Proceedings of the 1996 Winter Simulation Conference
ed. J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain

Distributed Interactive Simulation: It's Past, Present, and Future

Deb Fullford

MäK Technologies
185 Alewife Brook Parkway
Cambridge, MA 01238, U.S.A.

ABSTRACT

In the mid 1980's the U.S. Department of Defense, Defense Advanced Research Projects Agency (DARPA) launched the SIMNET project, which gave birth to a technology era and networking protocol presently known as Distributed Interactive Simulation (DIS). The DIS Protocol is a widely used IEEE Standard. At first it was primarily used by the military for mission rehearsal, training, and weapon evaluations. Today, other areas, such as transportation, medical care, entertainment, Internet commerce, and manufacturing, are using DIS complaint simulations to meet their needs. This paper will provide a detailed overview of DIS. It will focus on the fundamental concepts, such as peer to peer connectivity, protocol data units, and dead reckoning. It will also cover some advanced topics, such as smoothing for visualization and shared virtual environments. Finally, this paper will discuss the future of DIS. In particular, MäK will discuss the next generation DIS, otherwise known as the High Level Architecture (HLA), new protocols, and bandwidth reduction techniques.

1 INTRODUCTION

Over the past 30 years, the military has developed computer-generated synthetic environments to train people in very complex, dangerous, or financially risky tasks. However, since the technology for developing synthetic environments was very expensive, simulators were only developed for systems in which the cost of failed performance outweighed the cost of the training simulator.

The solution of this problem led to the invention of a technology called Distributed Interactive Simulation (DIS), a networking technique for connecting thousands of simulators together on a computer network. Instead of spending $50 Million for a single flight simulator, the demand was created for thousands of simulators to be networked in a synthetic environment for training groups of combatants in collaborative tasks.

The following sections will discuss the history of this cutting edge technology, provide high level technical discussions on how it works and provide some insight on where it is going.

2 HISTORY OF SYNTHETIC ENVIRONMENTS IN THE MILITARY

Until about 1983, the common perception of a training simulator in the Department of Defense was that of a stand-alone, high-fidelity device designed to train a single individual, or single crew, in the skills necessary to operate a weapon system; skills such as flying, shooting, targeting, etc. The goals of the engineers who designed and built these devices was to provide the highest possible degree of realism to the crew, accurately emulating the real-world cues that the crew would experience in an actual combat situation. These lofty goals tended to drive the performance requirements and prices of these systems through the roof.

Depending on performance capabilities, these simulators could cost anywhere from $1 Million to $50 Million each. Very often, the simulator would be more expensive than the vehicle it simulated. The cost per man-hour of training time was staggering because each simulator accommodated only a single crew (often only one person), while the training facility required several operators to run the simulation. If a $50 Million simulator were operated for 8 hours a day, 200 days per year, for 10 years, it would cost $3,125.00 per hour to amortize its purchase price, without considering hourly operating costs or repairs.

Moreover, while these high-fidelity simulations could be used for individual crew training, they could not address the need to train teams of soldiers in collective or collaborative tasks such as practice of tactics, strategies and procedures. To train these groups, it was necessary to transport all participating soldiers and equipment to an exercise site to rehearse battle scenarios. These live training exercises, often held in unpleasant and hard-to-reach locations such as the North Pole or the Sudan Desert, usually concluded with a small percentage of casualties, and even some fatalities.

Scheduling such a live exercise can be as daunting a logistical nightmare as scheduling an actual war, and may take several months to complete. This limits the number of collaborative training exercises to one or two per year.

Beyond their cost and logistics problems, live exercises are a security nightmare. To provide a realistic training environment, the exercises must be carried out in an area which closely resembles the area of conflict. These areas may be close to enemy land where an exercise may be easily observed, or on U.S. soil where the exercise may be observed via foreign satellites.

With the live training exercises' associated loss of life, high costs, poor security, and infrequent staging, the need for a better collaborative training solution was obvious. In 1983, Air Force Lieutenant Colonel Jack Thorpe, of the Defense Advanced Research Projects Agency (DARPA), decided to solve this problem.

The specific mission of DARPA was to fund high risk/high payoff advanced technologies which could be used by our armed forces to maintain an advantage on the battlefield. Almost all exotic technologies in the US arsenal, including Stealth radar avoidance, cruise missiles, laser guided weapons, Virtual Reality, and even the Internet, have all come from DARPA projects.

Jack Thorpe, along with his consultant, Colonel (retired) Gary Bloedorn, initiated one of the most influential technology revolutions ever to have occurred within the DoD. Jack and Gary envisioned a multi-user computer-generated synthetic environment where collaborative training, weapon concept development and tactics development could be performed in a very controlled, cost effective, safe, and secure manner. The outcome was SIMNET (SIMulator NETworking), a major paradigm shift for the training community. Instead of developing a small number of high-fidelity, multimillion dollar simulators for part-task training, DARPA developed inexpensive $250,000 simulators which had less fidelity than their older cousins.

"Selective fidelity" was a buzzword coined in the SIMNET program, meaning that a simulator need only be good enough to accomplish the training goal at hand. Hundreds of these "selective fidelity" simulators were networked together to provide an electronic world for collaborative tactics development and rehearsal. The tank simulators which SIMNET fielded had relatively low resolution (320 X 128 pixels per vision block), a relatively low frame rate (15 Hz), no motion base, an inexpensive fiberglass hull, and a reduced set of controls. For an 8 hour a day, 200 day a year, 10 year service life, the amortized cost of a SIMNET simulator costs only $15.63 per hour.

Proponents of high-fidelity simulators and live training exercises vilified the SIMNET system, claiming that it's cartoonish graphics and reduced fidelity man-machine interface rendered it useless as a training device. However, for its targeted intention, namely collaborative training, the SIMNET system ultimately proved to be an excellent value. Over its 10 year life span (now entering its twilight) SIMNET enjoyed widespread success and acclaim for its reduction of training costs, increased quality and quantity of tactical team training, and usefulness as a testbed for new weapon concepts.

SIMNET's biggest accomplishment, however, was the spawning of a standard networking protocol which enables heterogeneous simulations to interact in a shared synthetic environment. This standard, called the DIS Protocol, has forever changed the way the Pentagon trains our soldiers. DIS compliance is now mandatory for all new training devices procured by the DoD.

There is still decades of research and development to be done in the areas of networked simulation for collaborative training. The most notable present-day need is in the area of immersive virtual reality technology, used in the simulation of individual soldiers. Even with the mighty budget of the DoD, we still cannot field a helmet-mounted display with the field-of-view, resolution, and frame rate necessary to render a useable world. Current offerings leave the user legally blind, and without necessary visual cues to perform required tasks. Force and tactile feedback are also lagging far behind. State-of-the-art systems currently use exercise bicycles, stair climbers, or linear treadmills to simulate human locomotion.

It is clear, however, that the infrastructure and architecture for a global network of synthetic environment portals for training, entertainment and education is in place. This infrastructure is in constant use and is driving development of associated technologies.

3 OVERVIEW OF DIS

3.1 Technical Details

The DIS Protocol consists of some 29 different network packets, known as Protocol Data Units (PDUs) which are passed among simulation nodes. The structure of the PDUs are outlined in a DIS standards document. The most current document being the "IEEE Standard for Distributed Interactive Simulation - Application Protocols" which outlines the IEEE 1278.1 DIS Standard PDUs (IEEE 1995).

The application specific information is encapsulated in UDP/IP Ethernet frames and transported in broadcast mode (one to all). The packets can be sent over any network media, from voice grade phone lines to ATM switches. More bandwidth allows a greater number of entities to be supported. DIS has also been used over the Internet using IP multicasting with excellent results.

A DIS PDU can describe the state information of a dynamic entity (Entity State), a combat event (Fire and Detonate), a resupply interaction (Logistic), or an electromagnetic emission such as light, radar, or energy weapon emissions (Electromagnetic Emission). Because the heritage of the protocol is military in nature, versions up to DIS 1278.1 have retained many military specific fields in the packets. However, current proposals for the next generation protocols remove most of the military content and provide a much more efficient, reconfigurable, and general purpose packet structure.

DIS uses predictive algorithms, known as dead reckoning, that allow entities on the network to greatly reduce the frequency of rebroadcasts of state. Each entity broadcasts its type, location, velocity, acceleration, orientation, and angular velocity. All the receiving simulators can then propagate the sending entity into the future, relieving the sending entity of the responsibility to continually rebroadcast. When the error between the exact position of the entity and the predicted position exceeds a certain threshold, the sending entity will update the network with its new kinematic state.

Each entity sends out information on how it should be dead reckoned. The Dead Reckoning Models (DRM) notation consists of three elements:

DRM (F or R, P or V, W or B)

The first element specifies whether the coordinate system of the entity is fixed (F) or rotating (R). The second element specifies whether rate is constant (P) or varying (V). The third elements specifies the coordinate system to be used, either world (W) or body (B). (IEEE 1995). A common dead reckoning algorithm used is DRM (F,P,W), shown below in Equation 1.

$$P = Vo + V\mathbf{d}t \quad (1)$$

Networked latency and dropped packets can cause a jittering effect when visualizing the location of the entities. This occurs when there is a discrepancy between the values most recently received from a PDU and the calculated dead reckoned values. To counteract this, several visualization tools use translational and rotational smoothing to interpolate between the dead reckoned location and the location received from the PDU. As a cautionary note, high level smoothing (such as rotational) utilizes precious CPU resources.

DIS enables disparate simulators to interact in a shared virtual world. In order to ensure interoperability between these simulations a common coordinate system is used to transmit the entities' state within the shared synthetic environment. Each simulator needs to be able to transform between their local system and the DIS coordinate system, knows as the Geocentric Cartesian Coordinate System (GCS). It is a right-handed system with the origin at the centroid of the earth. The positive X-axis passes through the Prime Meridian at the equator, the positive Y-axis passes through 90 degrees East Longitude at the equator and the positive Z-axis passes through the North Pole (IEEE 1995).

The DIS Entity State packet, which makes up most of the network traffic in DIS 2.0.4, is about 140 bytes long, and is broadcast anywhere from once every 30 seconds, to four or five times per second. The next generation will further increase efficiency, using a trimmed-down Entity State packet that reduces network traffic by a factor of 3 or better (see Figure 1) (Taylor 1996b).

3.2 Benefits

The DIS architecture provides flexible tradeoffs between computational loading, positional error, and network bandwidth. If highly accurate position is required, the error threshold can be reduced, which will result in more network broadcasts of state. Conversely, if the only network bandwidth available is that provided by a 9600 baud modem, and 25 entities are required on the system, the error thresholds can be increased, and more compute-intensive prediction algorithms can be used on the receiving ends. State-of-the-art DIS systems currently in use by the government are running over 10,000 entities on Ethernets, long haul networked from several locations. Future developments, such as low-cost commercial data services (bi-directional cable TV, ISDN, etc.), next generation protocols, and support for multicast routing, will enable DIS exercises to reach the 100,000 player level.

Another important feature of DIS is its self-healing nature. When a new entity enters the world, he begins to broadcast Entity State packets. If recipients have never heard from this entity before, they simply add him to their remote entity database. If an entity is not heard from within five seconds, recipients will time him out, removing him from their remote entity database. Players can enter and leave at will without disturbing other participants, causing no effect other than their appearance and disappearance. Dropped packets don't cause system failure. In this architecture there is no central server, thus no single point of failure.

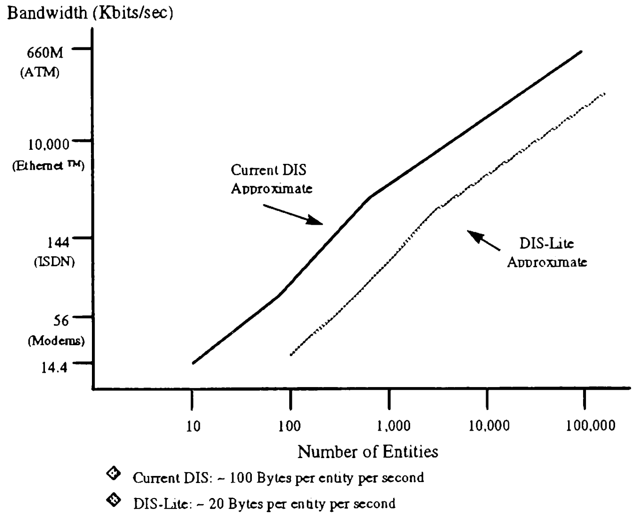

Figure 1: Bandwidth Requirements for Current and Future DIS Versions.

Heterogeneous nodes can interact with each other using DIS. The network protocol provides a standard mechanism for communication between simulators which may have radically different architectures. An entity broadcasting an Entity State packet is simply informing the network of what kind of entity it is, and its kinematic information. Two different recipients of the same Entity State packet may render the remote entity with very different levels of fidelity. For example, a PC may render a remote F-15 with only 10 flat shaded polygons, and may only have an internal simulation frame rate of 5 frames per second. A high-end Silicon Graphics workstation may render the same F-15 with 500 Gouroud shaded, phototextured polygons at a rate of 60 frames per second. In the entertainment world, this would allow a high-fidelity arcade-based system costing tens of thousands of dollars to interact, via the Internet, with a $200 home-based game system.

As compared to the computational power necessary to render out-the-window images in real-time, the DIS prediction algorithms do not represent a significant load. The most expensive prediction algorithm that DIS commonly uses consumes about 100 floating point operations per remote entity per simulation frame. Considering that the average PC can do about 10 million of these operations per second, and a simulation might run at 20 frames per second, a PC can handle several thousand remote entities using the most expensive algorithm. Naturally, this load must be balanced with other computational tasks, but this is a good indication as to the level of hardware necessary to accommodate fairly large DIS entity populations in a virtual world.

Each year, a demonstration of DIS interoperability is conducted at the annual Interservice/Industry Training Systems and Education Conference (I/ITSEC). Booths representing over 50 different organizations are networked together for a simulated exercise. DIS node systems, ranging from PCs to high fidelity flight simulators with multi-million dollar image generators, are connected for a week of interactive exercises. This is the third straight year such a demo was shown.

Though there are several commercial-off-the-shelf DIS developer's toolkits available, most of the industry uses a product called VR-Link™ produced by MäK Technologies of Cambridge, MA (www.mak.com). VR-Link, developed by some of the original SIMNET engineers, currently dominates the networked simulation market worldwide.

4 CURRENT NON-MILITARY USES OF DIS

It is no surprise that DIS technology has been finding its way into other government and commercial industries. For example, Air Traffic Control Simulations and Intelligent Vehicle Highway System research are among the government simulation communities using DIS.

In the commercial realm, DIS has been used for virtual reality systems, networking of games and entertainment systems, and most recently for multi-user virtual worlds on the World Wide Web. This section discusses some of the emerging non-military applications.

4.1 Government Applications

4.1.1 FAA Air Traffic Control Trainers

A DIS-based Air Traffic Control Training System was developed by the MITRE Corporation of Virginia. The system enables students to interact with simulated aircraft, either manned or automated, which could be coming from anywhere around the country. The aircraft are represented as DIS entities. Different situations can be scripted, repeated, and analyzed for performance. A national dial-in network for pilots in training is certainly possible, and has been proposed to the FAA.

4.1.2 Intelligent Transport Systems

The Department of Transportation has been sponsoring work in Intelligent Transportation Systems, formerly known as Intelligent Vehicle Highway Systems. The Analytical Sciences Corporation (TASC) of Massachusetts has developed a DIS-based automotive simulation to be used during their research. One experiment examined the effects of transportation on pollution. Like the FAA's pilot training network, a national network for driver training and city planning is also on the drawing board.

4.2 Commercial Applications

4.2.1 Networked Computer Games

Several game companies have adopted DIS as their networking mechanism for multi-user games on the Internet. Among them are Electronic Arts, Zombie, and Serum, three of the leading game developers for the new generation of 3-D first person interactive games. Several titles will be in the stores by Christmas of 1996.

4.2.2 Interactive Television

Telephone and cable TV companies are gearing up for high-bandwidth, bi-directional data services for the home, a move fueled by the passage of the new Telecommunications bill. The future promises such services as movies-on-demand, interactive video games, home shopping and more. By using DIS, the service provider can bring interactive virtual reality to users in their homes. The DIS bandwidth reduction techniques which allow thousands of entities to exist on Local Area Networks (LANs) will allow people in their living rooms to interact with each other in multi-participant virtual worlds. Each channel can be associated with a specific activity or experience. User's will be able to select from myriad sporting activities, games of chance and skill, and even fantasy adventures.

5 FUTURE DEVELOPMENTS IN DISTRIBUTED INTERACTIVE SIMULATION

DIS, being a relatively new technology, can still be improved. This section will discuss several research projects currently underway which will make DIS the logical choice for future Internet applications. The projects focus on the dynamic transfer of high-order data (VR-Protocol), accurate representation of contact forces (Newtonian Protocol), interoperability (High Level Architecture), and bandwidth reduction (DIS-Lite and Multicast).

5.1 VR-Protocol

In DIS, it is hard to add new behaviors or visual representations dynamically. This prohibits entities with experimental behaviors or visual characteristics from jumping into an existing exercise. A research project examining this problem has developed a VR-Protocol. The VR-Protocol enables the dynamic transfer of both behavioral and visual characteristics between systems (Taylor 1996a).

The VR-Protocol research project has demonstrated the ability to transfer a polygonal file for a new entity (used for visual representations such as avatars) on a query-response basis when needed. New players can now join the game and inform the pre-existing systems of their appearance. Though this capability is similar to Virtual Reality Markup Language (VRML), the transaction is not initiated by a person but is automatic and transparent to the users of the system.

The VR-Protocol also can transfer the behavior of an entity, via scripts. The scripts define the new behaviors based on a predefined set of basic primitive behaviors. For example, if a hundred virtual people are being simulated on a single workstation, a very efficient algorithm for wandering around can be defined and sent to the other players. The complex behavior can then be used by each player to simulate the entities without using any bandwidth. VR-Protocol is incorporating several scripting languages, such as JAVA. The VR-Protocol is really the best of VRML and Java rolled into one.

5.2 Newtonian Protocol

A drawback of DIS is its inability to transmit contact force information between two objects. This leads to low-fidelity emulations of collisions and logistical operations. The Newtonian Protocol is an extension to DIS which allows objects to exchange Newtonian forces. Using the protocol, systems can emulate tactile and force feedback, towing, and combat operations. The Newtonian Protocol uses a novel transfer function exchange paradigm to reduce the amount of network traffic required to maintain physical contact between two entities across a network (Taylor and Katz 1995)

5.3 DIS-Lite

Even with vast technological advancements in the communications arena, we are still far away from infinite bandwidth. In order to reduce the network bandwidth, a next generation DIS protocol, DIS-Lite, is being developed. DIS-Lite reduces bandwidth requirements by eliminating the transmission of redundant information. Figure 1 on page 4 compares the current version of DIS and DIS-Lite with respect to the required bandwidth versus the number of participants. DIS-Lite is approximately 30-70% more efficient than the current generation DIS (Taylor 1996b).

5.4 Multicast Traffic Management

Currently, DIS uses UDP Broadcast messages to send information across the network. This means that each participant receives every message that is sent whether or not they need or want the information. This unnecessary traffic is a waste of bandwidth. To alleviate this problem, a Multicast Traffic Management system is being implemented for the next generation of DIS. This capability will enable large scale DIS exercises to run over the Internet. The network traffic is separated into groups of interest, possibly by exercise ID, location, or even fidelity level. On a packet-switched internetwork, traffic will be routed only to those listeners who subscribe to a particular group of interest. This enables players to selectively send packets to subsets of listeners.

5.5 High Level Architecture

Although the existing DIS protocols work extremely well for certain applications, the DoD needs increased interoperability among simulations for which the DIS protocols may not be well suited, for example, higher order war game models. A fairly new research effort, funded by the Defense Modeling and Simulation Office (DMSO), called High Level Architecture (HLA) is concentrating on creating an interoperable architecture where the different types of systems can interact.

Under HLA all participating simulations are known as federates. A collection of simulations is known as a federation. An Object Model Template (OMT) is used to define simulation object models (SOMs) and federate object models (FOMs). The FOM is used by the participating federates to specify how information is sent. The OMT provides a standard framework for describing a simulation in terms of its objects and the interactions between objects (DoD 1996b).

HLA also encompasses a Runtime Infrastructure (RTI) which binds the federates together. Federates send information to the RTI which in turn sends the information to the appropriate parties. The RTI is comprised of five service groups:

- Federation Management
- Declaration Management
- Object Management
- Ownership Management
- Time Management

Federation management handles the creation, dynamic control, modification, and deletion of a federation execution. Declaration management enables federates to declare their desire to generate (publish) and receive (subscribe/reflect) object state and interaction information. Object management is responsible for the creation and deletion of objects as well as the transmittal of object and interaction data. Ownership management allows federates to transfer ownership of object attributes. Time management provides useful services for setting and modifying simulation clocks. (DoD 1996a)

Even though there has been a lot of hype about HLA, it is still a long way off. The release of a fully functional RTI has been pushed back to May of 1997.

6 CONCLUSION

DIS lets users interact with one another, and their environment, over a network. DIS' predictive algorithms dramatically reduces the amount of network traffic. Its self-healing property enables users to come and go without affecting the other participants. In addition, heterogeneous nodes can interact in the same synthetic environment.

The future is bright for DIS. The VR-Protocol will give DIS the ability to transfer visual representations, as in VRML, and behavioral information, as in Java. The Newtonian Protocol provides a mechanism to represent sustained contact force between entities over a network. DIS-Lite and Multicast Traffic Management Systems are aimed at reducing the required bandwidth making large scale DIS operations across the Internet possible.

REFERENCES

Department of Defense. 1996a. *High Level Architecture for Simulations Version 0.5 Interface Specification*.

Department of Defense. 1996b. *High Level Architecture Object Model Template Version 0.3.*

Institute of Electrical and Electronics Engineers. 1995. *IEEE Standard for Distributed Interactive Simulation - Application Protocols*. IEEE Publishing. NY.

Taylor, Darrin. 1996a. The VR-Protocol. In *Proceedings of the 14th DIS Workshop on Standards for the Interoperability of Distributed Simulations*.

Taylor, Darrin. 1996b. DIS-Lite and Query Protocol: Message Structures. In *Proceedings of the 14th DIS Workshop on Standards for the Interoperability of Distributed Simulations*.

Taylor, Darrin and Katz, Warren. 1995. Aspects of the Newtonian Protocol and its Application to Distributed Simulation. In *Proceedings of the 10th DIS Workshop on Standards for the Interoperability of Distributed Simulations*.

AUTHOR BIOGRAPHIES

DEBORAH A. FULLFORD is a member of the technical staff at MäK Technologies. She has a B.S and M.S in Aeronautical Engineering. She is a member of the Society of Computer Simulation and has served as both Session and Panel Chair for the Distributed Simulation Track of the Summer Computer Simulation Conference and the Multi-Simulation Conference. She is also a member of the DIS Standards Protocol Working Group.