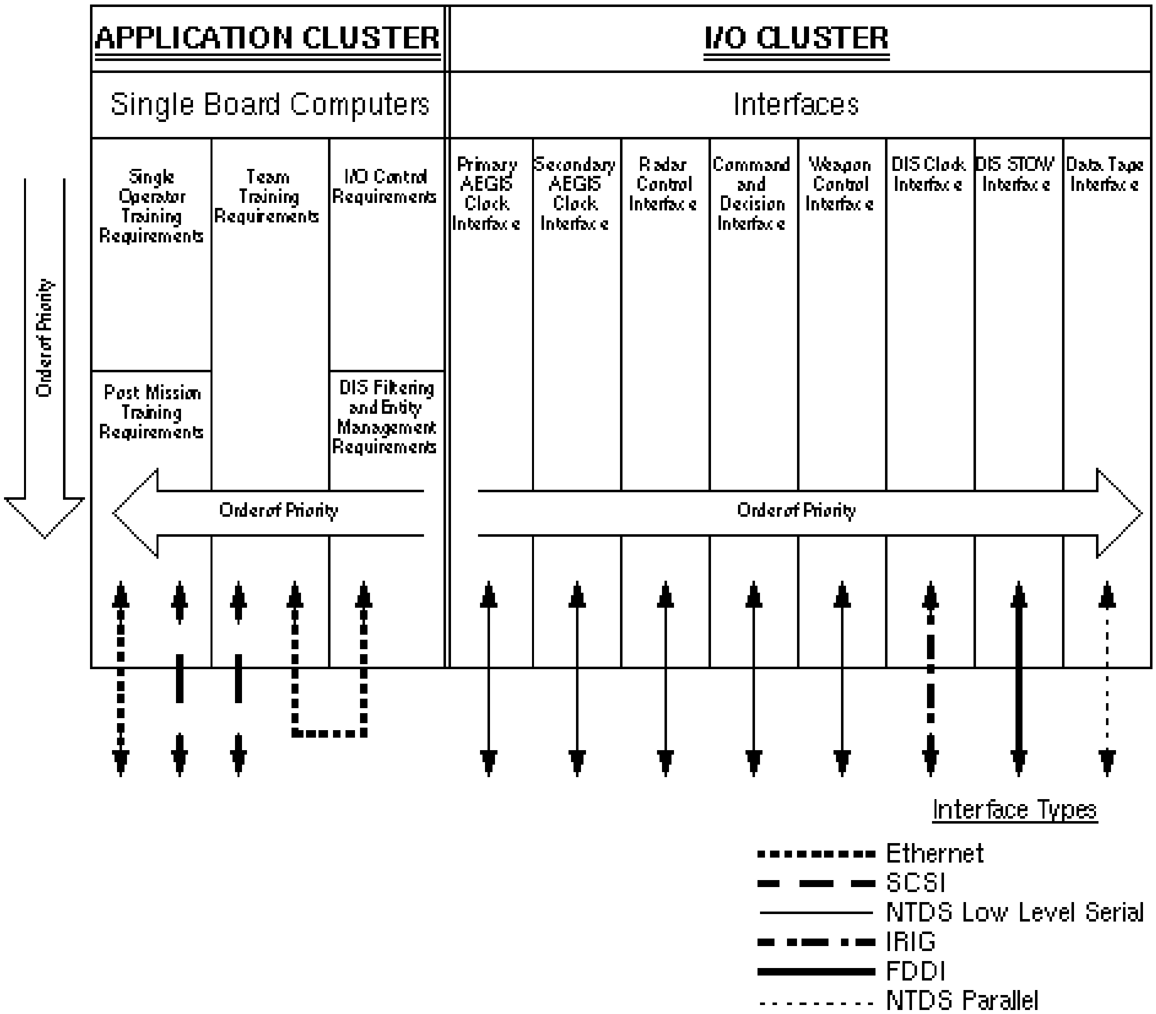

Figure 1. Training System Architecture

Throughout the development, the system was measured to ensure the required performance and throughput and to assess the original predictions and product assessments. In almost all cases, the measurements fell in line. Figure 2 illustrates this. The only notable problem was the cost of Ada's run-time checking. In certain time-critical areas, the checks prevented the fulfillment of performance requirements and were disabled. For instance, the management of AEGIS' internal 976 µs clock required 2.76 ms with checks and 80 µs without.

Thus, at the end of this stage, the new system was in place - with intraship training having been upgraded at each level.

2.2 DIS Implementation

To expand beyond intraship training, AEGIS implemented DIS - the lingua franca enabling the intership/multi-platform communications. (DIS is a set of IEEE standards that provides the specifications for linking together geographically dispersed simulations. The actual simulation information is encapsulated in PDUs inside standard IP packets. The packets are exchanged in real-time between conforming "simulation applications.")

AEGIS implemented DIS' "Communications Services" and "Applications Protocol" standards (IEEE 1278.2 and 1278.1-1995, respectively). This implementation was influenced by BFTT. The reason - BFTT already had developed a shipboard

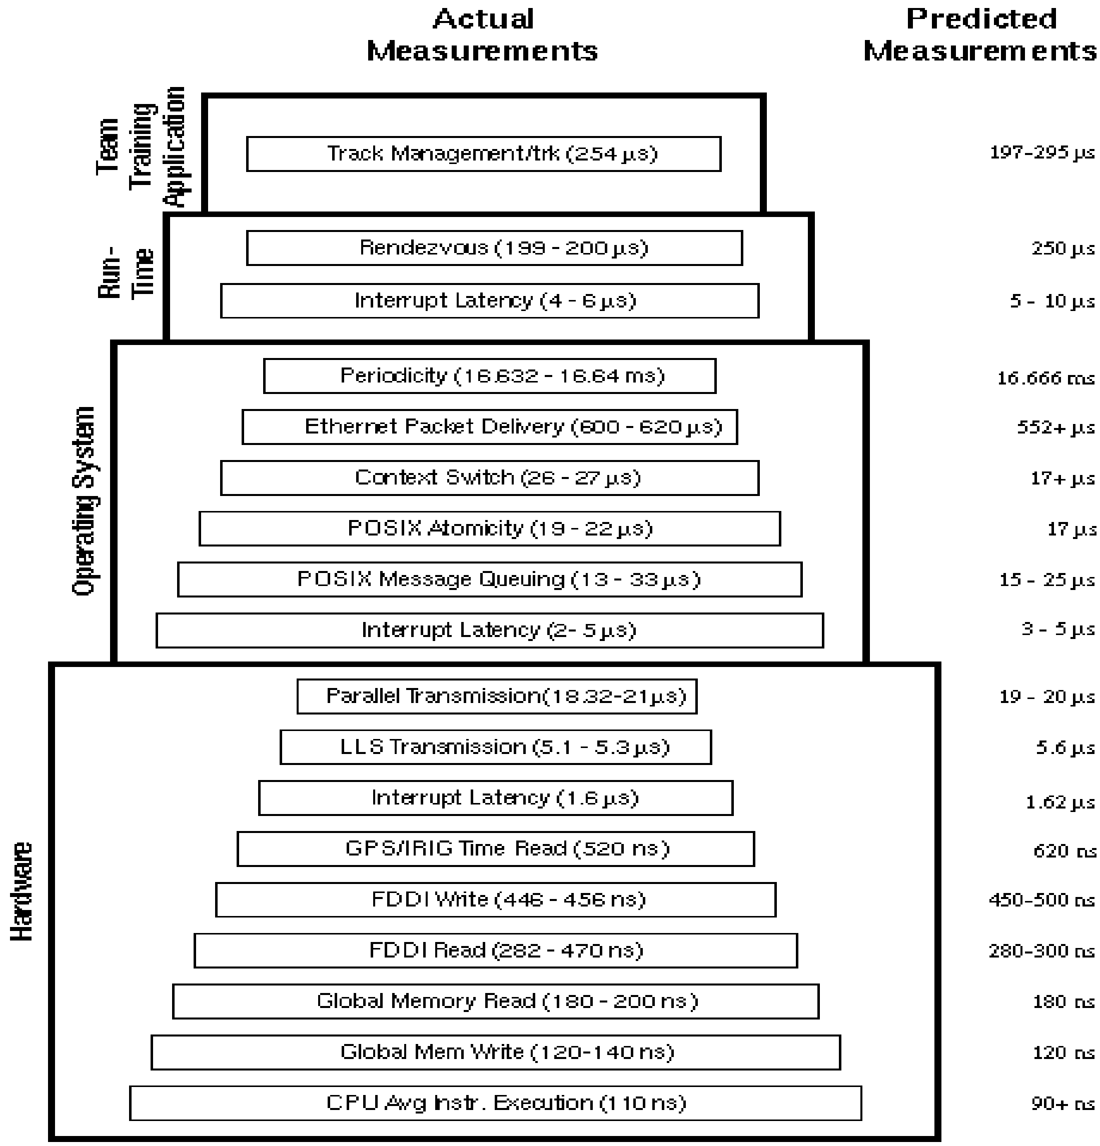

Figure 2. Training System Timing

DIS WAN gateway and a shipboard DIS STOW LAN which, together, were capable of linking together the "On-Board Trainers" (OBTs) for three other Navy platforms (carriers, frigates, and submarines). Consequently, the gateway and LAN were points of leverage for AEGIS.

At the communications services layer, the gateway and LAN were using network broadcasting and the UDP/IP protocol. AEGIS complied with both of these. On the one hand, compliance was straightforward: when installing the FDDI driver, AEGIS simply configured the broadcast and UDP/IP options. On the other hand, broadcasting meant that each PDU transmitted by each DIS participant would be entering the training system resulting in unwanted inputs. Given the up-front reengineering, including the FDDI performance optimizations, the volume was not a problem. However, when quantifying the amount of data that actually could be used, the number came in under 50%. Consequently, to minimize the impact of the extraneous data internally, two techniques were applied: PDU filtering and PDU repackaging.

The PDU filtering kept unusable PDU kinds, entity types, entity domains, colliding entity identifiers, and entity ranges from imposing themselves on the team training program. They simply were not forwarded. [This filtering is functionally similar to the Run-Time Infrastructure (RTI) of the High Level Architecture (HLA) - the Department of Defense (DoD) mandated successor to the DIS protocol.]

The PDU repackaging involved reporting only the usable portions of the PDUs. For instance, since the Combat System could only utilize 15 of the 36 words comprising the Entity State PDUs, that subset was re-grouped and forwarded to team training. This resulted in a 58% reduction in bandwidth per entity. Small PDUs were forwarded as received, since in almost all instances, team training could utilize the entire PDU.

At the application protocol layer, AEGIS implemented the DIS requirements for (a) initiating an exercise, (b) acting upon the PDUs pertinent to the Combat System, (c) modelling the position and movement of the input entities in the Combat System's reference, (d) modelling the output entities (i.e., Combat System-generated tracks) for positioning and movement by other simulation applications on the network, (e) managing time, and (f) terminating an exercise.

In initiating a DIS exercise, AEGIS implemented the Start/Resume, Acknowledge, Action Request, Action Response, and Create Entity PDUs. The real-world time field of the Start/Resume PDU contained the GPS time to start the exercise.

In acting upon the pertinent PDUs, AEGIS implemented entity appearance and movement, weapons firing, ordnance detonation, collision detection and assessment, entity interrogation, electronic warfare, and acoustic emissions. Table 1 summarizes the applicable PDUs along with their protocol family.

In modelling the input entities, AEGIS implemented mathematical conversions to address the differences between AEGIS' geodetic orientation and DIS' geocentric orientation. Dead reckoning algorithms were implemented to extrapolate the position of the entities, as parameterized in the Entity State PDUs. And, entity mapping functions were developed to address the differences between DIS' enumerations and AEGIS' tracking types.

In modelling the output entities, AEGIS

DIS Protocol Family	Protocol Data Unit Name
Entity Information/Interaction	Entity State IFF/ATC/NAVAIDS Collision
Warfare	Fire Detonation
Distributed Emission Regeneration	Electromagnetic Emission
Simulation Management	Start/Resume Stop/Freeze Acknowledge Action Request Action Response Set Data Data Comment Create Entity Remove Entity Surface Ship System Status

Table 1. Applicable PDUs

implemented the inverse of the input entity modelling. The geodetic position was transformed to geocentric, AEGIS' dead reckoning was parameterized, and AEGIS' tracking types were converted to DIS' enumerations.

In managing time, AEGIS accessed the GPS's time-pulse downlink through the shipboard GPS/IRIG LAN. For outbound PDUs, this time was mathematically converted into the encoded DIS format and made absolute via DIS' absolute timestamping option. For inbound PDUs, the encoded (absolute) timestamp was converted into AEGIS' time format and then mathematically "aligned" with AEGIS' clock. This was necessary to determine "when" in AEGIS an entity's data had been generated at the "sending simulation" so that AEGIS could fully extrapolate the entity's "distance."

In terminating an exercise, AEGIS implemented the Stop/Freeze and Acknowledge PDUs. As was true for initiating an exercise, the real-world time field of the Stop/Freeze PDU contained the GPS time to stop the exercise.

Lastly, several other DIS adjustments were made to accommodate AEGIS' sensors and OBTs. The experimental Surface Ship System Status (SSSS) PDU was implemented. In doing so, it was refined to include a "Permission To Train" field reflecting the state of AEGIS' Combat System (tactical or in training). And, DIS' Comment PDU was tailored to enwrap carry-over training messages between team training and the Anti-Subsurface Warfare (ASW) OBT, minimizing the impact to both sides. (As a leveraging of the STOW LAN, the legacy interface between team training and the ASW OBT - NTDS Parallel - had been replaced with FDDI. However, the ASW OBT itself had not migrated fully to DIS.)

Thus, at this point in the second stage, AEGIS possessed the ability to communicate in intership/multi-platform simulations, but needed theater coordination.

2.3 BFTT Integration

BFTT, as one of its missions, coordinates virtual theater training between geographically dispersed Navy platforms situated in their home ports. They accomplish this by providing for each platform the DIS WAN gateway, encryption equipment, STOW LAN, and synthetic theater console. The theater console is a graphical user interface providing control over a 4,000 nm by 4,000 nm by 300,000 feet gaming volume anywhere in the world.

Given what each project had to contribute, it was natural for AEGIS and BFTT to integrate. Recognizing this, the projects adopted a two-step integration plan - joined in the middle by DIS.

Step 1 comprised a basic communications and simulation experiment between the two systems. This took place on board the AEGIS cruiser U.S.S. Hue City in Florida, at AEGIS' testing laboratory at the Naval Surface Warfare Center Dahlgren Division (NSWCDD) in Virginia, and at BFTT's development laboratory at the NSWC Port Hueneme Division (NSWCPHD) in Dam Neck, Virginia.

BFTT contributed its equipment and used commercial land lines to link the three platforms. (The labs acted as virtual ships.) AEGIS contributed its Combat System. Since this step predated the completion of AEGIS' reengineering, AEGIS also contributed a temporary bridge between the STOW LAN and the Combat System. With this setup, DIS traffic was transmitted from BFTT's theater console, to the transmitting WAN gateway, over the land lines, to each platform's receiving WAN gateway, to the AEGIS bridge, into AEGIS' Combat System, onto AEGIS' Large Screen Displays (LSDs) for the ship's crew, and vice versa - for each platform. The experiment proved that multiple, distributed AEGIS training systems could be interconnected into an interactive virtual theater. As a proof-of-principle, though, only modest simulation traffic was exchanged.

Step two comprised fully integrating the two systems for BFTT's Initial Operational Capability (IOC) and increasing the simulation traffic to realistic levels. To these ends, BFTT contributed upgraded equipment and software capable of supporting 2000 DIS entities. AEGIS contributed its reengineered training system (including like support for 2000 entities) and its DIS implementation, supplanting the temporary bridge.

For AEGIS, connecting to BFTT's FDDI STOW and GPS/IRIG LANs proved straightforward, due to the up-front work conducted during the reengineering stage. In fact, it was as simple as running the cables and making the connections.

The integration of the DIS PDUs also went well, due to frequent engineering meetings with BFTT and a living Interface Design Document (IDD). The only notable problem was schedule incompatibility.

(AEGIS' schedule had an earlier deadline by six months.) To overcome this, AEGIS conducted unilateral tests using BFTT engineering releases. The testing was accomplished by examining BFTT's raw FDDI/DIS traffic and determining the issues requiring attention. Corrections then were made while AEGIS' schedule was still forgiving.

Thus, at the end of the second stage, AEGIS and BFTT had integrated with realistic, high stress DIS traffic flowing between the two. In doing so, AEGIS achieved its goal of participating in coordinated Fleet and Joint simulations thereby directly supporting theater preparedness.

Figure 3 illustrates the transformed training subcomplex.

3. BENEFITS

3.1 Training

The benefits to training were numerous and included the following:

a. AEGIS' training repositioned itself for the future and the succession of new requirements cycles.

b. Intraship training was upgraded for the Ticonderoga and Arleigh Burke classes, improving the relevance and responsiveness for the trainees.

c. New interactive, Fleet and Joint training was added, enabling AEGIS' participation in virtual theaters.

d. Three theater exercises already have occurred; other exercises are planned.

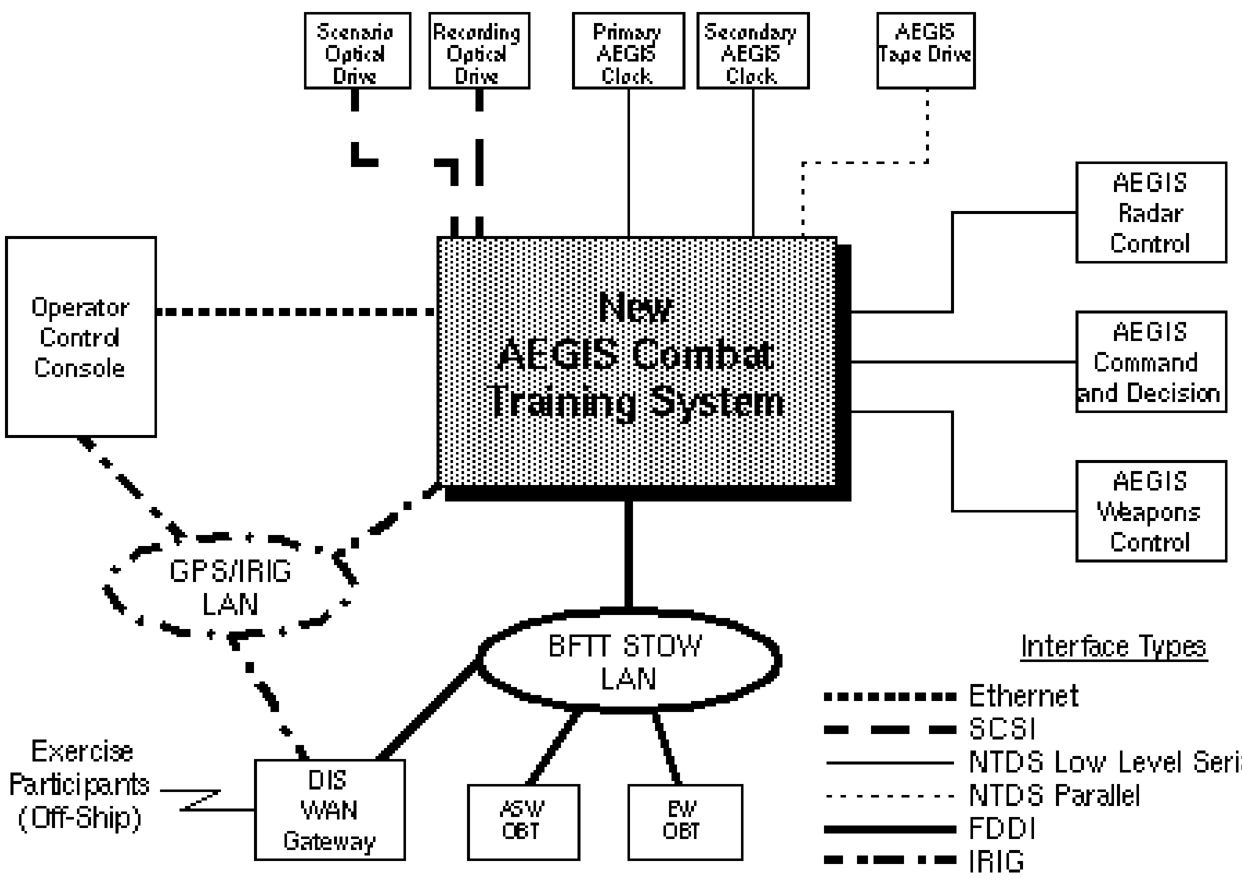

Figure 3. New Training Subcomplex

e. High stress entity environments were incorporated, making the simulations more realistic for the trainees.

f. Theater rehearsals and wargamings and simulated deployments and stagings are viable, enabling preparations prior to actual site penetration.

g. And finally, ship crews are able to *train where they fight* - wherever that electronic battlefield may be.

3.2 Cost Leveraging

The training expansion leveraged the following items:

a. Approximately 60 man-years of experience with AEGIS training and the legacy hardware and software.

b. BFTT's DIS WAN gateway and DIS STOW LAN.

c. Lessons learned from the early AEGIS/BFTT communications experiment such as realistic WAN and LAN throughput measurements.

d. Lessons learned from previous DIS exercises such as the nature and volume of the simulation traffic.

e. Certain BFTT source code, including the complex entity management processing.

f. The COTS commodities, resulting in the new training system costing a mere 15% of the legacy system.

4. CONCLUSION

Through its comprehensive modernization effort, AEGIS upgraded intraship training for single operators and the tactical team, and incorporated new intership/multi-platform training enabling Fleet and Joint simulations.

AEGIS joined the DIS infrastructure and forged a partnership with BFTT, allowing BFTT to move forward with its installation plan for 53 AEGIS ships through the year 2003.

AEGIS achieved its open architecture engineering goals and kept costs down through the use of COTS commodities.

And finally, in accomplishing all the above, AEGIS addressed the needs of the sailors who are called upon to man the stations.

Looking ahead, the DIS protocol is to be replaced with the HLA. Since AEGIS and BFTT are joined in the middle by the protocol, both will be assessing the full impacts. This much is known today: since BFTT is the WAN gatekeeper for AEGIS, AEGIS' link to HLA will be, at a minimum, "through" BFTT. A paper describing three options for satisfying HLA has been prepared by BFTT and is being studied by AEGIS. In the near future, AEGIS will be working with BFTT to engineer the best solution.

Acknowledgements

The authors would like to acknowledge the support of the AEGIS Training Center (ATC), the Navy management team at NSWCDD, the BFTT engineering team at NSWCPHD, and the rest of the training development team at CSC and NSWCDD. The tireless efforts of these teams helped make AEGIS' training all that it is.