Enhancing DIS to Support the Million Player Scenario

Christina Bouwens
Hughes Training Incorporated, Link Operations
Binghamton, NY

Mark Juliano
AEL Industries, Cross Systems Division

Ron Matusof
BroadVision, Inc.
Los Altos, CA

ABSTRACT

This paper discusses the limitations of the DIS standard with respect to large-scale exercises and proposes enhancements to overcome these limitations. It studies the question of whether the futures of DIS and various commercial interactive applications (such as interactive television and telemedicine) are inter-related in such a way as to provide a low-cost, high performance solution for large-scale simulation exercises. By using data obtained from large-scale exercises, such as I/ITSEC 1994 and STOW-E, this paper highlights the trends in network performance. Study of these trends is useful because it indicates current limitations of DIS implementations. This paper next provides an in-depth discussion of the interactivity problems involving system bandwidth, inefficient protocol definitions, and latency. These issues are discussed both in the context of existing applications as well as future commercial applications.

ABOUT THE AUTHORS

Christina Bouwens is a project engineer with Hughes Training Incorporated, Link Operations, Binghamton NY. She has over six years experience in simulator networking. She currently serves as systems engineer for several Link projects that use Distributed Interactive Simulation (DIS) and Aggregate Level Simulation Protocol (ALSP). Before joining Link, Ms Bouwens was the program engineer for the DIS standards program at the Institute for Simulation and Training. Ms Bouwens has a Masters of Science in Mathematical Science from the University of Central Florida and a Bachelor of Science in Mathematics from Geneva College. She has published several papers on the subject of DIS and interoperability. Ms Bouwens chairs the DIS working group for Communication Architecture and Security.

Mark Juliano is a Software Engineer with AEL Industries, Cross Systems Division where he is studying the impact of missile/aircraft interactions on DIS. Prior to joining AEL Industries, Mr. Juliano was a DIS Consulting Engineer with Encore Computer Corporation, Orlando, Florida. He has several years of experience in distributed simulation and real-time computer systems. At Encore he was the technical lead for DIS and the Assistant Principal Investigator in a Co-operative Research and Development Agreement (CRDA) between Encore and the Naval Air Warfare Center Training Systems Division (NAWC-TSD). Mr. Juliano has a Bachelor of Science degree in Liberal Studies from the University of Central Florida (UCF), and is currently pursuing his Masters degree. He has published two papers on DIS specific issues, and is an active member of the DIS working groups.

Ron Matusof is an applications engineer with BroadVision, Inc., in Los Altos, California. He has over twelve years of experience in analyzing, designing, and implementing large-scale, real-time interactive applications for both commercial and military customers. Prior to joining BroadVision, Mr. Matusof was part of the research and development staffs of both Interactive Network and CAE-Link. He has been involved in simulation networking since 1987 and has been an active participant in the DIS standards development process since the first DIS workshop in 1989. Mr. Matusof holds a Bachelor of Science in Electrical Engineering from the University of Pittsburgh. He has published numerous papers on the subjects of interoperability, mission rehearsal, image generation, radar simulation, and advanced DIS concepts.

Enhancing DIS to Support the Million Player Scenario

Christina Bouwens
Hughes Training Incorporated, Link Operations

Mark Juliano
AEL Industries, Cross Systems Division

Ron Matusof
BroadVision, Inc.
Los Altos, CA

INTRODUCTION:

Darwin's Laws of Technological Evolution

e-volve (i vlv') **vt., vi**. [<L. e-, out, + volvere, to roll] To develop gradually.

All ideas evolve. In general, stronger ideas flourish while weaker ideas become extinct. Evolution is usually subtle, but occasionally ideas radically and fundamentally change. Einstein's theory of relativity, for example, redefined the relationship between space and time and made all previous relationships obsolete. On a less cosmic scale, ARPA's creation of SIMNET altered how the training community views and uses simulation. Distributed Interactive Simulation (DIS) is an outgrowth of SIMNET, and it continues to evolve training and simulation concepts. DIS is not the end of the simulation evolutionary chain -- merely a step toward its future.

At SIMNET's birth in 1983, the underlying architectural concepts were not new -- they had been around since at least the mid 1970's -- but they did require new enabling technologies (Alluisi, 1991). In recent years, DIS has embraced the basic concepts of the SIMNET distributed architecture. The danger is that SIMNET concepts, while highly useful, are based on ideas that are over twenty years old. The tradeoffs which were made by SIMNET architects because of technological constraints do not necessarily apply today.

It was envisioned that non-military applications might leverage DIS development for commercial use. DIS was therefore developed as a set of commercial standards (IEEE, 1995). Its use, however, has been studied and rejected for a number of commercial applications because DIS was not considered commercially viable. The largest envisioned DIS exercise involves 100,000 entities -- a number derived from ARPA requirements for the Synthetic Theater of War (STOW) program (Calvin, Seeger, Troxel, & Van Hook, 1994). While 100,000 entities sounds daunting for the training and simulation community, it is a number generally considered below the "critical mass" necessary for commercial success. As an alternative to DIS, the convergence industry (the industry formed by the convergence of various computer and telecommunication industries) is developing its own capabilities to allow large-scale, real-time interactivity ("Video Games", 1995). This has resulted in creation of products which have a potential for significantly improving the price-performance curve for DIS.

As an example, Catapult Industries currently markets devices which provide large-scale, real-time interaction of video games for a cost of less than seventy dollars per player. It may seem somewhat odd to equate the networking of video games with solving large scale DIS problems, but reviewing commercial game technologies to help reduce cost and improve quality of training exercises is not an original idea. DARPA's Cybernetics Technology Office studied the use of video games technology for application toward tank gunnery training as early as 1981(Alluisi, 1991).

Our research is centered around how new and existing commercial technologies can be used to improve the quality of the DIS network experience, increase the quantity of interactive players to commercially viable levels, and reduce the cost of the interface to consumer pricing levels. We have set a target of 1 million interactive entities -- a number which is considered to be necessary to develop a competitive market -- and a target date of 1998. We have further set a goal of maximizing the ease of transition for existing

DIS applications to participate in the million player scenario. By developing a commercially viable DIS solution, with a defined migration path for existing DIS applications, the cost of performing training and system evaluation using DIS can be substantially reduced. By concurrently solving the problems associated with a million player scenario, we believe that the cost of large-scale military exercises, such as STOW, can be significantly reduced while the quality of the exercise is simultaneously improved.

Our research has three distinct phases to it. The first phase is described in this paper and involves analysis of the existing DIS to characterize large-scale applications and determine the deficiencies which can and should be corrected by the application of new technology and architectural enhancements. The second phase is to design an architecture which can support a minimum of a million simultaneous players in a cost-effective manner. The final portion of this research will simulate the architecture to determine its overall effectiveness.

PROBLEM DEFINITION

The most important part of any research program is defining the scope of the problem. This involves defining bounds that do not unnecessarily restrict the problem space. The goal of our research is to provide an architecture which can support one million simultaneous entities. We bounded this problem by stating that this architecture must be both technologically and economically feasible in calendar year 1998 and must support a migration path for existing DIS applications. Therefore, in addition to defining bounds for the study, we have to make some assumptions concerning the available technology, its associated pricing, and emerging training requirements as they will exist in 1998.

Making assumptions for this research is difficult because the problem space is dynamic and unpredictable. New Protocol Data Units (PDU) are proposed at each DIS workshop. This trend is likely to continue for the foreseeable future. We have limited the scope of our analysis to the protocols that had been approved by the DIS steering committee as of June 1995 -- namely, DIS version 2.0.4. The differences between version 2.0.4 and IEEE-1278.1-1995 are minor and are insignificant to this analysis. We have also defined the 1998 DIS requirements to be those which have been documented in DIS version 2.1.1 (Draft) as of June 1, 1995 (DIS Steering Committee, 1995). This is a reasonable assumption, since version 2.1.1 will likely be the basis of the next update to IEEE-1278.1 sometime in 1997. DIS version 2.1.1 introduces thirteen new PDUs to DIS. Since there is no data concerning these PDUs, we have estimated the frequency of their occurrence.

We are modeling 1998 technology based on Moore's Law -- a relationship between price and performance of processing hardware first espoused by Gordon Moore, founder of Intel, in 1979 (Gilder, 1993). Moore's Law states that the processing price-performance ratio of computers doubles every 18 months. This law has held true since 1979 and all indications are that it will continue for some time to come. We assume a factor of four improvement in processing price-performance over the next 36 months. Bandwidth performance is expected to increase at 5 to 100 times the rate of processing improvements, so we estimate a factor of ten bandwidth improvement between now and 1998.

There are technological and economic impacts on the future of DIS from various activities within the convergence industry. Many companies are realizing that the missing 'killer application' for the entertainment software world is real-time, low-cost connectivity to support popular game platforms (Grosser, 1994). The centerpiece of this connectivity is a device known as an information appliance. The information appliance is the result of the convergence of computer systems, televisions, telephones, and other devices. Market projections indicate that a television set-top box with the approximate power of today's Silicon Graphics Indy workstation will be available in 1998 for less than three hundred dollars (Lambert, 1995). We assume for the purposes of this research that these market projections are correct, and that such an appliance may also be used, unmodified, for DIS applications. One such approach has already been proposed at DIS workshops (Guckenberger, Guckenberger, Green, & Andrews, 1995).

ANALYSIS

A review of the literature shows a relatively limited set of issues concerning the implementation of a large-scale DIS network. An excellent summary of these issues is contained within the proceedings of the 11th

DIS Workshop (Calvin & Van Hook, 1994). Interestingly, the literature shows that these issues apply equally to a wide variety of large-scale, non-DIS applications such as interactive television (Sweeney, 1994) and telemedicine (Murakami, Shimizu, Mikami, Hoshimiya, & Kondo, 1994). The characteristic issues described in the literature are channel capacity (bandwidth), real-time operation (latency), rapid obsolescence of technology, reliability, system size, and electromagnetic interference. Although this list of issues accurately describes the technological problems associated with large scale networking, we wanted to frame the problem in terms of human interaction. Our analysis centers on three issues of interoperability: entity representation, physical interconnection, and information flow.

We analyzed in detail the data from I/ITSEC 1994 collected by Encore Computer Corporation and compared it with data collected during other exercises to determine trends in simulation networking. We analyzed over six gigabytes of data collected over a six-day period beginning on November 26, 1994. We hoped to compare this data against STOW-E data but this was impractical, since our research is currently unclassified. Instead, we have some "ballpark" data from STOW-E to help validate the I/ITSEC 1994 data. Our analysis centered on the following:

1. Number of active entities.
2. Amount of traffic.
3. Amount of traffic generated per site.
4. Amount of traffic per exercise type.
5. Amount of traffic by domain
6. Amount of traffic by DR algorithm.
7. Average Packet Size.
8. Average load.

Our intent is to study the data from a number of different viewpoints to help define a "generic" DIS exercise which could be successfully scaled to one million entities.

ENTITY REPRESENTATION ISSUES

This portion of our research is concerned with how well each entity represents itself to other entities in the distributed environment. This is different from studying the fidelity of the entity. While fidelity is concerned with how the entity represents the real world, we are concerned with how the network limits or enhances the representation of the entity. Our goal is to determine trade-offs which may be made in the network design without limiting the particular types of tasks to be performed.

Entity representation is one of the most studied aspects of DIS. A review of the literature shows major concerns about positional and temporal accuracy for entity representation. These concerns are also documented for a wide range of non-DIS applications including telemedicine, interactive television, video games, and distance learning. Accuracy is a description of how closely the representation of an entity on the network corresponds to the position and time sequence of events computed locally by the entity.

In their current draft of the Protocol Architecture Development Plan, the DIS Special Task Group for Protocol Architectures (STGPA) points out that precision of current PDUs, in most cases, cannot be measured or observed, and that such measurements are not needed (STGPA, 1995). Rather than concentrate on the PDUs themselves, we looked at measurable data points from which we could determine the accuracy of an entity's representation. Experiments have found positional errors to vary with the number of entities involved in an exercise. In smaller exercises (less than 700 entities), an error of about 0.25 meters was found while in larger exercises (1100 entities) the average error exceeded 15 meters (Dille & Swaine, 1993). This research showed that the error was caused by data latency, which increased as data packets were lost during ethernet collisions. A study of the I/ITSEC 1994 data shows similar results, although the delay in this case was due to latency in the long-haul networks. The average latency on one long-haul node of the I/ITSEC 1994 network was measured to be 115.2 milliseconds when data was transmitted over the Defense Simulation Internet (DSI) and 162.1 milliseconds when the data was transmitted from the same node using a BReeze Bridger/Router over a commercial land line (Woodyard, 1995). This yields theoretical errors of approximately 0.5 meters for a vehicle moving at 10 kilometers/hour and 18 meters for a vehicle moving at 300 knots.

The accuracy of entity representation is affected by the choice of dead reckoning algorithm. Dead reckoning is an estimation technique used to reduce the frequency at which entity state PDUs are transmitted. Data from I/ITSEC 1994 shows that most simulations (51.11 percent) chose to use dead reckoning algorithm 2, a method of extrapolation where the future position of an entity is predicted based on the entity's current position and linear velocity (that is, no acceleration terms or rotational components are used). Another 29.21 percent of all entity state PDUs used dead reckoning algorithm 4, which includes both acceleration terms and rotational components. Dead reckoning was not employed by 13.54 percent of all entity state PDUs and the remaining 6.14 percent were dead reckoned using some other algorithm.

An examination of traffic sorted by entity domain shows how the exercise is affected by different entity groups. Entity representation differs between entity domains. During I/ITSEC 1994, air entities generated four times the number of entity state PDUs and electronic emissions PDUs than land entities. Surface entities and sub-surface entities generated these PDUs at roughly the same rate as land entities. For other PDU types, there was no significant difference found between entity domains.

INTERCONNECTION ISSUES

DIS has been designed to allow computers to communicate information between themselves in what is essentially a conversation. As such, it assumes that the physical exchange of data between computers is symmetric (that is, the amount of data transmitted is approximately the same as the amount which is received). By contrast, information systems which communicate directly with humans are asymmetric and are characterized by a large difference between the amount of data transmitted and the amount received Experiments traditionally place the human's information processing at less than 100 bits per second (where a bit represents a Boolean decision) while communication channels range from 105 to 109 bits per second (Flannigan, 1994). Large scale exercises are also highly asymmetric. For very large exercises, there will generally be many more external entities than internal entities. The data from the 1994 I/ITSEC demonstrations highlights the asymmetric nature of the data. For this exercise, a single site broadcast a minimum of 1 packet every five seconds, or 31 bytes/second ("the DIS "heartbeat"), and a maximum of 61.2 packets/second (9364 bytes/second). On the average, 4.8 packets/second (734 bytes/second) were broadcast by an entity, while it received an average of 79.65 packets/second (16,812 bytes/second). This asymmetry is consistent throughout the entire six day I/ITSEC 1994 recording, although the asymmetry becomes more pronounced as the number of entities increases.

During I/ITSEC 1994, the number of entities varied from 5 to 905, yet the average packet size remained almost constant, varying between 139 and 179 bytes. This consistency is apparent throughout the I/ITSEC 1994 data, independent of the exercise type, the number of entities involved, or the total amount of traffic on the network. It is also consistent with the data gathered from other distributed exercises, dating back to the SIMNET program.

There has been a fair amount of research concerning effect of various PDU types on overall bandwidth. Most of this research shows that packets for PDUs other than entity state (such as fire and detonate) occur so relatively rarely that their effect is statistically insignificant. Our review of the I/ITSEC 1994 data confirms this to be the case when all data is reviewed as an aggregate. This is one possible explanation for the consistency we found in average packet size. However, when we separated the data by exercise ID, we found some interesting results, as shown in Table 1.

EXERCISE PDU	Simulation Managemen t	Computer Generated Forces	Naval Battle	Helicopte r Armed Recon	Air to Air	Land War	Preci-sion Strike	Above Real-Time	Joint Ops	Air Defense	I/ITSEC 1994 TOTAL
Exercise Number	4	5	7	9	10	11	12	13	14	15	
Entity State	99.00	14.50	83.00	74.80	99.99	81.00	61.00	99.90	98.00	97.60	88.04
Fire	0.00	0.50	0.00	0.06	0.50	3.00	0.08	0.05	0.10	0.10	0.27
Detonate	0.00	0.50	0.00	0.06	0.50	3.00	0.08	0.05	0.05	0.05	0.27
Create Entity	0.03	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Remove Entity	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Start	0.03	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Stop	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Acknowledge	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Action Request	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Action Response	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Set Data	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Data	0.03	0.00	0.00	0.00	0.00	0.00	0.6	0.00	0.00	0.00	0.07
Electromagnetic Emissions	2.00	0.00	0.00	0.06	0.00	0.00	0.00	0.00	0.80	1.70	0.70
Laser	0.00	85.00	0.00	5.20	0.00	0.00	0.00	0.00	0.20	0.50	2.61
Transmitter	0.00	0.00	17.00	7.10	0.00	0.00	4.90	0.00	0.10	0.10	2.49
Signal	0.00	0.00	0.00	12.20	0.00	0.00	33.00	0.00	0.10	0.10	1.85
Receive	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Stealth	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 1 I/ITSEC 1994 PDU Distribution by Exercise Type (percent)

It is interesting to look at the bandwidth utilization for DIS. In a previous study performed using an ethernet simulation, bandwidth utilization had distinct characteristic differences when the total number of entities exceeded 1000. For these larger exercises, the law of diminishing returns applied. The percentage of useful bandwidth (that is, bandwidth which contains data actually received by another entity) actually dropped as total bandwidth utilization increased (Dille & Swaine, 1993). It was assumed that this was caused by additional packet collisions occurring due to a higher number of entities. We were unable to reproduce this result with the I/ITSEC 1994 data, but believe that such effects will occur for very large exercises.

In a study presented at the 16th I/ITSEC (Swaine & Stapf, 1994) it was shown that for a typical processing board (Motorola MVME-197), approximately 320 microseconds of processing time is required to transmit a single PDU while 72 microseconds is required to receive a single PDU. This data was derived for a particular processor, but it typifies the processing required by a DIS interface.

The highest traffic rate at the 1994 I/ITSEC was approximately 230 PDUs/second (averaged over a one hour time-frame). A typical processor like the MVME-197 requires approximately 73.5 milliseconds/second of processing to handle all of the Entity State PDUs on the network. This is highly significant since the amount of processing time required for a typical single CPU system is greater than one 15 Hz computational frame. This implies that for simulations with an update rate of 15 Hz or better, a separate DIS interface computer is required.

Dead reckoning and coordinate conversions have been shown to be the largest contributors to host CPU usage (Flannigan, 1994). Coordinate conversion is a tradeoff between precision and computational load (Long, Anshuetz, & Smith, (1994) while dead reckoning has been considered a trade-off among three factors: network communication traffic, the computational load, and precision (Miller, Pope, & Waters, 1988). In I/ITSEC 1994, the majority of simulations used extremely simple dead reckoning, thereby optimizing computational load at the expense of precision and network bandwidth reduction.

Latency is an important part of intercommunication. Unfortunately, we were unable to obtain hard latency data due to the use of relative timestamps. Once absolute timestamps and time synchronization are in use, analysis of latency can be properly performed.

INFORMATION FLOW ISSUES

Information flow refers to the movement of information between simulation hosts and the subsequent processing of that information by the hosts. The goal of managing information flow is to reduce the amount of network traffic while minimizing any impact on the task being conducted. In order to understand information flow, we studied how much of the data placed on the network is actually used.

A number of information flow techniques are described in the proceedings from the 11th DIS Workshop (Van Hook, Calvin, Newton, & Fusco, 1994) and can be used to manage information flow. These techniques may be quite useful in limiting the amount of data required for large scale interoperation. Prior research has shown that filtering techniques can be quite effective. A 90% reduction in bandwidth can occur if filtering on visual range is accomplished (Gehl, 1994). At I/ITSEC 1994, the play area was very small and was composed of a high resolution area (where most of the entities congregated) and a low resolution area. Entities inside the high resolution area tended to be within visual range of sixty to seventy percent of all entities, while those in the low resolution area tended to filter out eighty to ninety percent of the entities. It would appear that filtering effectiveness based on visual range is a function of the entity density. However, relevance filtering can be performed on any number of parameters, and appropriate filtering would provide significantly improved results.

DISCUSSION

In this first phase of our work we used the observations discussed in this paper to develop a generic model of how DIS is used and determine the deficiencies which should be corrected by the application of new technology and architectural enhancements.

Can DIS be characterized in generic terms? Are all DIS exercises statistically equivalent? Our review of the literature indicates that there are striking similarities between all large-scale interoperable applications. Additionally, our analysis indicates that the general characteristics of DIS vary little between exercises, regardless of the number of entities involved or the intended purpose of the exercise.

We have developed four models for the "generic" DIS exercise: entity representation, physical interconnection, information flow, and cost. Each model helps us develop an appropriate set of trade-offs to be applied as we develop an enhanced DIS architecture.

Entity Representation: Less accuracy is demanded for entity representation in DIS exercises than for comparable stand-alone simulations. This is perfectly acceptable, since the intended goal of large-scale interoperability is to realistically represent large numbers of "thinking" entities.

How accurate does DIS have to be? The DIS community has strongly resisted the urge to mandate minimum levels of interoperability, leaving that determination to the exercise planner. There are, however, some practical limits to accuracy, based on the temporal characteristics of the network and the precision of the locally computed entity state.

We found that over seventy percent of I/ITSEC 1994 participants applied simple dead reckoning to their representation. This implies that for these simulations, conservation of processing power outweighs accuracy and bandwidth considerations. Similar observations can be made for other PDU types as well. For example, emitter beams were rarely regenerated from the data contained in the Electromagnetic Emissions PDU. Instead, most implementations used the emitting entity ID field to point to a model stored locally in a look-up table. This method compromises the accuracy of the signal reconstruction for improved processing performance.

DIS version 2.1.1 introduces thirteen new PDUs to DIS. We believe that they will not change the nature of the "generic" DIS exercise. Although they have different names and functions, the performance impact of the new PDUs is already characterized by the existing PDU set. For example, the new Newtonian Protocol Collision PDU is roughly equivalent to the current Collision PDU, the Underwater Acoustic and IFF/ATC/Nav Aids PDUs are equivalents to the existing Electromagnetic Emissions PDU, the Simulated Environment PDUs are similar to the current Entity State PDUs, and so on. We have no reason to believe that the entity representation issues will change substantially for this group of PDUs.

Physical Interconnection: DIS is by nature asymmetric and asynchronous. It is therefore inappropriate to think of the physical interconnection between entities in terms of an even, regular flow of data, as one would find on most computer networks. Instead, DIS is more like a distribution system, and is modeled more accurately as a distribution topology.

One interesting possibility is to use the standard telecommunications "tree and branch" topology. Tree and branch is roughly analogous to the human circulatory system, where arteries and capillaries act as the major and minor transport pipelines for the blood. In DIS, the artery is the so-called "firehose" from which each host must filter appropriate PDUs for internal use and the capillary is a set of PDUs generated locally by the host. This opens a number of potential solutions for transport of DIS, including Asynchronous Transfer Mode (ATM) and Asymmetric Data Subscriber Loop (ADSL). It also opens the possibility of bundling DIS packets within another transport mechanism, such as MPEG-2, which is currently finding favor within the telecommunications industry.

Information Flow: Information in DIS follows the paradigm of an asynchronous data flow system. Asynchronous data flow systems are those where all data is passed between processes in the form of messages, and all synchronization is contained within the data.

There are a number of methods available to improve the information flow of DIS. However, the greatest improvement to information flow is to reduce the amount of information which needs to flow between entities. Reductions in individual fields of PDUs, while highly encouraged, do little to improve information flow unless they impact the amount of information in the Entity State PDU (and in the future, in the Environment PDU) since these PDUs represent greater than 85 percent of the traffic in every published DIS data set to date. Based on the I/ITSEC 1994 data, it seems that allowing multiple data representations within a particular PDU would be desirable and could reduce bandwidth requirements significantly. For example, if acceleration data were passed only for those Entity State PDUs which use higher order dead reckoning, overall I/ITSEC 1994 bandwidth would have been reduced by almost ten percent. Recent tests of multiple data representations yielded improvements in Entity State PDU bandwidth of almost thirty percent.

Another useful improvement to information flow is the elimination of both redundant and useless data. For example, the velocity field of the Entity State PDU is redundant whenever the entity is known to be a stationary object (such as an air traffic control tower). As another example, the current DIS uses 48 bits to describe the Entity ID, while 24 bits is sufficient to uniquely address over 16 million entities. The draft of the STGPA Protocol Architecture Development Plan indicates that a reduction of up to ten to one can be achieved by culling the data (that is, removing all data which is irrelevant to other entities) before it is transmitted.

There is also a fair amount of research is currently aimed at data compression. It is generally considered feasible to get a four to one reduction in data by applying appropriate data compression techniques.

Cost: What are the appropriate cost/performance trade-offs for 1998. If we assume the pricing model of the information appliance ($300, or $5/MIPS), what will the entity gateway to the million player scenario cost?

Assuming that the exercise scales linearly, a one million player exercise will require one thousand times the bandwidth of I/ITSEC 1994. At its peak, I/ITSEC 1994 traffic showed almost 300 PDUs/second being transmitted. For a linearly scaled million player exercise, 300,000 PDUs/second would be issued. Using data published at the 16th I/ITSEC (Swaine & Stapf, 1994), it requires about 500 instructions/PDU, including filtering, to receive data from the network. This equates to approximately 150 MIPS to process the network traffic. At $5/MIPS (in 60 MIPS increments), this means an interface that can handle one million entities will be achievable for $1000 to $1500.

The bandwidth requirement for one million entities, if scaled linearly from I/ITSEC 1994, is again 1000 times greater than the worst case we recorded. The worst case we recorded was approximately 50.6 Kilobytes/second. For one million players, the raw bandwidth requirement would be 506 Megabytes/second. However, if we assume four to one reduction in bandwidth due to data compression, this required bandwidth is 126.5 Megabytes/second, which is achievable with 1995 technology.

CONCLUSION

A million player DIS exercise is technologically feasible in 1998. It is also economically feasible, given the rapid pace at which an interactive infrastructure is being deployed.

What trade-offs should be made in the architecture which supports one million entities? Based on the DIS data to date, it would seem appropriate to trade the accuracy of the exercise against the cost of processing and bandwidth. But application of Moore's law suggests that in 1998 the cost of achieving higher accuracy will not be a significant barrier to all but the largest exercises. The architecture of the future should be developed by balancing five factors: processing, bandwidth, precision, cost, and number of entities. Unlike the current DIS, the enhanced architecture must not assume that the trade-offs can be fixed into the architecture. As we develop a million player architecture in phase 2 of this research, it is our goal to allow DIS developers to choose the appropriate trade-offs which make sense for a given application.

We are encouraged by the direction that the DIS STGPA has taken in improving the information flow of DIS, and believe that these types of initiatives should continue in the DIS community. We are also encouraged by the rapid progress being made in the commercial world, and hope that innovations in large-scale interoperation can be applied to DIS.

ACKNOWLEDGMENTS

The STGPA is currently developing enhancements to improve scalability and interoperability in future DIS applications. They run an unmoderated e-mail reflector which we have been closely monitoring. We have based some of our research on unpublished preliminary test results which Russ Gminder from IST has placed on the reflector. Also, some of the concepts for data compression are based on the work of Danny Cohen and Moshe Kirsch at Perceptronics.

REFERENCES

Alluisi, E.A. (1991). The development of technology for collective training: SIMNET, a case history, Human Factors, 33(3).

Calvin, J.O., & Van Hook, D.J.. (1994). AGENTS: An architectural construct to support distributed simulation. Proceedings of the 11th DIS Workshop, (pp 357-362). Orlando, FL.

Calvin, J.O., Seeger, J., Troxel, G.D., & Van Hook, D.J. (1995). STOW real-time information transfer and networking system architecture. Proceedings of the 12th DIS Workshop (pp 343-353). Orlando, FL

Dille, J.W., & Swaine, S.D. (1993) Discrete Event Simulation and Analysis of DIS Network Architectures. Proceedings of the 14th Interservice/Industry Training Systems and Education Conference, (pp 167-176). Orlando, FL.

DIS Steering Committee. (1995). Standard For Distributed Interactive Simulation--Applications Protocols Version 2.1.1 (Working Draft) (IST-CR-95-06). Orlando, FL: University of Central Florida, Institute for Simulation and Training.

Flannigan, J.L. (1994). Technologies for Multimedia Communications. Proceedings of the IEEE, 82, (4).

Gilder, G. (1993, March). George Gilder's Telecosm The Bandwidth, Forbes ASAP, pp 163-177.

Grosser, A. (1994, December). Connectivity -- The Future of Interactive Electronic Entertainment, The Red Herring, pp 46-48.

Guckenberger, D., Guckenberger, L., Green, J., & Andrews, B. (1995). DIS Synthesis with Interactive Television (DISSIT): Revolution of the Educational Paradigm. Proceedings of the 12th DIS Workshop, (pp 165-174). Orlando, FL.

Institute of Electrical and Electronics Engineers. (1995). Standard for Distributed Interactive Simulation -- Application Protocols (IEEE Standard 1278.1-1995). New York, NY: IEEE Standards Department.

Lambert, P. (1995, May). The Set-Top Design Conundrum, On Demand, pp 24-31.

Long, R.A., Anshuetz, E.E., & Smith, L.R. (1994). Performance Limitations of the DIS interface, Proceedings of the 16th Interservice/Industry Training System and Education Conference, (pp 4-23). Orlando, FL.

Miller, D.C., Pope, A.R., & Waters, R.M. (1988). Long Haul Networking of Simulators, Proceedings of the 10th Interservice/Industry Training Systems Conference, (pp 577-582). Orlando, FL..

Murakami, H., Shimizu, K., Yamamoto, K., Mikami, T., Hoshimiya, N., & Kondo, K. (1994). Telemedicine using mobile satellite communication IEEE Transactions on Biomedical Engineering, 41, 488-497.

Special Task Group for Protocol Architectures. (1995). Protocol Architecture Development Plan, (Report Number DRAFT). Orlando, FL: Institute for Simulation and Training, University of Central Florida.

Swaine, S.D., & Stapf, M.A. (1994). Large DIS Exercises -- 100 Entities Out of 100,000, Proceedings of the 16th Interservice/Industry Training System and Education Conference, (pp 4-13). Orlando, FL.

Sweeney, J.O.P. (1994). An introduction to interactive television. Proceedings of the International Broadcasting Convention, (ICE Conference Publication No 397), (pp 503-508), London, England.

Van Hook, D.J., Calvin, J.O., Newton, M.K., & Fusco, D.A. (1994). An Approach to DIS Scaleability, Proceedings of the 11th DIS Workshop (pp 347-355). Orlando, FL.

Video game platforms get interactive. (1995, January). Interactive Content. pp8-9.

Woodyard, J.M. (1995). BReeze bridge/router timing data from I/ITSEC 1994. Proceedings of the 12th DIS Workshop, (pp 35-40). Orlando, FL.