ANALYSIS OF DUAL–MODE MULTICAST FOR LARGE SCALE DIS EXERCISES

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

keywords: Communications Architecture, Multicast, Scalability

ABSTRACT

We are currently investigating a scalable approach to DIS communication architecture called dual-mode multicast that combines the benefits of implementing grid filtering using multicast technology 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 grid-based multicast groups at the local (LAN) level to minimize simulator output.

We have constructed simulations that allow us to study the network level characteristics of this approach on both small (500 entities) and large (50,000 entities) simulations. These simulations have given us insights into the computational and communications demands of this technique, including the rate at which incoming PDUs are filtered at the application gateways, the PDU traffic expected on the individual LANs, and the effect of the grid size used for the multicast groupings. These simulations are currently being extended to allow us to consider how this approach can be used to allow sites connected at lower data rates to participate fully in DIS simulations.

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 the DIS world, simulations support various entities that interact by causing and perceiving events or entity–state PDUs in the overall system. Figure 1 shows what we envision a large exercise looking like. Backbone sites are connected via a wide area network. These backbone sites support smaller sites where the simulators themselves are located.

Initially, the DIS simulations used a broadcast approach to PDU transmission. However, as the number of entities being simulated in DIS exercises has grown, this approach has placed unacceptable demands on both the WAN and the individual simulators. As result, scalable approaches to reducing WAN 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 with the increasing presence of wide-area and global viewers associated with exercises. These factors motivate recent approaches to handling issues of scale in DIS environments by implementing a number of information flow management techniques in application gateways[1] (AG) at the LAN/WAN boundary of each participating network. In terms of scalable architectures, the information management technique of greatest interest currently is geographic filtering based on multicast technology.

Geographic filtering is a type of relevance filtering designed 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.

The standard multicast approach to implementing geographic filtering requires support for a large number of multicast groups at the WAN level. The DSI currently supports only a small number of static multicast groups,

[1] Or in processes associated with the gateways[CaHo94].

Figure 1: Large Scale DIS exercise.

reducing the effectiveness of using standard multicast techniques to implement geographic filtering schemes. An additional issue is managing the assignment of individual simulators to multicast groups. LAN adapters typically support at most 16 groups with hardware filtering; software filtering in the simulators would reduce the benefits of multicasting. Therefore a multicast manager would be needed at each site, to monitor all traffic and coordinate assignment of simulators to groups. Managing the large number of multicast groups at the WAN level is likely to prove difficult.

We are currently investigating an approach called dual-mode multicast 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.

In dual-mode multicast, we can have an agent per LAN called a MultiCast Group Agent (MCGA) to map between the global multicast layer and the local multicast. Each MCGA is associated with a particular exercise and has knowledge of the overall geographic-based scheme and the entities associated with simulators on its LAN. The MCGA sends all messages generated by its local simulators to the WAN. It also receives all messages from the WAN associated with a particular exercise and discards any messages that arrive and are not needed by any local simulators. One MCGA at a backbone site may server several LANs there, plus regional sites connected by lower capacity circuits, such as T1. (For more details on the algorithm, see [PuWh95].)

In addition to the advantage of simplicity at the global level, an advantage of the dual–mode multicast approach is that, unlike single-mode multicast approaches, it enables the system to bundle messages on a per-exercise basis. Although bundling is not advantageous in some network technologies (ATM, for example), overhead can be reduce in this manner in some cases.

To study the feasibility of this approach, we are interested in various issues. Both standard multicast and dual–mode multicast have the same estimated traffic on the LAN, the WAN and at the individual simulators. However, in a single-mode multicast approach, an application gateway only receives and processes PDUs for multicast groups that are needed on its LAN. In contrast, the MCGA for dual–mode multicast receives all PDUs sent to the WAN, many of which will not be forwarded to the MCGA's LAN. Section 2 describes the results of our initial simulations that allow us to further quantify the amount of PDU filtering that is done by the MCGA in dual–mode multicast.

We were also interested in addressing the issue of grid size. Intuitively, the smaller the individual regions, the better the filtering process will work. Section 3 describes the simulations that support this intuition and describe how smaller grid sizes can be used to support small simulation sites with low connectivity.

2. INITIAL SIMULATION RESULTS

To give us insights into this technique, we turned to simulation techniques. For the purpose of the initial simulations, we assumed a 20 X 20 grid was imposed on a 100 KM x 100 KM terrain. We assumed standard PDU generation rates for the entities: 2 PDU/sec for an active entity and a lower 0.2 PDU/sec "heartbeat" rate for inactive entities. We also assumed that entities stayed in clusters of size 50 and that these clusters moved between regions at a rate less than once a minute. The three simulations we built focused on looking at different DIS simulation sizes: 500, 5000, and 50,000 entities. For each simulation the entities are distributed across some given number of sites in local simulators. The general characteristics of each size simulation are shown in Figure 2.

We used the CSim[Schw90] tool to build the simulations. Each simulation consists of two CSim processes per site. The first process for a given site was responsible for

- generating event messages at the proper rates given the overall activity rate in the simulation;
- updating MCGA data structures to keep track of what regions the local simulators are interested in; and
- broadcasting the event messages to all sites in the simulation.

To simulate an actual exercise in progress, we needed to have the entities move around the terrain. Our first simulations were run using a random movement function; however, we noticed that this tended to spread out the entities over the entire grid in what seemed to be an unnatural way for an exercise. We are currently using a biased random algorithm. If an entity needs to move, it will choose (randomly) an adjacent square where some other entities are located or have an interest. If there are no adjacent square with this property, the entity moves randomly. This function tends to keep the entities in groups, while permitting interesting and varying interactions.

The second process per site receives incoming messages from the wide-area-network, first checking to see if the message originated locally, and then checking to see if the event was of interest to any local simulations. For

Number of Entities	500	5000	50,000
Activity rate	100%	50%	20%
Overall PDU rates (PDU/sec)	1000	5500	28,000
Total Sites	2	10	50
Distribution of entities across sites (#entities x #sites)	1 x 100 1 x 400	6 x 100 2 x 200 2 x 2000	20 x 100 20 x 400 10 x 4000

Figure 2: General information about the characteristics of the three simulations.

Site	Percent of Messages/second filtered									
Size	0–10	10–20	20–30	30–40	40–50	50–60	60–70	70–80	80–90	90–100
100	0.0	0.0	0.7	3.4	8.7	15.6	21.2	21.0	18.1	11.4
400	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	79.0	21.0

Figure 3: Distribution of filter rates for various sites in simulation of 500 entities.

these simulations we assumed low latency WAN technology that does not insert delay beyond speed-of-light transmission time.

The parameter of interest to us was the amount of filtering done by the MCGA using dual–mode multicast. The higher the filter rate for incoming PDUs, the lower the multicast traffic on the individual LANs. When computing the amount of filtering done by a site in the simulations, we included both PDUs discarded because they were generated locally and PDUs discarded because they were not of interest to any of the local simulators. Computing whether a given PDU is of interest to local entity was accomplished by keeping information about the location of the local entities (extracted from outgoing simulation PDUs) and deciding if each PDU was an update for an entity located either in the same grid square as some local entity or in one of the eight grid squares adjacent to some local entity. During the simulations, we recorded the percent of incoming messages filtered during every simulation second. This allows us to compute not only the average filter rates, but the distribution of these rates as well.

First consider a simulation consisting of 500 entities distributed across 2 sites of sizes 100 and 400, respectively. In this case, we have assumed that all entities are active at any given point in time. Figure 3 shows the distribution of the filter rates for the two sites for a simulation of duration 50,000 seconds. The small site, A, has a wide variability in the filter rates. However, between 60 and 90 percent of the incoming PDUs are filtered more than half of the time (21.2% + 21.0% + 18.1%). These filter rates depend upon how many entities associated with site B are located in an area of interest to entities of site A. Site B has less variability in filter rates due to the fact that it is generating 80% of the PDUs in the simulation and these are discarded by the MCGA because they originated locally.

The simulation of 5000 entities consisted of 10 sites: 6 sites with 100 entities, 2 sites with 200 entities, and 2 sites with 2000 entities. For this simulation, we assume that approximately half of the entities are active at any given point in time. Figure 4 shows the distribution of the filter rates for each of the sites (ordered by site size). This distribution is interesting in that we can see that small sites tended to have fairly high filter rates, greater than 80% 85-90 percent of the time during simulation. The filter rate for smaller sites is lower; typically between 40 and 60% of the PDUs were being filtered. At first glance, this seems to be the inverse of the results of the 500 entity simulation; however, it makes sense in light of two facts:

1. given more entities at a site, there are more regions of interest to that site, and
2. the number of PDUs produced at a large site is a smaller percentage of the PDUs overall.

The simulation of 50,000 entities was of the most interest to us because of the importance of scalability issues. This simulation consisted of 50 sites: 20 of size 100, 20 of size 400 and 10 of size 4000. Figure 5 gives the distribution of the filter rates for each of the sites (ordered by site size). This simulation clearly highlights the fact that the degree of filtering done by a MCGA depends upon several properties of the simulation site including the number of entities and different sites in contact with entities of this site. The filter rates for both 100 and 400 entity sites varied dramatically; the

Site Size	Percent of Messages/second filtered 0–10	10–20	20–30	30–40	40–50	50–60	60–70	70–80	80–90	90–100
100	0.0	0.0	0.0	0.0	0.0	0.1	0.5	10.8	54.5	34.0
	0.0	0.0	0.0	0.0	0.0	0.1	0.5	11.0	55.0	33.4
	0.0	0.0	0.0	0.0	0.0	0.1	0.6	12.5	56.4	30.4
	0.0	0.0	0.0	0.0	0.0	0.0	0.4	9.2	55.1	35.2
	0.0	0.0	0.0	0.0	0.0	0.0	0.6	11.3	54.5	33.5
	0.0	0.0	0.0	0.0	0.0	0.0	0.4	11.0	55.2	33.3
Average	0.0	0.0	0.0	0.0	0.0	0.1	0.5	11.0	55.1	33.3
200	0.0	0.0	0.0	0.0	0.1	1.9	18.4	48.7	29.1	1.9
	0.0	0.0	0.0	0.0	0.1	1.4	15.3	45.1	34.5	3.5
Average	0.0	0.0	0.0	0.0	0.1	1.7	16.9	46.9	31.8	2.7
2000	0.0	0.0	0.1	10.9	62.5	25.2	1.3	0.0	0.0	0.0
	0.0	0.0	0.1	10.6	61.7	26.1	0.0	0.0	0.0	0.0
Average	0.0	0.0	0.1	10.8	62.1	25.7	0.7	0.0	0.0	0.0

Figure 4: Distribution of filter rates for various sites in simulation of 5,000 entities.

sites that stayed in close contact with many other entities had lower filter rates than those that were more isolated. Large sites consistently had low filter rates due to the increased number of regions of interest associated with those sites.

3. GRID SIZE SIMULATIONS

The size of the grid to be imposed on the terrain strongly influences the filter rate at a particular site. Intuitively, the smaller the individual regions, the better the filtering process will work. Figure 6 demonstrates this intuition using two different grid sizes. Some entity (given by the filled circle) has a defined sensor range, shown in both (A) and (B) by a dashed circle. There are two other entities close by (shown as non-filled circles); however only one of these entities is in sensor range. If the grid size is as in (A), the simulator for this primary entity will receive PDUs for both entities since they are both located in grid squares adjacent to that of the primary entity. The information about the further entity would eventually be discarded by the local simulator. By making the grid squares smaller (closer in area to that of the sensor range), the further entity moves out of an adjacent grid square and therefore, the primary entity will not receive update PDUs from it.

To test this intuition and to see whether smaller grid squares would enable site with lower connectivity to interact with a large DIS simulation, we took the 50,000 entity simulation described in the previous section and changed the size of the grid imposed on the terrain. The original exercise used a 20 X 20 grid on top of a 100km X 100km terrain. For this new simulation, we used a 20 X 20 grid on the same size terrain and used the same initial entity positions.

Figure 7 shows the average filter rates for these two simulations. For each site size, the largest and smallest average filter rates define the range. For the 100 entity sites, the 20 x 20 grid simulation has average filter rates ranging from 83.3 to 92.6, while the 40 x 40 grid has average filter rates ranging from 92.9 to 96.8. This improvement in average filter rate, also seen in 400 and 4000 entity sites, is substantial.

In addition to recording filter rates, we extended the data collected to include the amount of PDU traffic on each LAN each simulation second. The average filter rate results agree with the difference in the average PDU traffic sent to the LANs from the MCGA, shown in Figure 8. For each site size, the largest and smallest average PDU/sec rates define the range. For the 100 entity sites, we saw 4151.0 to 9706.0 average PDU/sec for a 20 x 20 grid and 1590.3 to 3961.1 average PDU/sec for a 40 x 40 grid. The figure also shows the largest PDU/sec rates for each site size during the simulation. These numbers are important because they represent the overall required capacity of the LANs.

Obviously, the choice of grid size in a geographic filtering technique is important. A grid square must be large enough to allow entities to get information about all other entities in sensor range, but small enough to filter out as much extraneous information as possible. Different entity types may have different sensory requirements and the choice of grid size would have to be driven by the requirements of the entity types that needs the most information.

One interesting property of dual-mode multicast is

Site	Percent of Messages/second filtered									
Size	0–10	10–20	20–30	30–40	40–50	50–60	60–70	70–80	80–90	90–100
100	57.0	42.4	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	60.1	39.2	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	57.9	41.6	0.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0
	59.4	40.1	0.4	0.1	0.0	0.0	0.0	0.0	0.0	0.0
	0.0	0.1	4.6	46.3	46.2	2.7	0.0	0.0	0.0	0.0
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	3.1	89.4	7.3
	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.6	64.1	35.0
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.4	40.3	59.2
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	39.6	60.3
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	35.7	63.9
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	34.4	65.6
	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	31.2	68.8
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	42.9	56.8
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.3	38.3	61.3
	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.3	51.1	48.6
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.3	27.4	72.1
	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2	33.7	66.0
	0.0	0.0	0.0	0.0	0.0	0.1	0.1	0.3	58.2	41.3
	0.0	0.0	0.0	0.0	0.1	0.1	0.1	0.4	65.5	33.8
	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.2	62.2	37.3
Average	11.7	8.2	0.3	2.3	2.3	0.2	0.1	0.3	35.7	37.3
400	0.0	0.0	0.0	0.2	0.1	0.9	37.7	57.2	4.0	0.0
	0.0	0.0	0.1	0.1	0.2	0.6	38.3	58.2	2.6	0.0
	0.0	0.0	0.0	0.1	0.1	0.8	38.0	57.2	3.7	0.0
	0.0	0.0	0.1	0.1	0.1	3.1	62.1	34.1	0.4	0.0
	0.0	0.0	0.0	0.1	0.1	1.7	51.6	44.7	1.8	0.0
	0.0	0.0	0.1	0.1	0.1	0.2	20.6	70.2	8.6	0.0
	0.0	0.0	0.1	0.1	0.1	0.3	17.1	72.1	10.2	0.0
	0.0	0.1	0.1	0.1	5.6	56.6	36.1	1.3	0.0	0.0
	3.1	50.4	24.3	11.6	5.4	4.8	0.3	0.0	0.0	0.0
	25.2	49.2	7.3	4.7	5.5	5.3	2.9	0.0	0.0	0.0
	1.1	43.0	25.7	11.0	3.9	10.8	4.3	0.2	0.0	0.0
	0.0	17.1	55.9	11.5	81.	6.1	1.3	0.0	0.0	0.0
	18.5	41.6	15.2	7.9	6.7	5.1	4.9	0.1	0.0	0.0
	15.0	45.5	18.4	12.9	4.0	3.3	1.0	0.0	0.0	0.0
	46.9	25.7	10.1	5.4	6.0	2.5	3.3	0.1	0.0	0.0
	0.3	42.2	27.6	10.7	8.5	7.4	3.2	0.1	0.0	0.0
	4.1	46.3	20.4	7.8	9.6	9.5	2.3	0.0	0.0	0.0
	37.5	25.2	13.8	10.0	4.4	5.5	3.4	0.1	0.0	0.0
	48.3	20.2	10.9	6.3	7.0	4.9	2.3	0.0	0.0	0.0
	36.3	30.8	10.1	9.6	6.1	3.1	3.4	0.6	0.0	0.0
Average	11.8	21.9	12.0	5.5	4.1	6.6	16.7	19.8	1.6	0.0
4000	58.5	39.9	1.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	43.4	54.5	2.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	60.0	39.2	0.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	43.6	54.5	1.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	49.0	49.5	1.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	58.5	40.0	1.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	42.9	55.2	1.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	57.3	41.3	1.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	49.0	49.8	1.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	49.5	48.2	2.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Average	51.2	47.2	1.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Figure 5: Distribution of filter rates in simulation of 50,000 entities.

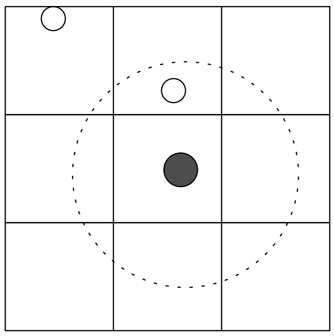

(A)

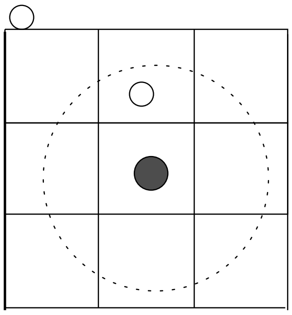

(B)

Figure 6: Smaller grid sizes may eliminate unusable PDUs.

that it allows each MCGAs to use a different grid imposed on the terrain. This is possible because there is only one group at the global level and each MCGA does not need information about how other MCGAs map the incoming messages to multicast groups. This choice of grid size could be driven by the needs of the entities with the greatest information needs at that individual site (as opposed to that of the entire exercise).

4. CONCLUSION

We are currently investigating a scalable approach to DIS communication architecture called dual-mode multicast that combines the benefits of implementing grid filtering using multicast technology 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 grid-based multicast groups at the local (LAN) level to minimize simulator output.

In this paper we describe simulations of small and large scale exercises designed to give us insights into the computational and communications demands of dual-mode multicast. Using simulation techniques, we have demonstrated the rate at which different size sizes filter incoming PDUs. We have also shown how the choice of grid on top of the terrain influences the filter rates and PDU/sec traffic on the individual LANs. This is of particular interest to small sites with low connectivity.

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.

[HoRC] Hook, D., S. Rak, J. Calvin, Approaches to Relevance Filtering *11th DIS Workshop on Standards for the Interoperability of Distributed Simulations*, September 1994, pp. 367-369.

[MZPB94] Macedonia, M., M. Zyda, D. Pratt and P. Barham, Exploiting Reality with Multicast Groups: A Network Architecture for Large Scale Virtual Environments, *11th DIS Workshop on Standards for the Interoperability of Distributed Simulations*, September 1994, pp. 503-510.

[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.

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

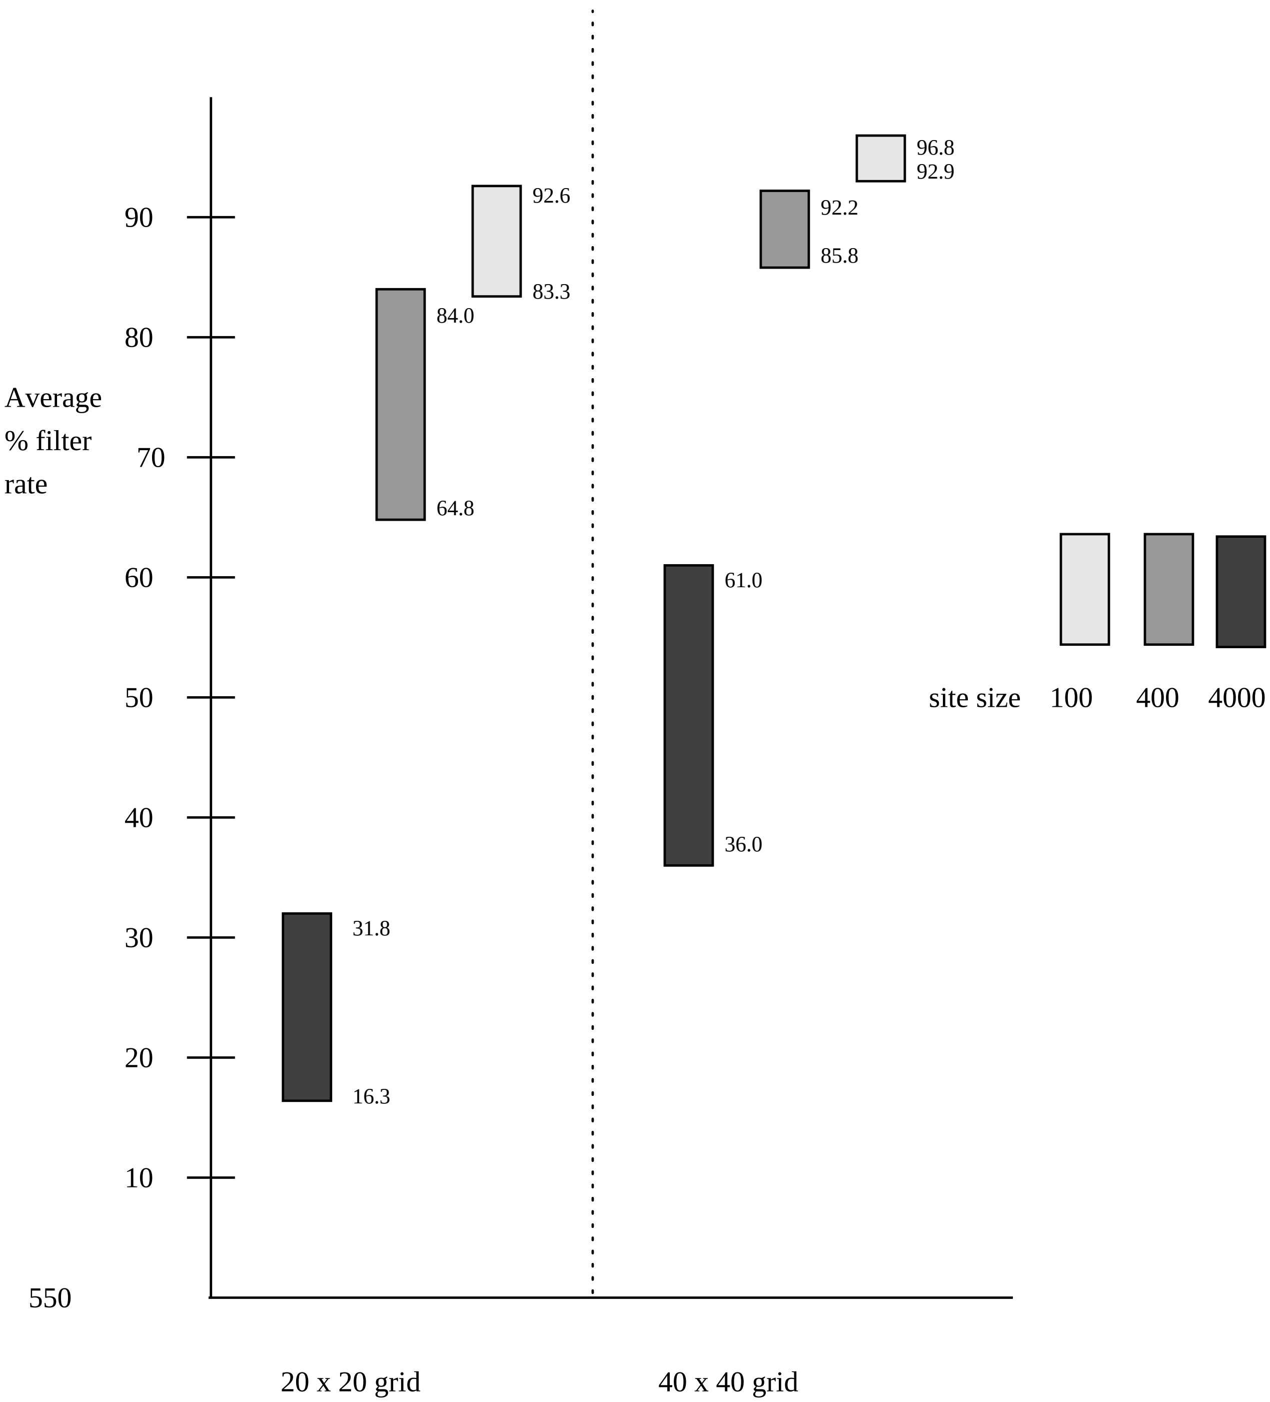

Figure 7: Average filter rate ranges for simulations using 20 x 20 grid and 40 x 40 grid on 100km x 100km terrain (shown by site size).

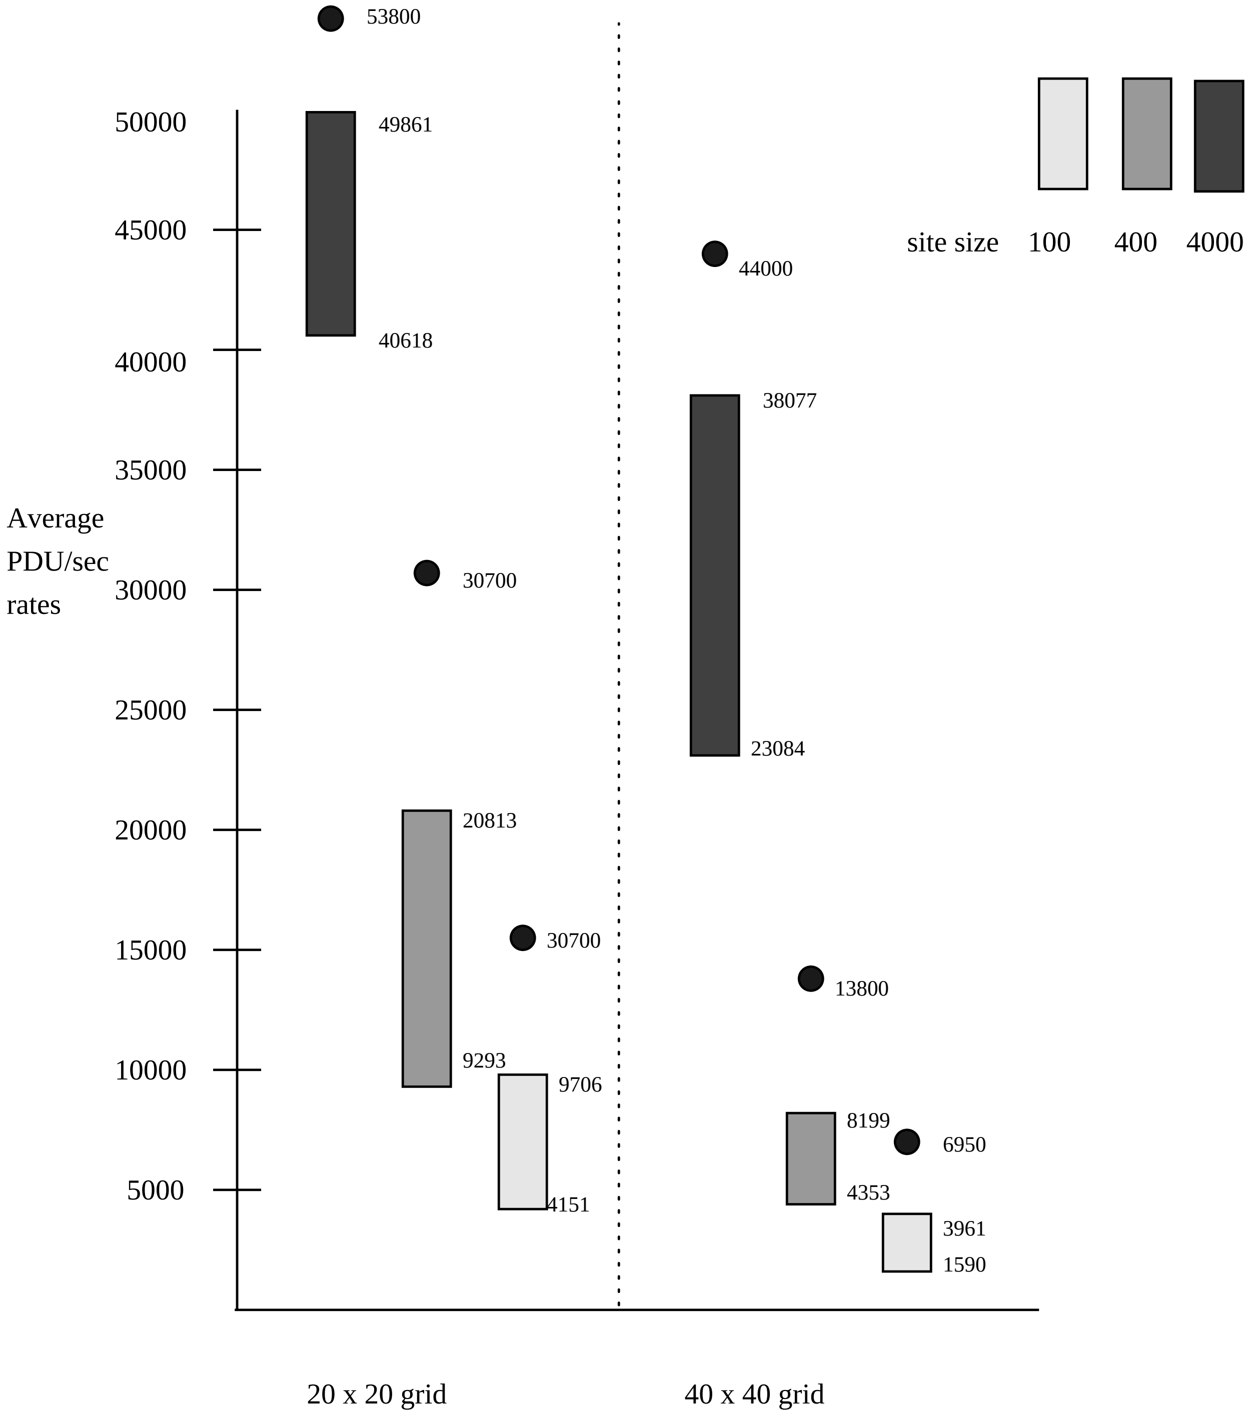

Figure 8: Average and Maximum PDU/sec Rates for simulations using 20 x 20 grid and 40 x 40 grid on 100km x 100km terrain (shown by site size).