Simulation of Dual-Mode Multicast Using Real-World Data

J. Mark Pullen Elizabeth L. White
Department of Computer Science and
Center for Excellence in
Command, Control, Communications and Intelligence
George Mason University

keywords: Communications Architecture, Multicast, Scalability

ABSTRACT

In large DIS exercises, many geographically separated backbone sites are connected via a Wide–Area Network (WAN). These backbone sites support local simulators and satellite sites where smaller numbers of simulators are located. Dual-mode multicast is a scalable approach to DIS communication between simulator sites that uses exercise-level multicast groups at the global level to simplify WAN communication, while using grid-based multicast groups at the local (LAN) and satellite site level to minimize simulator input. This is implemented by extending the application gateway functionality to filter and locally multicast PDUs from other sites.

Our prior work used simulation techniques to study the amount of filtering done when using dual-mode multicast communications; however, the simulations built were theoretical. In the work described in this paper, we have extended our simulation to model different types of entities and basic terrain information and have run these simulations using data from an actually planned large-scale DIS exercise. The goal of this work is to assess what the performance impact of using this technique would be on a realistic exercise.

1. INTRODUCTION

The primary mission of DIS (Distributed Interactive Simulation) is "to define an infrastructure for linking simulations of various types at multiple locations to create realistic, complex virtual worlds for the simulation of highly interactive activities."[DIS94] In large DIS exercises, many geographically separated backbone sites are connected via a Wide–Area Network (WAN). These backbone sites support local simulators and satellite sites where smaller numbers of simulators are located.

As the size of these distributed simulations grows, scalable approaches to reducing WAN traffic, LAN traffic and simulator input have been sought. The claim has been made [HCNF94] that even if it were possible to transmit all of the data generated to each participating simulator, 99% of the information would not be used by the receiver. However, simply limiting the amount of data that is sent to the WAN is not a good solution due to the increasing presence of wide-area and global viewers associated with exercises.

Geographic filtering is a type of relevance filtering designed to reduce the amount of unnecessary traffic received at each simulator. The idea behind geographic filtering is that each entity has some region of interest based upon the inherent sensor limitations of the entity such as viewing range or available sources of input (radar, radio, human vision, etc). Each entity belongs to a multicast group to which it sends PDUs. It also is interested in and receives messages from a set of multicast groups associated with *regions of interest*. As the exercise progresses, the membership in the various multicast groups will constantly change as entities change their regions of interest.

Dual-mode multicast is a scalable approach to DIS communication that combines the benefits of implementing geographic filtering using multicast technology with simplicity at the global (WAN) level of the architecture. What we propose is to use multicast at the global (WAN) level on a group-per-exercise basis to minimize WAN traffic, while using multicast at the local (LAN) level to minimize simulator input. This is implemented by having a filtering process at each gateway, shown in Figure 1, that 1) forwards *all* PDUs generated locally to an exercise–level multicast group; and 2) processes all PDUs coming from the rest of the exercise, removing or filtering any PDUs that are not needed locally, and multicasting those remaining PDUs to the local simulators. (For more details on the algorithm, see [PuWh95].)

Previous work[PuWh95a] has used simulation techniques to study the characteristics of this approach; however, the simulations built were concerned with issues of scale, rather than exercise accuracy. In this work, we have taken data (terrain, entity distributions

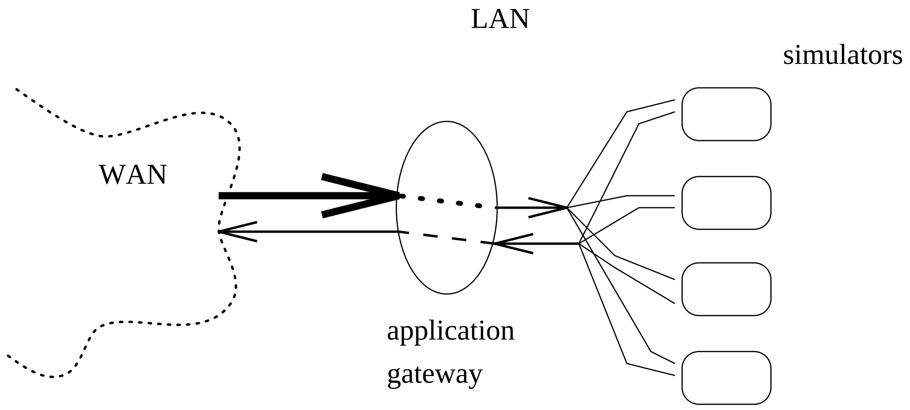

Figure 1: Dual–mode Multicast Application Gateway. All messages generated by the LAN are sent to the WAN. The messages from the WAN are filtered and multicast to the local simulators.

and physical distributions) from a planned DIS exercise and have used this data to get a more accurate picture of what happens in an exercise that uses dual–mode multicast. For the most part, we have assumed worst–case conditions in order to study peak loads on the backbone and at the individual sites.

In Section 2, we describe the characteristics of these simulations. Section 3 gives the results of our simulations. Finally, Section 4 addresses what we have learned about this technique and where we plan to go from here.

2. SIMULATION OF REAL–WORLD ENTITIES

To analyze this approach further we have constructed simulations of a DIS military exercise from the STOW program consisting of 12 sites with approximately 15,500 entities[Meng95]. The exercise terrain covers a 2200KM x 1500KM area. Several simulations were run with either a 44x30 or a 22x15 grid imposed on the terrain. Each grid cell represents the region of interest for the entities within a given cell. We only simulated the entity–state PDUs generated by this exercise. Omitting other PDU types such as emission and collision PDUs reduces the realism in the simulations; however, as entity state PDUs typically constitute 85 to 90% of the PDUs[HCNF94] in the system, we believe the numbers are convincing. The PDUs are generated by the entities at the rate of 2 PDU/second for an active entity and .2 PDU/second for inactive entities.

The simulations described in this paper differ from the simulations described in [PuWh95a] in several significant ways.

- The simulations use entity types that would be present in a real DIS exercise and the characteristics of these different entity types are simulated (relative speed, groupings and terrain constraints).
- The entities are distributed across physical sites in a manner similar to that of a real DIS exercise.
- Movements of entities can be specified by defining a route.
- All entities are active throughout the simulation. This pessimistic assumption allows us to look at the worst case scenario.

In the next two sections, we describe the differences in simulation entities and in their physical and logical distributions. In Section 2.3, we describe some of the important issues that remain to be addressed.

2.1 SIMULATION ENTITIES

In our previous simulations (described in [PuWh95a]) we modeled different size DIS exercises in terms of generic entities. All of the entities had the same probability of being active and the same probability of moving between adjacent grid squares. A biased random movement pattern was assumed, meaning that entities were more likely to move into a grid square that was already occupied.

In our current simulations, we are using entity types and distributions that are planned for use in a real DIS exercise, including aircraft, ground forces, and naval vessels. Although we are still using a biased random movement algorithm for some of the ground and naval forces, we have added the ability for an entity to move along a route. This allows us to more accurately depict some defined movements of a cluster of entities.

One complexity of using real entity types in a simulation is the need to accurately model their behavior. Aircraft move across the terrain at a rate significantly faster than ground forces. To model this behavior, faster entities have a larger probability of movement between grid squares than do slower entities. Figure 2 shows the entities, with their relative speeds and other associated attributes. The relative speed values given in the table were gathered from similar simulations using the Modular Semi-Automated Forces (ModSAF) tool [Loral95].

Some terrain characteristics are also modeled. Each grid cell is provided an attribute that determines if it is land or water. Likewise, each entity, as described next, is characterized by the terrain it is able to travel upon. While aircraft can travel anywhere on the grid, ground entities can only travel between "land" grid squares and sea entities can only travel between "water" grid squares. To do this, we had to refine the movement algorithms to use entity type and terrain information when deciding where to move an entity.

In our initial simulations, we assumed entities moved in groups, or clusters. When designing this new simulation, we still made this assumption; however, we used information on military doctrine to decide what size the groups should be for different entity types.

Entity	Relative Speed	Cluster	Terrain
Infantry	1	28	land
Tanks	10	15	land
Trucks	10	15	land
F15	200	3	air
AH64	30	15	air
Ship	20	2	surface
Submarine	20	1	subsurface

Figure 2: Simulation Entity Attributes

2.2 ENTITY DISTRIBUTIONS

The physical distribution of the entities in this exercise is similar to that of a DIS exercise. The sites labeled A–F in Figure 3 correspond to different geographic locations. The figure shows with the number and types of entities at each site.

The initial layout on the exercise terrain and the logical distribution of entities across the exercise terrain is also similar to that of a real exercise. Figure 4 gives a rough picture of what the exercise layout is. The labels correspond to the sites in Figure 3. For the purposes of our simulation, the aircraft are assumed to be distributed evenly across the associated airbases and the ground forces are distributed with the combat vehicles at the front lines and the auxiliary assests at the rear. The sea forces are distributed randomly over the large area of water (shaded in the diagram).

2.3 ISSUES NOT ADDRESSED

There are many important factors that we have not addressed with this simulation. One is the issue of other PDU types in the simulation; in particular tactical communications which is the communication between the entities of the simulation. This communication can be a significant overhead in a realistic exercise[Meng95]. However, we do not have accurate data on which to base simulation of inter–entity communication.

We also did not simulate the interaction of entities via weapons such as missiles. In real DIS simulations, the firing of a missile causes a burst of PDU traffic. Simulating this realistically would have entailed extending the simulation with the ability to add and remove entities dynamically.

Another issue is that of entity realistic movement and interaction. In a real exercise, the behavior of one entity can cause behavior of another entity (i.e. a plane taking off can cause planes from the opposing side to be launched). In a real exercise, entities may leave the simulation as they become disabled. Simulating these behaviors is beyond the scope of this work. To study how this type of action/reaction impacts PDU rates in the system, we will have to look at real exercise data.

3. SIMULATION RESULTS

The simulations were run in two sets. The first imposed a grid of 44x30, or 50KM square grid cell. The second used a grid of 22x15, or a 100KM square cell. For each grid size, we ran a simulation of 36,000 and 72,000 seconds (10 and 20 hours).

We define the filter rate at an application gateway to be the fraction of PDUs coming from the WAN that are discarded by the gateway because they are not of interest to any of the local sites. In these simulations, the filter rate is observed to be dependent on the amount of interaction of entities at one site with entities at other simulation sites rather than to the size of the site. In Figure 5, sites B–F have lower filter rates at the application gateway. This can be explained by the fact that entities simulated by sites B–F are initially located in close proximity to each other on the exercise terrain and remain this way during the simulation.

The remaining sites, mostly aircraft and naval vessels, have mean filter rates well above 90%. These sites spent most of the simulation not in direct contact with other sites' entities. For example the surface and subsurface ships are constrained to the water cells, eliminating them from interaction with the land cells since no tactical communication or firing PDUs were simulated. The routes specified for these entities do increase their interaction with other sea entities which yields a small percent of filtered messages.

As was found in [PuWh95a], the selection of a grid

Site Name	Entity Types					Total Entities
	Tracks	Wheels	Indiv	Ships	Aircraft	
A	300	300	196			796
B	105	60	140			305
C	60	210				270
D	1845	4425	3360		270	9900
E		15			30	45
F	840	1800	1120		435	4195
G				6		6
H				50		50
I				3		3
J					33	33
K					66	66
L					15	15
Total	3150	6810	4816	59	849	15684

Figure 3: Distribution of entities across simulation sites.

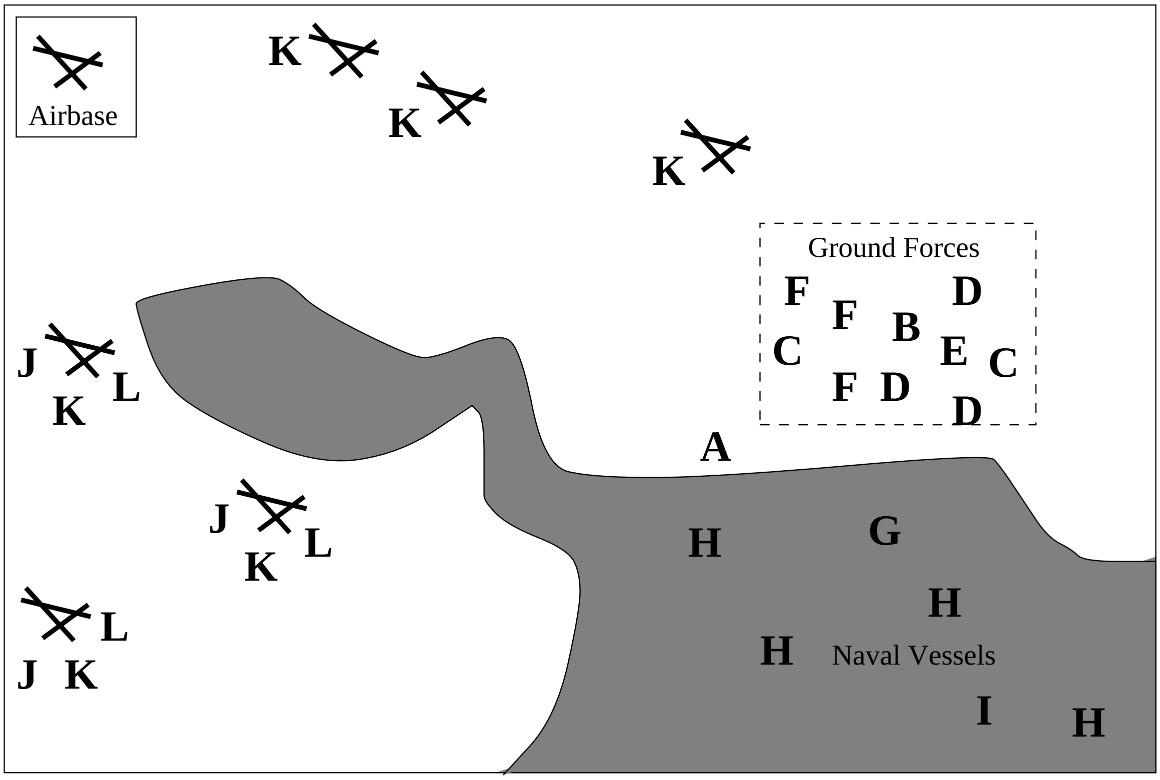

Figure 4: Distribution of entities across exercise terrain. The labels correspond to sites from Figure 3.

size is important, as it will determine how much information is perceived by an entity, but also how much irrelevant data can be filtered. Figure 5 shows the comparison of the 20 hour runs with both grid sizes. As the grid cell size increased, the PDU filter rates at the sites tended to decrease. This result is intuitive since increasing the grid size means enlarging the region of interest.

Of more interest are the actual PDU rates on the backbone and at the individual sites. Figure 6 shows the maximum PDU (MBit) per second rates for the incoming traffic on each of the sites. Because these numbers result from worst–case scenarios, they show what could happen to the system during peak activity. The site loads varied greatly, with sites C and D having the largest loads. Some of these loads are in excess of the capacity typically provided on a LAN; clearly other techniques must be used as well if acceptable loads are to be achieved. We have not included a compression factor in the numbers reported here, thus our numbers represent a bound on network activity rather than what would be expected in a finished system. Simulation results in [HCNF94] show a factor of 10 reduction in packet rate when a combination of grid-filtering (similar in spirit to our technique), compression, bundling, overload management and fidelity management are used.

Traffic on the backbone itself is an important issue. The simulations consisted of more than 15,500 active entities, resulting in more than 58 MBits/sec traffic during the exercise, assuming no compression. To assess how much of this traffic was really required to be on the backbone, we ran a final set of simulations where a given site only sent a PDU onto the backbone if some other site would require this information. This "global knowledge" scheme did result in lower overall traffic on the backbone. For 20 hour simulations using grid sizes 50KM and 100KM, we recorded peak PDU/second rates of 15,842 (30.2 MBits/sec) and 26,440 (50.4 MBits/sec) respectively. Details for sites are summarized in Figure 7. Such a global knowledge scheme could be implemented with a small number of messages between the application gateways.

These simulations were based on data from a single projected DIS exercise. We expect that the general pattern would be similar for other simulations with different entity distributions. However, it will be necessary to simulate more real–world exercises to confirm this.

Site	50KM		100KM	
Name	PDU/sec	MBit/sec	PDU/sec	MBit/sec
A	278	0.5	256	0.5
B	3140	6.0	11910	22.7
C	5290	10.1	21930	41.8
D	7770	14.8	8040	15.3
E	3720	7.1	3810	7.3
F	2614	5.0	6656	12.7
G	42	0.1	98	0.2
H	934	1.8	1440	2.8
I	26	< 0.1	1196	2.4
J	1668	3.2	1666	3.2
K	1580	3.0	1656	3.2
L	90	0.2	390	0.7

Figure 6: Maximum Incoming PDU and MBit per second rates on LANs at sites during the 20 hour simulations for the 50KM and 100KM grid squares.

4. CONCLUSION

Dual-mode multicast is a scalable approach to DIS communication that combines the benefits of grid filtering with simplicity at the global (WAN) level of the architecture. Dual-mode multicast uses exercise-level multicast groups at the global level to minimize WAN traffic, while using geographic multicast groups at the local (LAN) level to minimize simulator input.

In this work, we have considered simulation with more detail in terms of terrain description, entity distributions, and physical distributions, as well as realistic movements, groupings, and speeds or entities from a real exercise and have used this input to get a more accurate picture of what happens in an exercise that uses these techniques. Using this newly built model, it can be seen to what degree filtering irrelevant messages can relieve LAN traffic. The model also reinforces the importance of selecting the proper grid size. However, these techniques alone will not reduce the traffic to a manageable level and must be used in conjunction with other reduction techniques.

Additional simulation work includes addressing the impact on PDU rates of dynamically adding and removing entities over the course of the simulation. This would allow us to study more complex entity interactions. There is also an expectation that tactical communications have a significant impact on DIS performance and should be studied via simulation techniques. Finally, we would like to show the feasibility of using different size grids at different application gateways to allow the level of filtering to be tuned to the entities at each of the backbone sites. This would allow sites that have entities requiring large regions of interest (such as aircraft or sea vessels) to use a coarse grained grids, while sites with entities requiring small regions of interest (infantry for example) could use fine grained grids.

The data flows that resulted in this simulation highlight the difficulties that can be expected in a large DIS exercise. While few exercises are likely to demand the

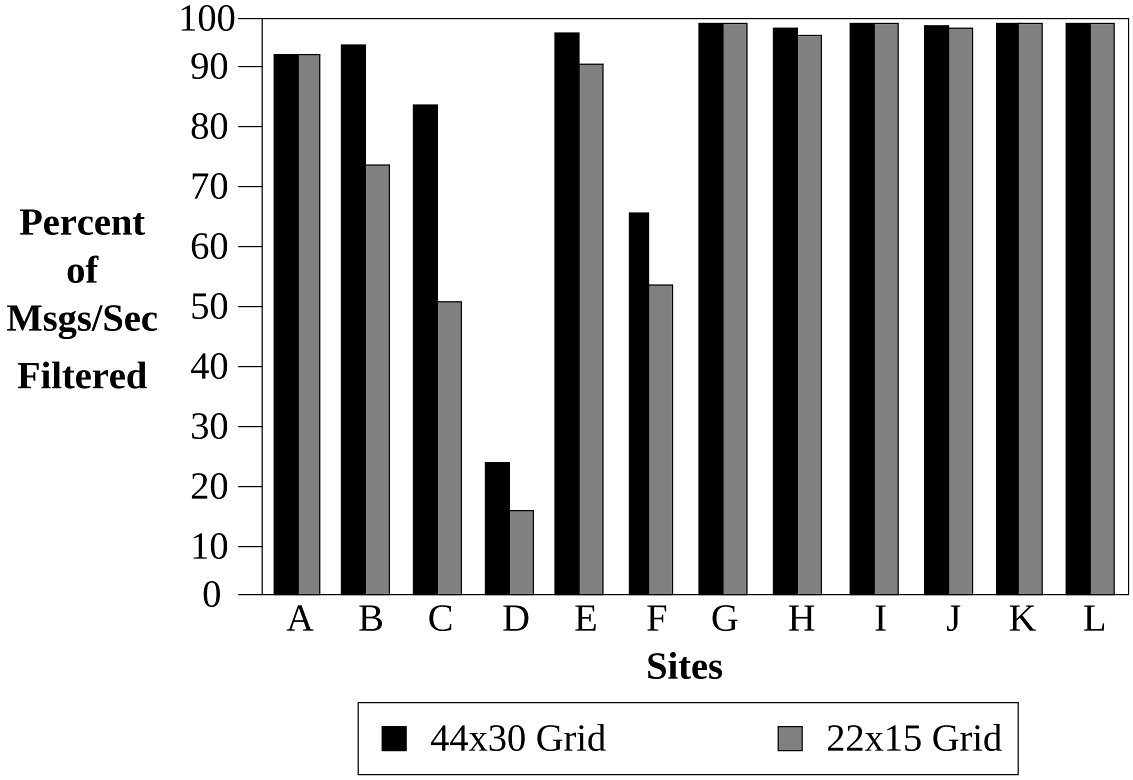

Figure 5: Comparison of average filter rates at simulation sites for different grid sizes.

Site Name	No Global Knowledge		Global Knowledge 50KM		Global Knowledge 100KM	
	PDU/sec	MBit/sec	PDU/sec	MBit/sec	PDU/sec	MBit/sec
A	1592	3.0	1592	3.1	1592	3.1
B	610	1.2	610	1.2	610	1.2
C	540	1.0	330	0.6	540	1.0
D	19800	37.8	7970	15.2	18240	34.8
E	90	0.2	90	0.2	90	0.2
F	8390	16.0	6800	13.0	6800	13.0
G	12	< 0.1	8	< 0.1	10	< 0.1
H	100	0.2	52	0.1	92	0.2
I	6	< 0.1	4	< 0.1	4	< 0.1
J	66	0.1	66	0.1	66	0.1
K	132	0.3	42	0.1	78	0.2
L	30	0.1	30	0.1	30	0.1

Figure 7: Maximum Outgoing PDU and MBit per second rates for 20 hour simulations assuming 1) No global knowledge and 2) global knowledge. When no global knowledge is used, the rates are not influenced by the size of the grid squares.

number of battlefield entities seen in the STOW data, it is clear that much work remains to develop practical, affordable network capabilities for large DIS configurations.

References

[CaHo94] Calvin, J. and D. Hook, AGENTS: An Architectural Construct to Support Distributed Simulation, *11th DIS Workshop on Standards for the Interoperability of Distributed Simulations*, September 1994, pp. 357-362.

[DIS94] Dis Steering Committee, The DIS Vision: A Map to the Future of Distributed Simulation, May 1994.

[HCNF94] Hook, D., J. Calvin, M. Newton, and D. Fusco, An Approach to DIS Scalability, *11th DIS Workshop on Standards for the Interoperability of Distributed Simulations*, September 1994, pp. 347-355.

[Loral95] Loral Advanced Distributed Simulation, User's Manual for ModSAF: Version 2.0, September 1995.

[Meng95] Mengel, Larry L., Personal communication. September 1995.

[Pull94] Pullen, J. M., Networking for Distributed Virtual Simulation. *Proceedings of INET '94*, paper 243.

[PuWh95] Pullen, J. M. and E. White, Dual-Mode Multicast for DIS, *12th DIS Workshop on Standards for the Interoperability of Distributed Simulations*, March 1995.

[PuWh95a] Pullen, J. M. and E. White, Analysis of Dual–Mode Multicast for Large Scale DIS Exercises, *13th DIS Workshop on Standards for the Interoperability of Distributed Simulations*, September 1995.

[Schw90] Schwetman, Herb. CSim Users' Guide (Revision 14), Microelectronics and Computer Technology Corporation Technical Report ACA-ST-257-87, 1990.