SCALEABILITY TOOLS, TECHNIQUES, AND THE DIS ARCHITECTURE

Daniel J. Van Hook, Deborah J. Wilbert, Richard L. Schaffer
Loral Advanced Distributed Simulation
Cambridge, Massachusetts

Walter Milliken
BBN Systems and Technologies
Cambridge, Massachusetts

CDR Dennis K. McBride, PhD.
Advanced Systems Technology Office
Advanced Research Projects Agency

ABSTRACT

"Scaleability" is the concept of dramatically increasing, or scaling up, the number of simulated vehicles (and other "entities"), human participants, and geographically dispersed sites involved in a Distributed Interactive Simulation (DIS). ARPA is developing a scaleability toolset which will allow the development and evaluation of DIS exercise scenarios, scaling and network algorithms and techniques, and network topologies in a simulation environment with the goal of developing solutions to the DIS scaling problem. In addition to giving an overview of the architecture and capabilities of the scaleability toolset, this paper presents an example of, and representative results from, employing the toolset to characterize a specific promising scaling technique, the grid-based geographic filtering algorithm.

ABOUT THE AUTHORS

Daniel J. Van Hook is a senior scientist with Loral Advanced Distributed Simulation. He holds an S.M. in Electrical Engineering and the Electrical Engineers' degree, both from the Massachusetts Institute of Technology, and a B.S. from Rutgers University, also in Electrical Engineering. Mr. Van Hook has worked in the area of distributed simulation for the past eight years, including the SIMNET Large Scale Simulation project. His focus has been on simulation system architectures with an emphasis on software development. He has participated in developing vehicle, sensor, and radio simulations, digitized voice communication systems, network interfaces and protocols, graphics algorithms and system software.

Richard L. Schaffer is a scientist with Loral Advanced Distributed Simulation. He holds a B.S. in Physics from the Massachusetts Institute of Technology. Richard has worked in the area of distributed simulation for the last eight years, including the SIMNET Large Scale Simulation project. Mr. Schaffer has been actively involved in the Distributed Interactive Simulation standards process since its inception in August, 1989. As part of that effort, he has chaired the DIS Articulated Parts and Radio Communication subgroups.

Deborah Wilbert is a scientist with Loral Advanced Distributed Simulation. She holds a Bachelor of Science in Computer Science and another in Science and the Humanities with a concentration in writing, both from the Massachusetts Institute of Technology. Ms. Wilbert has worked in the area of distributed simulation for the past five years, including the SIMNET Large Scale Distributed Simulation project. She has a background in network management, has developed a simulator manager, and is currently a member of the Simulation Management subgroup working on Distributed Interactive Simulation standardization. She has also participated in the development of C^3 simulation systems, semi-automated forces, and protocol definitions.

Walter Milliken is a senior scientist with BBN Systems and Technologies' Network Technology Department. He holds an M.S. in Electrical Engineering from Stanford University and B.S. degrees in Electrical Engineering and Computer Science, both from Washington University in St. Louis. Mr. Milliken has worked in the areas of computer networking and parallel processing systems for the last fourteen years, including work on the Wideband Satellite Network, the Butterfly parallel computer and its Chrysalis operating system, the Modular Tactical Gateway project, and network research in time synchronization and multicast protocols. His current focus is on real-time multicast networking technologies.

Commander Dennis K. McBride is Program Manager for Warfighting Simulation, Advanced Systems Technology Office, Advanced Research Projects Agency, Arlington, Virginia. He directs the development of innovative simulation technology to enable manned, force-on-force warfighting preparation and weapon system concept evaluation. Dr. McBride received a Direct Commission in 1979 following completion of the Ph.D. program in Mathematical Learning Theory at The University of Georgia. After completion of Navy Primary Flight Training/Flight Surgeon Program at the Naval Aerospace Medical Institute, he was designated an Aerospace Experimental Psychologist in 1980, and has since completed several tours in Navy RDT&E. His assignments included Principal Investigator for Maintainability Design, Naval Air Development Center; Division Chief for Basic Research in Aviation Man-Machine Interface Design, Naval Aerospace Medical Research Laboratory; Head, Manned Systems Evaluation Lab, providing Simulation T&E Support for F-14D and EA-6B, Pacific Missile Test Center; and Chief Simulation Officer, Manned Flight Simulator - directing T&E Support for TACAIR and various Avionics Programs, Naval Air Test Center. Dr. McBride earned additional post-Doctoral Masters Degrees in Public Administration (Science, Technology & Govt.); and in Systems from the University of Southern California; and completed the Flight Test Engineering Program (University of Tennessee Space Institute) while assigned to NATC. He has over 75 technical publications and reports in systems, aeromedical-human factors, psychological science, and flight test engineering, earned several military decorations, including Navy Achievement, Commendation, and Joint Service Commendation (Desert Storm Impact Award) medals, as well as the L.P. Coombes (Australian Defence Science Organization) Medal. He was chosen by the U. S. Navy as a NASA Astronaut Candidate. His principal flying responsibilities have been in fighter aircraft Test & Evaluation.

SCALEABILITY TOOLS, TECHNIQUES, AND THE DIS ARCHITECTURE

Daniel J. Van Hook, Deborah J. Wilbert, Richard L. Schaffer
Loral Advanced Distributed Simulation
Cambridge, Massachusetts

Walter Milliken
BBN Systems and Technologies
Cambridge, Massachusetts

CDR Dennis K. McBride, PhD.
Advanced Systems Technology Office
Advanced Research Projects Agency

MOTIVATION

The ability of DIS networks to support large scale battle simulations is critical for many anticipated applications of DIS exercises, such as the tri-service Louisiana Maneuvers and ARPA's Advanced Technology Demonstration One. As the size of DIS exercises is scaled up from hundreds to thousands and eventually tens and even hundreds of thousands of simulation entities interacting in simulation exercises, it will become uneconomical and impractical to adequately scale up the supporting communication and computation resources, given the existing broadcast-like DIS communication model. While the DIS standard in principle supports exercises of any scale, limitations will continue to manifest themselves and grow in at least three system components: wide area networks, local area networks, and simulation hosts and their network interfaces. Techniques and approaches are called for to obtain the maximum level of performance at whatever level of communication and computation technology is applied to the problem. Recognizing this need, ARPA is developing a "scaleability" toolset.

SCALEABILITY TOOLSET OVERVIEW

The scaleability toolset is intended to provide a testbed for developing and evaluating solutions to the DIS scaleability problem. The goal is to provide a system that allows "simulation of a simulation" so that alternative approaches may be evaluated at relatively low cost and risk before committing to an architecture or design. The toolset provides a flexible environment for examining the parameters of and interaction between the options available to developers and planners of DIS architectures, exercises, and protocols, including:

- network topologies
- characteristics of links, gateways, etc.
- network algorithms
- exercise and simulation scenarios
- scaling algorithms and techniques for reducing bit and packet per second requirements

With the toolset, a user may develop an exercise scenario and evaluate that scenario in the context of its network resource requirements. The toolset provides a framework to help determine answers to the following important types of questions:

- "To what extent does a proposed scaling algorithm or technique affect network performance?"

- "What approaches to network-level algorithms for congestion control, routing, bandwidth allocation, and delay guarantees are most advantageous for distributed simulation systems?"

- "To what extent will a contemplated exercise scenario utilize available network resources and topology, or conversely, what network resources and topology are required to support a contemplated exercise scenario?"

- "What is the effect of a proposed protocol change on network resources?"

Figure 1 illustrates the functional structure of the scaleability toolset. The major components of the toolset are:

- scenario creation and simulation tool
- scaling algorithm tool
- network simulation and display tool
- scenario viewing and replay tool
- data analysis tools

Scenario Creation and Simulation Tool

The scenario creation and simulation tool, based the ARPA ODIN Semi-Automated Forces (SAF) system [1] is used to develop and simulate exercise scenarios. The toolset user specifies the scenario during an off-line *scenario creation phase*. The kinds of data needed to specify a scenario include: participating units, unit composition, initial unit positions, distribution of units to network sites, objective(s) for each unit, and specification of the terrain the exercise will be conducted on. The scenario is simulated by the SAF to produce a Scaleability Logger Format (SLF) file that represents in an abstracted form the data packets that would flow between network sites as a result of entity activity and interactions in the exercise. Additional data including timestamps and unit state are added to the output SLF file to provide information needed for other components of the toolset. The SLF format has been carefully designed to optimally trade off storage space, runtime performance, and the data requirements for scaling algorithm development and network simulation. A full discussion of the SAF as modified and used in the scaleability toolset and of the SLF data format and the tradeoffs involved may be found in reference [2].

Scaling Algorithm Tool

The scaling algorithm tool is a flexible testbed for developing and controlling scaling algorithms within the context of the scaleability toolset. It provides facilities for integrating new algorithms into the system according to well defined and documented interfaces. Algorithms are applied to a stream of SLF data (produced by the simulation tool) in order to reduce bit and packet per second rates. Algorithms may be applied in controllable configurations, switched in/out, and parameterized at run time under control of a keyboard interface. The scaling algorithm tool aggregates individual packets into traffic flows that form the input data for and conform to the input interface definition of the network simulation tool.

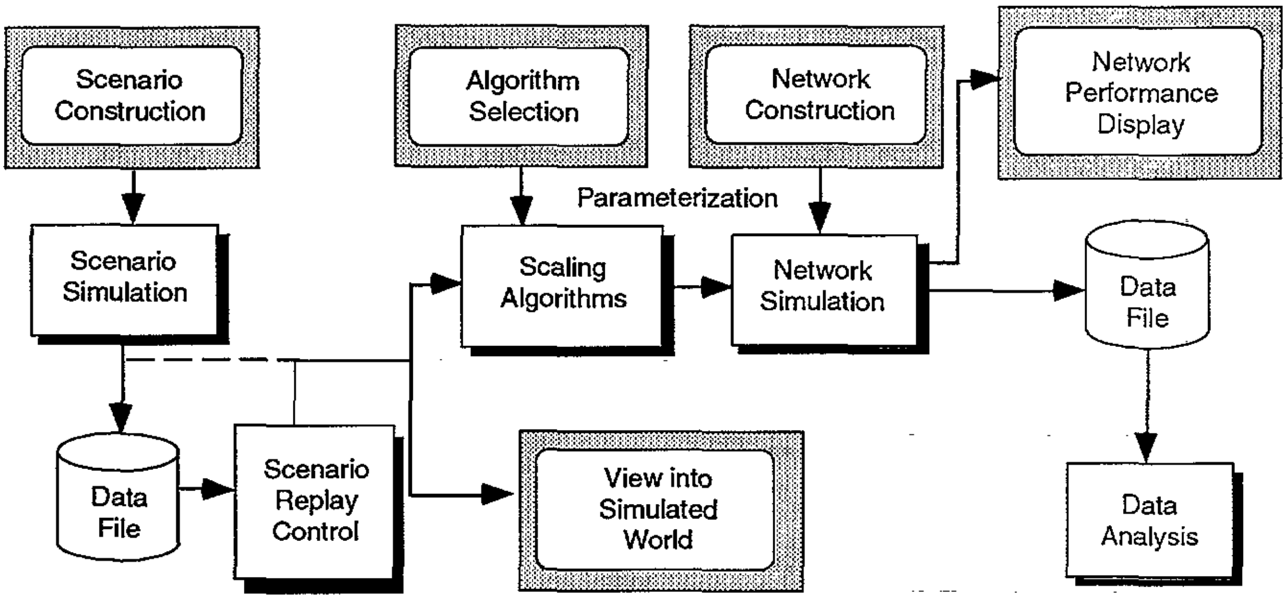

Figure 1 Scaleability Toolset Architecture

Network Simulation and Display Tool

The network simulation and display tool is a flow-based simulator that supports flexible and modular facilities for specifying network topologies and performance characteristics. Networks are modeled as interconnections of network elements such as switches, links, gateways and security devices. The network simulation models aggregated flows through the network as opposed to modeling the transmission of each each individual packet. As such, it produces a medium resolution prediction of network behavior that is consistent with what real physical networks exhibit. A flow-based approach was adopted in order to achieve fast (close to or preferably faster than real-time) processing so as to facilitate a responsive experimentation and evaluation environment. Since multicast capabilities are desirable for several proposed scaling algorithms, the network simulation incorporates support for multicast groups and routing schemes, a feature unavailable in commercially available network simulators. The input interface is a set of standardized messages consisting of timestamped flow and multicast group operations. The network simulator also calculates and gathers extensive performance data, including:

- offered and actual loads for switches and lines
- line transmission and queuing delays
- site LAN outgoing and incoming loads
- delay and drop rates by source site and destination group
- delay histograms of received traffic
- multicast group usage

Run-time performance monitoring indicators include histogram and strip chart displays of delay and loading data and color variation to show utilization levels. In addition, the network simulator includes facilities for outputting performance data to a file for post analysis by commercial data analysis software packages. Both a human-readable ascii/tab-separated format and a more compact self-describing binary format are supported, allowing flexibility in choosing an analysis package. The example analysis described later in this paper made use of BBN/Probe™ [3], which is well suited to flexibly processing and presenting the large volumes of data that are generated. We have made use of the self-describing nature of the binary statistics file format to construct tools that automatically generate analysis functions and scripts to control BBN/Probe processing.

Scenario Viewing and Replay Tool

The scenario viewing and replay tool, based on the ODIN SAF Plan-View Display (PVD) [1], provides a window onto the simulated battlefield. This tool presents the user with a display of the battlefield terrain and the entities participating in the battle, depicted with either unit symbols or as individual entities. By use of this tool, the scaleability toolset user can observe and correlate activity in the simulated world with network performance data reported by the network simulation and display tool. Facilities available to the user include:

- map zoom and scroll
- query an entity for information
- color code forces based on simulation site
- display as units or entities
- terrain analysis: selective terrain feature display, intervisibility, cross
- sections, measurements, coordinate conversions

In addition, a VCR-like control panel is provided for controlling SLF file replay. This control panel supports functions such as pause, play, seek, stop, filename specification, and simulated time display.

The following section illustrates a use of the toolset in detail and the process of analyzing a particular algorithm using a specific network and scenario definition.

SCALEABILITY TOOLSET APPLICATION EXAMPLE

This section provides an example of how the scaleability toolset has been used to investigate a particular scaling algorithm, called grid-based geographic filtering. The methods and results presented here are not an exhaustive characterization of this algorithm but are intended to illustrate how the toolset can be applied to algorithm analysis and by extension used for analysis of networks and simulation scenarios. The series of steps we took to use the toolset for our investigation of the grid algorithm is illustrative of its intended use. Before proceeding to the example, however, it is useful to highlight some of the issues and approaches in grid-based geographic filtering.

Grid-based Geographic Filtering

Geographic filtering is a general technique for reducing the demands a simulation exercise places on the bandwidth, processing, and memory requirements for networks, network interfaces, and simulation hosts. The fundamental concept of geographic filtering is that a simulated entity only needs information about other entities that lie within its region of interest. The region of interest can be thought of as the geographic region of the simulated world that lies within range of the supported sensors, e.g., visual, radar, acoustic, etc. Geographic filtering algorithms arrange to selectively deliver state and event data to simulation hosts based upon the region of interest of locally simulated entities.

One way of implementing the concept of geographic filtering is to associate multicast groups with terrain regions. The multicast routing and addressing capabilities of the underlying network are used to minimize the amount of traffic delivered to each simulator which falls outside its entities' regions of interest. Perhaps the simplest method of assigning terrain regions to multicast groups is to impose a uniform grid system on the simulated terrain and assign a multicast group for each grid cell. This results in each multicast group covering a rectangular solid with a grid cell as its base and infinite height.

Multicast routing and addressing is a developing technology; those networks which support multicasting often impose severe limits on the number of groups and on the rate at which group subscriptions may be changed. While it is appropriate to investigate an algorithm which requires the use of this developing technology, the algorithm should try to optimize its usage and a comprehensive analysis of the algorithm should include an estimate of the requirements for multicast services.

State updates for each entity are sent to the multicast group associated with the cell in which the entity lies. Each simulator joins the multicast groups for the cells that intersect its entities' regions of interest, so that traffic from entities in only those cells are received. As an entity moves across the battlefield, some grid cells will fall out of its region of interest while new grid cells come into range.

Our grid algorithm simulation manages subscriptions to cell multicast groups on a per site basis, whereby an agent notes the state of the locally simulated entities and endeavors to subscribe to the cells overlaid by the union of the regions of interest of the entities simulated at that site. For the purposes of our algorithm analysis, we model an entity's region of interest as an infinite height cylinder with a radius equal to the entity's viewing range. For more optimal cell group subscription, more detailed information about the internal state of each simulated entity than is available to an external entity is required. For this case, it is appropriate for each simulator to determine which multicast groups it should be joining or leaving (rather than being managed externally) since each simulator has the most specific knowledge of the range and scope of its own entities' sensors. Simulators supporting sophisticated implementations of the algorithm could model regions of interest as conic sections for each sensor (by tracking sensor orientation) to reduce the number of unnecessary cell subscriptions. Cell subscription could be further minimized by factoring in range reductions due to terrain or other obstacles (e.g, other entities, weather, chaff).

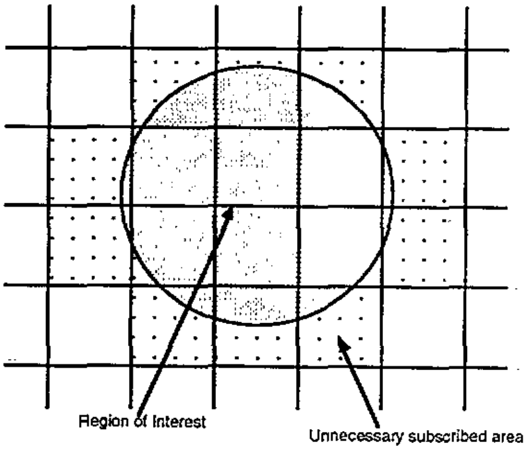

Figure 2 Region of Interest

An important parameter for this algorithm is the cell size. The tradeoff to be considered is that smaller cells result in more efficient filtering, but increase the number of multicast addresses which must be supported by the network and the rate at which the subscriptions to those groups must be changed. As figure 2 shows, a number of cells are only partially overlaid by the region of interest, indicated by the circle. Even though only part of these cells lie within the region of interest, traffic from entities in any part of the cell (including those areas outside the region of interest) is forwarded. Smaller cell sizes result in less unnecessary terrain being included in the set of cells whose groups are subscribed to, with the potential benefit of receiving less traffic from outside the region of interest. One should also note that in addition to increasing the number of multicast groups which must be supported, smaller cells require each simulator to add and drop multicast groups more rapidly as the entity traverses the terrain.

Enhancements to this basic algorithm may prove to yield substantial benefits. For example, dynamically assigning groups to cells, so that only cells that have entities in them or that lie in the region of interest of an entity actually consume a multicast group. This approach may result in conserving multicast addresses for battle scenarios that have entities widely dispersed across the simulated terrain area.

Another possible extension is that multiple grid systems may allow better control of filtering and more optimal consumption of scarce multicast resources for certain battle scenarios. Yet another possibility is to use a non-uniform grid so that the cells are smaller in regions only where finer grained filtering is required and larger where it is not. Finer grained filtering can be useful in areas heavily populated with entities, so that when another entity's region of interest borders on that area, a large amount of traffic from adjacent entities can be eliminated. Fine-grained filtering could also be used to support particularly heavily loaded network components or simulation hosts by clustering around the regions of interest of the entities served by those components. Optimization to this level would require either carefully coordinated scenario construction and grid cell design or a sophisticated dynamic mechanism for determining the optimal grid assignment on the fly.

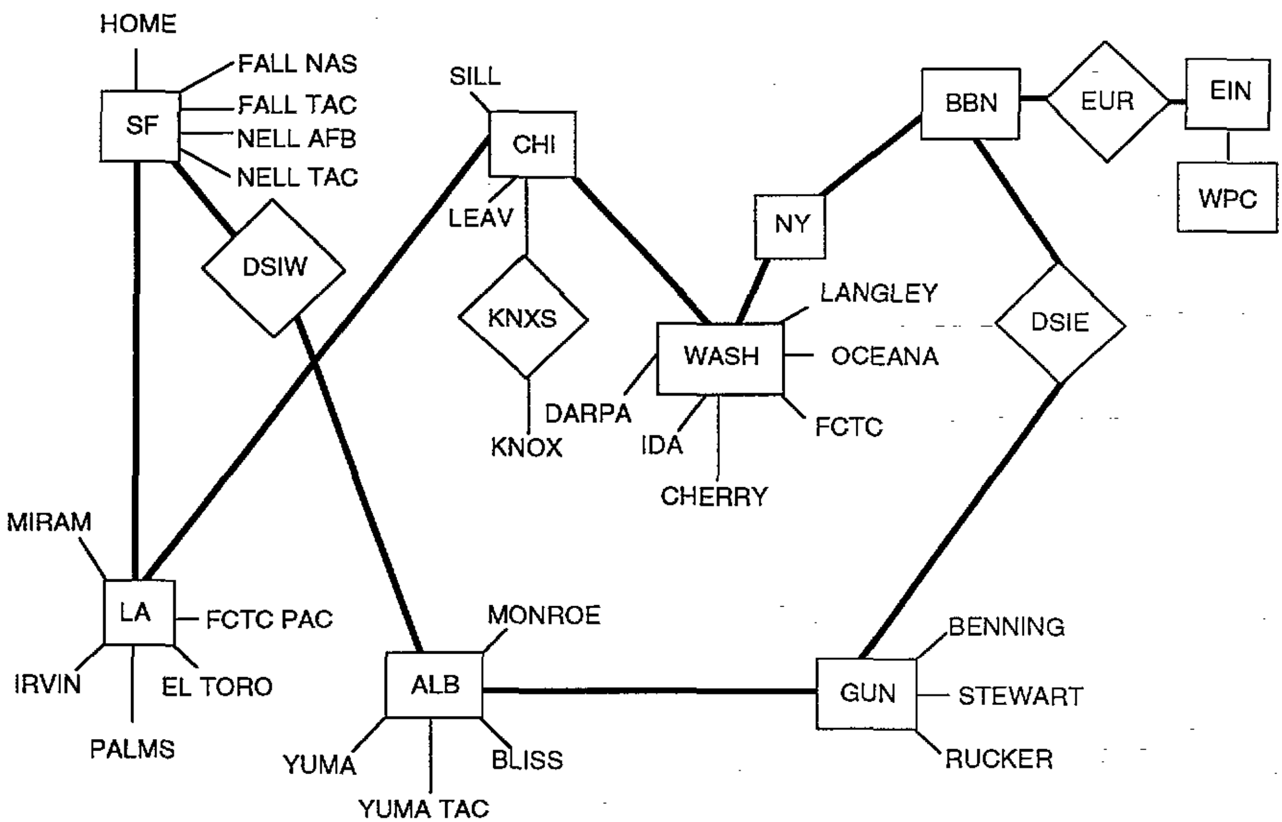

Figure 3 Experimental Network Model

Experimental Parameters and Procedures

Our approach to investigating the effectiveness of the grid-based geographic filtering algorithm consisted of comparing the performance data logged by the network simulator toolset component while the scaling algorithm was in operation (the test cases) with that logged while the scaling algorithm was not in operation (the baseline case). Two test cases were employed: one with a grid cell size of 2.5 by 2.5 kilometers, the second with a grid cell size of 10 by 10 kilometers. The viewing ranges for both test cases were 3.5 kilometers for ground entities and 7 kilometers for air entities (these settings determine the entity region of interest). For both the test and baseline cases, a single battle scenario and network model were employed. For the baseline case (the current DIS communication model), data are broadcast throughout the network, so significantly higher loads were observed.

The network model used in our investigation of the grid-based algorithm is illustrated in figure 3. This hypothetical network is loosely based on the DSI net with the following characteristics and enhancements:

- 27 connected sites
- T1 (1.54 Mbps) tail circuits connecting sites to the backbone
- FDDI LANs at each site
- T3 backbone (45 Mbps)
- switches commensurate in capacity with the T3 backbone

We selected the existing DSI net as the basis of our hypothetical network because it is the only existing network currently being used to host large-scale interactive simulation exercises. We chose to increase the capacity of some network components in order to provide a feasible base for the type of scenario we were interesting in investigating, involving 10,000 or more simulated entities. Figure 3 does not represent an actual planned expansion of the DSI network.

The battle scenario used in the evaluation was developed using the scaleability toolset scenario construction tools with inputs from military subject matter experts so as to insure that it accurately reflected a typical, large scale military exercise. The scenario consisted of five orange divisions (three motorized rifle, two tank) attacking two blue divisions (one tank, one mechanized infantry) across a broad front. Orange and blue aircraft and helicopters carry out missions at intervals throughout the scenario. The orange forces outnumber the blue forces by a factor of approximately three to one, with a total of approximately 10,000 entities represented. The battle takes place within a 50 by 75 kilometer region on simulated Ft. Knox terrain and takes approximately 4 hours of simulated time. The final result is that the orange forces push the blue forces back but the blue forces manage to prevent a complete breakthrough by the orange forces.

The performance measure employed in this study consisted of a measure of the reduction in offered load from the network to the connected sites as a result of applying the grid-based geographic filtering algorithm. This measure may be described as

$$reduction(n) = \frac{\sum_{m=1}^{M} load_b(n, m) - \sum_{m=1}^{M} load_t(n, m)}{\sum_{m=1}^{M} load_b(n, m)} \quad (1)$$

where M is the number of sites and $load_b(n, m)$ and $load_t(n, m)$ are respectively the loads offered to site m in timestep n for the baseline (no scaling algorithm) and test cases (scaling algorithm in operation). For this measure, the mean reduction over one minute segments spaced at ten minute intervals throughout the battle scenario was gathered and plotted. This measure provides a reasonable high level view of the effectiveness of a scaling algorithm in reducing bandwidth requirements.

Of note in interpreting this performance measure is the fact that the aggregate offered output load will generally contain multiple contributions from each packet that enters the network. Without any traffic reduction algorithm (the baseline case), the network will attempt to deliver every packet to every site. Particularly effective algorithms will reduce (or completely eliminate) the number of sites to which a packet must be delivered.

Experimental Results and Interpretation

Figure 4 graphs the reduction of aggregate output load reduction, described in the previous section, achieved by the grid algorithm as a function of simulated time. Results are shown for two grid cell sizes: 2.5 by 2.5 kilometers, and 10 by 10 kilometers. The more effective filtering action of the smaller cell size is evident from the figure. As described previously, this effect can be attributed to the reduction in area that is part of the cell multicast groups subscribed to that lies outside the circular region of interest with a consequent reduction is unnecessary received traffic. Of note is the drop off of effectiveness as the battle scenario progresses. This drop can be attributed to intermixing of the simulated military units which occurs later in the battle scenario - more entities from different sites come into contact (are positioned within each other's regions of interest) later in the exercise, requiring more traffic to be forwarded.

Two cautions should be noted when interpreting these results. First, while this measure of effectiveness yields useful insight into how well a scaling algorithm will perform in the aggregate, individual network elements may still be overloaded even when this measure predicts outstanding performance. This effect was in fact observed while conducting these experiments. Therefore, it is important that any full characterization of an algorithm requires that additional and diverse performance measures and approaches be applied. Relying on only a small set of characterizations would be unwise. We are currently investigating other ways of calculating measures of effectiveness for algorithms, which include considerations of network congestion, dropped packets, delay and individual network element performance. Second, we believe that the network topology and battle scenario may strongly influence the results obtained for any particular measure of effectiveness. Therefore, it is important that any full characterization of an algorithm requires experience with diverse scenarios and networks in order to determine their impact on scaling algorithm performance.

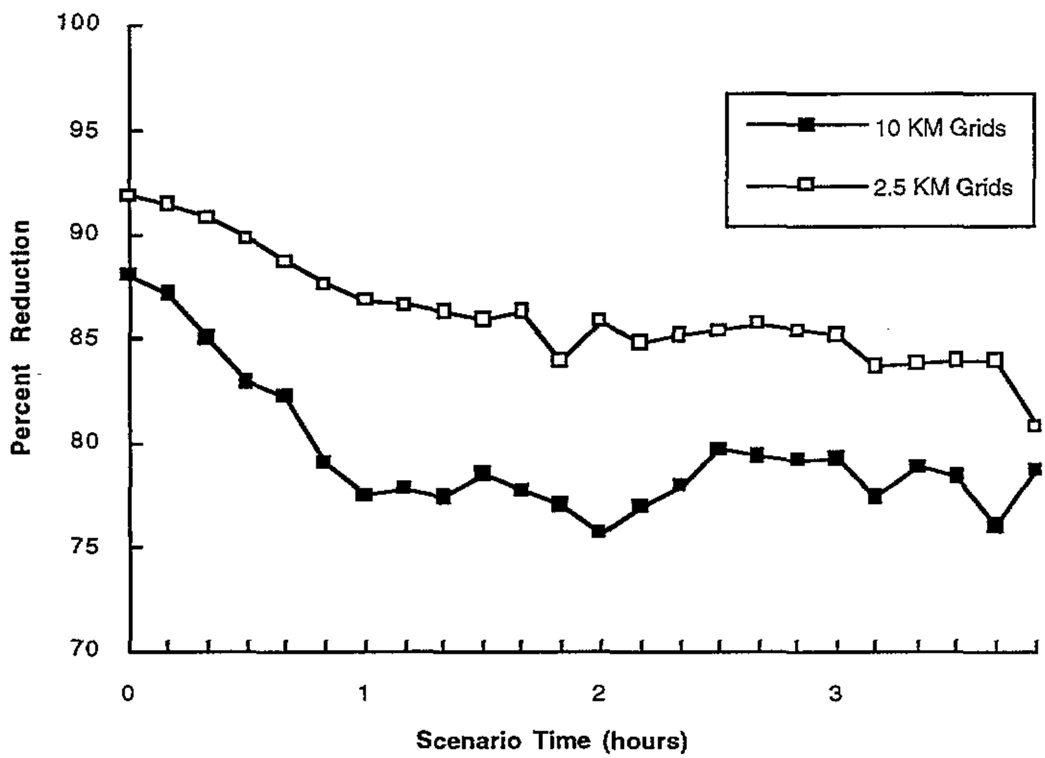

Figure 4 Reduction in Aggregate Offered Output Load

CONCLUSIONS

A scaleability toolset is being developed for the purpose of providing a testbed for developing solutions to the DIS scaling problem. This toolset enables development and evaluation of DIS exercise scenarios, scaling and network algorithms and techniques, and network topologies in a simulation environment, i.e., a simulation of a *simulation*. Not only can the toolset be applied to algorithm analysis, as illustrated by the example given in this paper, but it may also be applied to battle scenario development and network development.

As our analysis shows, grid-based geographic filtering algorithms that use multicasting can reduce network traffic dramatically over broadcast schemes. Therefore, it is fruitful to pursue further investigation of this promising family of algorithms for solving the scaleability problem and also to consider enhancing currently available multicast facilities. We are currently pursuing the analysis of enhancements to the simple single-grid algorithm. In addition, we are also considering other forms of geographic filtering that do not make use of grid-based techniques as described in this paper and also other families of algorithms which could be coupled with geographic filtering algorithms to form a comprehensive approach to solving the scaling problem.

Additional useful measures for characterizing scaling algorithm performance in terms of delay reduction, dropped packets, congestion characteristics, and individual link and switch loading have been developed or are under development. Our experience with these new approaches will be reported in the future.

More research, both in terms of possible algorithms and ways of measuring their effectiveness, will provide a good foundation for selecting appropriate traffic reduction algorithms. Good algorithm selection will be needed for DIS exercises to meet the goal of very large scale tri-service training on networks of the foreseeable future.

ACKNOWLEDGEMENTS

The development of the scaleability toolset is being sponsored by the Advanced Research Projects Agency, Department of Defense and monitored by the Topographic Engineering Center under contract DACA 76-91-C-0006. CDR Dennis K. McBride, USN, is the ARPA Program Manager and has provided the guidance. Debbie Deutsch, Josh Seeger, and Shawn Smith of BBN Systems and Technologies and Barbara Perry, Maureen Saffi and Rob Vrablik of Loral Advanced Distributed Simulation have all made significant contributions, as has Mr. David A. Fusco of the Naval Command, Control and Ocean Surveillance Center, RDT&E Division (NRaD), San Diego, CA.

REFERENCES

1. Saffi, M. "The OBG/SAF Interface User Guide," version 4.2.0, Bolt, Beranek, and Newman, Inc., August 14, 1992.

2. Vrablik, R. and Wilbert, D. "The Use of Semi-Automated Forces to Simulate a 10,000 Entity Exercise," presented at the Third Computer Generated Forces and Behavioral Conference, March, 1993, Orlando, Florida.

3. "BBN/Probe Command Reference Manual", BBN Systems and Technologies, 1992