

ENERGY CONSUMPTION OF HLA DATA DISTRIBUTION MANAGEMENT APPROACHES

SaBra A. Neal
Richard M. Fujimoto

School of Computational Science and Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

ABSTRACT

Energy and power aware data distribution methods are essential for using these approaches in energy constrained devices and environments. Data Distribution Management (DDM) is a set of services defined in the High Level Architecture (HLA) that aims to efficiently propagate distributed simulation state information. This paper describes an empirical study of the energy consumption of computation and communication components of region based and grid based DDM approaches in mobile devices. Experimental data illustrate that region based approaches tend to consume more energy than grid based approaches for computations, but less for communications. These results also show that the choice of grid cell size and grid cell constraints on publication regions can play an important role in the energy efficiency of grid based approaches.

1 INTRODUCTION

Data distribution management (DDM) is a set of services defined in the High Level Architecture to distribute information in distributed simulation environments (IEEE 2010). DDM services are implemented by Run-Time Infrastructure (RTI) software. Several different approaches to implementing the DDM services have been proposed including grid-based implementations, region-based implementations, and hybrid approaches that utilize a combination of ideas from the grid and region-based approaches (Tan et al. 2000). These approaches have certain computation and communication requirements that are necessary to perform DDM operations. To our knowledge no work has been conducted to date examining DDM services from the standpoint of energy consumption. This is the focus of the work described here. We examine the well-known region based and grid based approach along with two other DDM approaches to gain an understanding of the energy consumption properties of using DDM in energy constrained environments.

Dynamic Data Driven Application Systems (DDDAS) are applications that continuously monitor, analyze, and adapt operational systems in order to better assess and/or optimize their behavior (Darema 2004). Applications arise in many areas such as natural disaster management, transportation and manufacturing, and others (Chen et al. 2012; Long 2017; Mandel 2005). Many DDDAS applications involve sensing and computation on mobile devices utilizing communications through wireless networks. Energy consumption in these applications is a major concern because battery life often limits the effectiveness of DDDAS applications utilizing mobile platforms.

With the growing use of mobile devices research in the area of mobile computing has increased. Mobile devices have the ability to provide real-time information that can be used as input to real-time applications such as DDDAS. Sensors such as GPS, cameras, accelerometers, and environmental sensors are becoming more widely deployed. The dependence on battery power in mobile devices makes it vital to understand the energy consumption properties of running applications on mobile

systems. We are concerned with understanding how to implement and use DDM to provide real-time information to run embedded DDDAS applications within mobile devices. These DDDAS applications are used to make predictions and analyze systems such as traffic networks.

The organization of this paper is as follows: Section 2 reviews DDM concepts and implementation approaches. Section 3 outlines details of the energy tradeoffs between the region based, distributed region based, grid based, and grid filtered region based approaches. Section 4 describes the scenario used for the basis of our experiments to evaluate energy consumption. Section 5 describes the experimental setup used here. Section 6 presents measurement results of energy consumption of DDM computation and communication components. Section 7 summarizes the results, presents conclusions, and briefly discusses areas of future work.

2 DATA DISTRIBUTION MANAGEMENT APPROACHES

Data Distribution Management services are used to reduce traffic flow over the network. DDM services are defined in the High Level Architecture Interface Specification. The services are implemented by Run-Time Infrastructure software. DDM utilizes an N-dimensional coordinate system called a routing space to represent, for example, a geographical area. Federates express their interest by defining subscription regions that characterize the information they are interested in receiving. Each message is associated with a publication region to characterize the content of the message. If an overlap is detected between a message publication region and a federate's subscription region, the message is sent to that subscribing federate. The region based and grid based DDM approaches are the most well known approaches to implementing DDM services. Several other approaches have been proposed and implemented to overcome the drawbacks of these approaches.

2.1 Region Based Approach

The region based approach manages interests by performing a matching computation between all publication and subscription regions defined within the routing space. This approach incurs a $O(N^2)$ computation cost where N is the number of subscription / publication regions. Although costly this approach is efficient in communicating messages. A multicast group is defined for each publication region within the routing space. Once an overlap is detected between a publication region and a subscription region, the subscribing federate joins the multicast group associated with that publication region. If a subscription region changes, the new subscription region must be compared against each other publication region to determine which groups it should leave or join. Similarly, if a publication region is changed, it must be compared against every other subscription region to determine the new composition of its multicast group. This approach is sometimes called the brute force approach because it is a direct implementation of the DDM services.




Figure 1: Region Based Approach

An example illustrating this approach is shown in Figure 1. Two multicast groups are created, one for P1 and one for P2. Federates subscribing the region S2 join the group for P1 and those subscribing to S1 subscribe to the group for P2. Messages associated with publication region P1 are thus routed to federates subscribed using region S2.

2.2 Grid Based Approach

The grid based approach implements DDM by partitioning the routing space into fixed size grid cells. A multicast group is defined for each grid cell within the grid structure. The grid cells are used to determine overlap between publication regions and subscription regions within the routing space. Federates determine those grid cells for which their publication regions overlap and send each message to the associated multicast group(s). Similarly, federates subscribe to groups associated with cells that overlap with their subscription regions. The grid based approach is considered to be more computationally efficient and scalable than the region based approach.

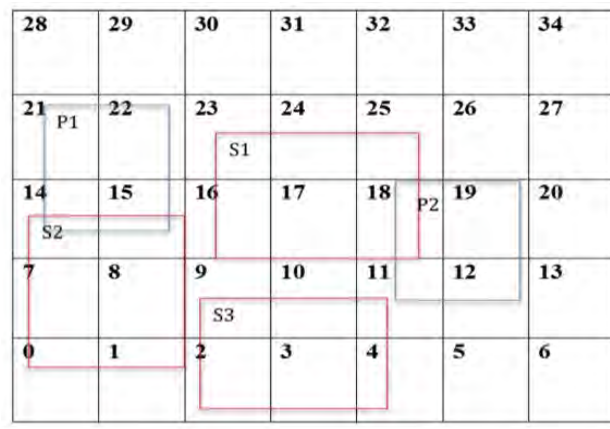


Figure 2: Grid Based Approach

The drawback to the grid based approach is it lacks communication efficiency. Subscribing federates can receive irrelevant and duplicate messages under this approach. The example in Figure 2 illustrates the grid based approach. A message sent with publication region P1 is sent to groups 21, 22, 14, and 15. Two of these messages are sent to federates subscribed to region S2, resulting in duplicate messages. Similarly, messages sent with publication region P2 are sent to federates subscribed to S3 through multicast group 11, even though P2 and S3 do not overlap.

2.3 Other Approaches

To overcome the communication inefficiency of the fixed grid based approach and the computation inefficiency of the region based approach, other approaches have been proposed. The hybrid based approach reduces the communication cost of the grid based approach and the matching cost of the region based approach by performing direct matching between publication and subscription regions within grid cells (Tan et al. 2000). Boukerche and Roy proposed the dynamic grid based approach whose goal is to reduce the number of multicast groups created under the grid based approach (Boukerche and A. Roy 2002). Under this approach multicast groups are only allocated to a cell if an intersection is detected between a publication and subscription region. The grid filtered region based approach was also proposed (Boukerche et al. 2005). This approach combines the grid based and the hybrid based approach by using a threshold parameter to determine when matching needs to be performed between regions within a grid cell whose area of coverage of the grid cell falls below the specified threshold.

The effect of grid cell size on multicast group assignment and communication cost has been explored in (Rak and D. VanHook 1996). Their findings concluded that irrelevant and duplicate messages are increased with the use of larger grid cells. Many other approaches and mechanisms have been created to overcome the drawbacks of purely region based and fixed grid based approaches. The goal of adaptive data distribution management is to optimize DDM time, matching, and multicast group assignment (Raczy et al. 2002). Raczy also proposed the sort based DDM matching algorithm. This algorithm sorts the coordinates along each dimension. An extension of the sort based algorithm was also created to promote efficiency in large spatial environments (Pan et al. 2007). Pan also proposed the dynamic sort-based approach. This approach efficiently matches selective region modifications (Pan et al. 2011). The parallel matching algorithm uses the divide and conquer technique among intersections of regions and creates region bound sub-lists to speed up the matching process. This approach reduces the computational overhead of region matching (Liu et al. 2015). Marzolla and D'Angelo proposed a sort based matching approach for data distribution management on shared memory multiprocessors that solves the region intersection problem for data distribution management (Morzolla and D'Angelo 2017).

3 ENERGY TRADEOFFS

Tradeoffs arise concerning energy consumption in different data distribution management approaches. Every DDM approach must perform DDM operations such as initial region matching, processing changes to regions, and communicating distributed simulation messages. The difference among DDM approaches concerns how they perform and handle these operations and the effects they have on energy consumption.

3.1 Region Based Approach

The region based approach typically uses a central controller to handle DDM operations. The central controller is responsible for comparing all publication regions against all subscription regions in order to determine which multicast groups each subscribing federates should join. The initial establishment of multicast groups using this approach is computation intensive. The matching computation is $O(N^2)$ where N is the number of publication or subscription regions. Each publication region must be compared against all subscription regions in the routing space to determine which regions overlap. Subsequent changes to publication or subscription regions require computation of complexity of $O(N)$ to compare the new publication (subscription) region with all of the subscription (publication) regions in the routing space. Therefore, the frequency at which a federate changes its subscription regions has a significant impact on energy consumption. Although the region based approach is computationally intensive it is efficient in terms of communications. The region based approach does not create irrelevant or duplicate messages that often occur in DDM grid based approaches, as discussed earlier. Subscribing federates only receive messages for publication regions that overlap with their subscription regions.

3.2 Distributed Region Based Approach

The use of a central controller can limit scalability. To address this issue we propose a distributed region based approach. This approach uses R controllers created by distributing the responsibility of managing federates among the controllers. Each controller is responsible for a fixed area of the routing space and manages publication and subscription regions that lie within their assigned area. Subscription regions that overlap more than one area require communication among the corresponding controllers to determine the multicast groups a subscribing federate must join and which controllers must be notified when a region changes. This alleviates the scalability issue in the original region based approach. Since direct matching still occurs under this approach the communication of distributed simulation messages using the DDM services is still just as efficient as the original region based approach; federates do not send or receive duplicate or irrelevant messages improving energy consumption relative to the grid-based approaches.

3.3 Fixed Grid Based Approach

The use of a grid structure greatly reduces the computation needed for DDM operations by avoiding the matching computations required by the region-based approaches. But, it leads to irrelevant and duplicate messages, as discussed earlier. The number of irrelevant and duplicate messages depend on the size of the grid cell. Thus, grid cell size plays an important role in energy consumption for grid based DDM approaches. Small grid cells can lead to many duplicate messages. While large grid cells reduce the number of duplicate messages, they can increase the number of irrelevant messages.

3.4 Grid Filtered Region Based Approach

The grid filtered region based approach increases scalability and reduces the number of irrelevant and duplicate messages. Like the grid based approach, this approach utilizes a grid structure to perform interest management operations. This method uses a threshold parameter that indicates the percentage of area a region must cover within the grid cell before it joins the grid cell's multicast group. This threshold triggers matching computations between regions within grid cells. Those publication regions whose area exceeds the threshold of the grid cell are placed in a full coverage list for the cell. The same is done for subscription regions; those exceeding the coverage threshold are placed in a full coverage list for subscription regions. These regions need not perform matching within the cell. Those regions whose area of coverage is below the threshold are placed in partial coverage lists for publishers and subscribers. All publication regions within the partial coverage list are matched against all subscription regions within the partial coverage subscribers list. This reduces the number of irrelevant message by using direct matching computations.

4 SCENARIO

These experiments focus on a vehicle traffic application. The envisioned system is a Dynamic Data Driven Application System (DDDAS) using mobile devices. These devices, e.g., smart phones, utilize traffic volume data from sensors embedded within the arterial traffic network. A two-dimensional routing space is used that corresponds to a traffic network measuring 50 by 50 kilometers. The embedded sensors are placed at each intersection of the traffic network and detect vehicles traveling in the North and South direction. Sensor devices publish data such as measurements of traffic volume; thus, their publication region corresponds to an area surrounding the location of the sensor. Subscription regions represent areas of interest for a mobile vehicle and are typically distant from the location of the vehicle. Specifically, vehicles express interest in receiving data from sensors that are located 805 meters (0.5 miles) away from their current location in their current direction of travel, which is either North or South bound in the arterial traffic network. The embedded sensors have a sensor range of approximately 400 meters. Each moving vehicle is a subscribing federate and the embedded traffic sensors are publishing federates. The arterial road network contains vehicles that are assigned a random position within the arterial road network and travel either north or south, moving with a constant velocity of 20 meters per second.

5 EXPERIMENTAL SETUP

These experiments utilized the LG Nexus 5x cellular phone with a Qualcomm Snapdragon 808 processor, 2GB memory, and 16GB storage as the mobile computing platform. The phone runs the Android version 6.0.1 (Android Marshmallow) operating system. Computation experiments are performed on the Jetson TK1 development boards with ARM A15 32-bit CPU with 4 cores operating at 2.3 GHz and 2 GB memory. Energy and power measurements were performed using a PowerMon2 power measurement system for measurements that were evaluated on the board (Bedard et. al. 2009). The Trepro profiler application was used to perform measurements on the Android phone (Trepro Profiler 2016).

To evaluate the computation energy of the four DDM approaches – region based, distributed region based, fixed grid based, and grid filtered region based – we perform all computations on the Jetson development board. Each approach’s matching component was run on the board and the energy consumption was measured to determine the average amount of energy consumed.

To evaluate the energy needed for communication we assumed that one RTI client communicates with the development board through TCP sockets over a WLAN network. The client for this experiment was the Nexus 5x mobile device that performed all DDM operations. This was implemented by creating an Android application using Android Studio. All communication was performed between the mobile device which acted as the central controller in the case of region based approaches and the publishing federate in the case of grid based approaches. Wireless communication was performed between the mobile device and the board with the device receiving 1000 byte messages and the board sending 1000 byte messages. The mobile device mimics the full operation of what would happen during communication under the traffic scenario described earlier. In the case of the region based approaches the device sends a message to the server indicating its subscription region. The server (board) then performs matching between the client’s region and the known other regions within the routing space to initially set up the multicast groups. The server then communicates with the client to indicate what multicast groups it should join. In the case of the grid based approach the client performs its own interest matching computation by determining which grid cells its subscription region overlaps, and then joins the designated multicast groups. Energy consumed by the client for messages sent and received under each approach was then measured.

All experiments assume that two central controllers are used for the distributed region based approach and the threshold for the grid filtered region based approach is set to 0.6. The grid cell size is set to 20m x 20m, publication regions are 40m x 40m and subscription regions are 80m x 80m.

6 MEASUREMENTS

6.1 Energy for Computations

We define computation energy to be the energy used to compute the matching component of the DDM approaches. The energy used as the number of subscribing federates increase are shown in the Figure 3.

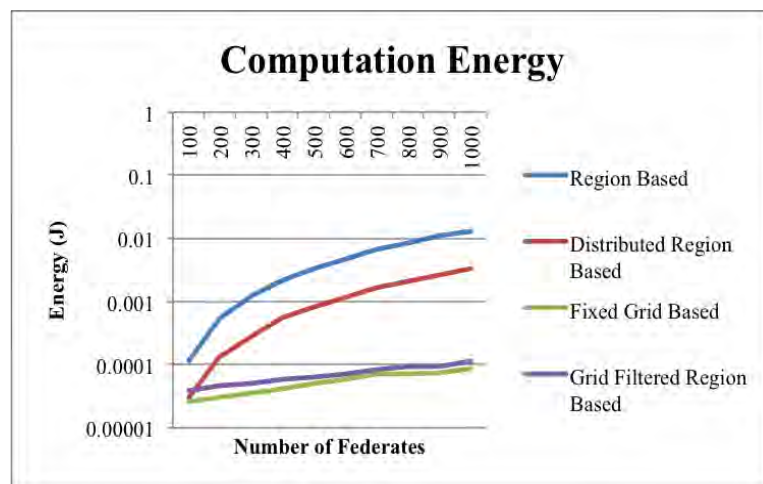


Figure 3: Computation energy as federates increase

The results show the energy cost of using a purely region based DDM approach in comparison to grid based DDM for multiple independent runs. The region based matching cost is $O(N^2)$ to initialize DDM operations since, every publication region is compared against every subscription region.

The distributed region based approach performs matching in the same manner as the region based approach. It designates a subset of regions to each central controller in the routing space based on the coordinates of the region. Each controller performs matching on the regions that are in their domain as well as those that cross multiple domains. Utilizing multiple central controllers allows the number of regions that must be compared to be reduced which in turn reduces the energy consumption cost compared to the centralized region based approach.

The fixed grid based approach consumes the least amount of energy out of all the approaches compared in this study. The matching operation in this approach involves federates determining the grid cells with which their regions overlap. The increase in energy is due to each federate determining the overlapping grid cells and adding themselves to the list of federates in the multicast group.

In the grid filtered region based approach computation energy is similar to that of the fixed grid based approach. Federates again determine those cells with which their regions overlap and only perform matching computations on those cells. Using the threshold parameter matching need only be done on regions in the partial coverage lists.

6.2 Energy for Communications: Update Messages

Update messages are defined to be data messages sent by publishers to subscribers to transmit simulation data. Figure 4 shows the average amount of energy consumed by one RTI client (federate) in receiving messages for each update message sent by a publisher, including irrelevant and duplicate messages. The average energy consumption over ten independent runs is shown.

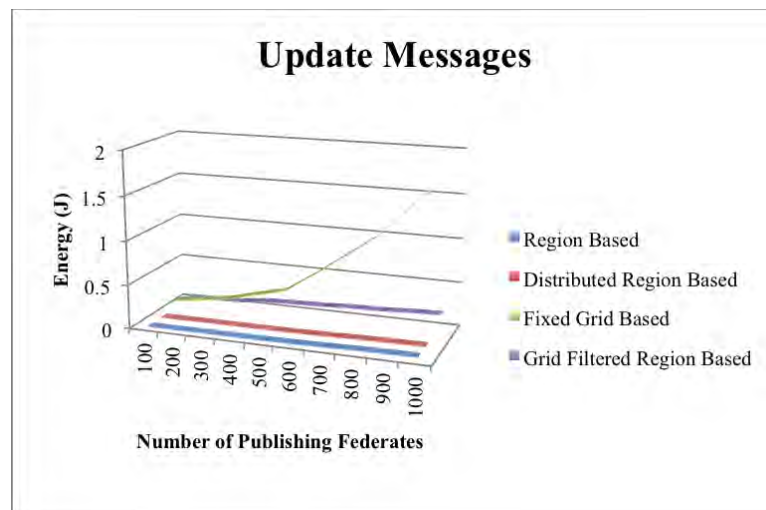


Figure 4: Energy of communicating update messages

The results show the communication tradeoffs that occur when utilizing region based approaches compared to grid based approaches. The fixed grid based approach consumes the most energy among the four approaches while the region based and distributed region based approaches consume the least amount of energy.

The region based and distributed region based approaches consume more energy for matching computation but use less energy for communications. The matching computation ensures that no

irrelevant or duplicate messages occur in the region based approaches. This reduces energy consumption. Sending fewer messages reduces the time needed to send and receive messages.

The grid region based approach consumes less energy than the fixed grid approach but still uses more energy than the region based approaches. The use of the grid structure results in duplicate messages, but adding the threshold parameter eliminates the irrelevant message problem. The use of direct matching among those regions that do not exceed the threshold avoids irrelevant messages. In turn eliminating irrelevant messages leads to a reduction in the number of messages received during an update which leads to a reduction in energy consumption. This experiment suggests that this approach achieves a balance between computation and communication with respect to energy consumption. It consumes little energy for computation and a small amount of energy for communication.

6.3 Energy for Communications: Constraining Publication Regions

Constraining the publication region to one grid cell will allow grid based DDM approaches to gain efficiency in federate interest matching. By constraining the publication regions irrelevant and duplicate messages are eliminated. Under this constraint subscription regions cannot overlap grid cells that partially contain a publication region due to the fact the publication region occupies the entire grid cell. This eliminates the irrelevant message problem. Duplicate messages also will not occur because publication regions only occupy one grid cell so update messages are sent to only one multicast group.

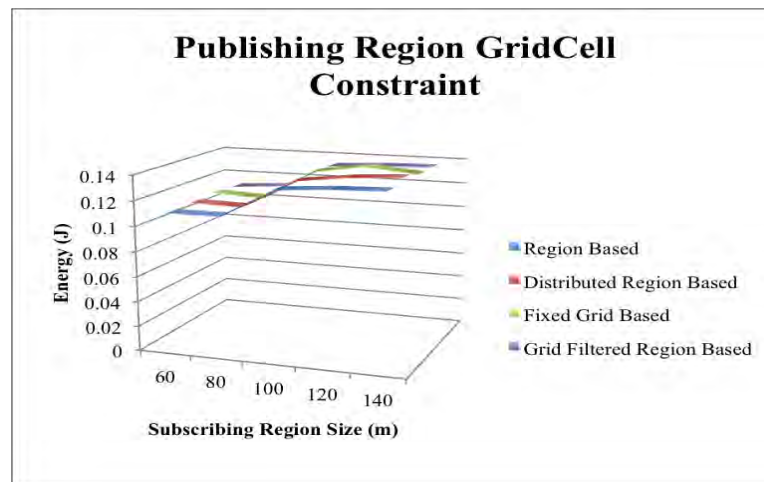


Figure 5: Publication region Constrained to One Grid Cell

Figure 5 shows the energy consumption when constraining publication regions to one grid cell. We see that as we vary the size of the subscription region we increase the number of update messages; but, the effect is the same regardless of using the region based or grid based approach. This result is due to the fact that there are no irrelevant or duplicate messages. Subscription regions under grid based approaches can only overlap any publication region once since it only contains one grid cell. This experiment illustrates that placing constraints on DDM regions enables one to gain the communication efficiency of region based approaches and the computation efficiency of grid based approaches. The region based and grid based approaches consume the same amount of energy. Exploitation of this constraint requires flexibility in the distributed simulation application to impose such constraints, and is not suitable for all applications.

6.4 Communication: Constraining Publication region to N Grid Cells

Varying the number of grid cells that the publication region encompasses allows one to understand the effects of grid cell size on the number of update messages a federate receives under different DDM approaches. We evaluate this effect by varying the number of grid cells covered by a publication region. Figure 6 shows the energy cost averaged over ten independent runs.

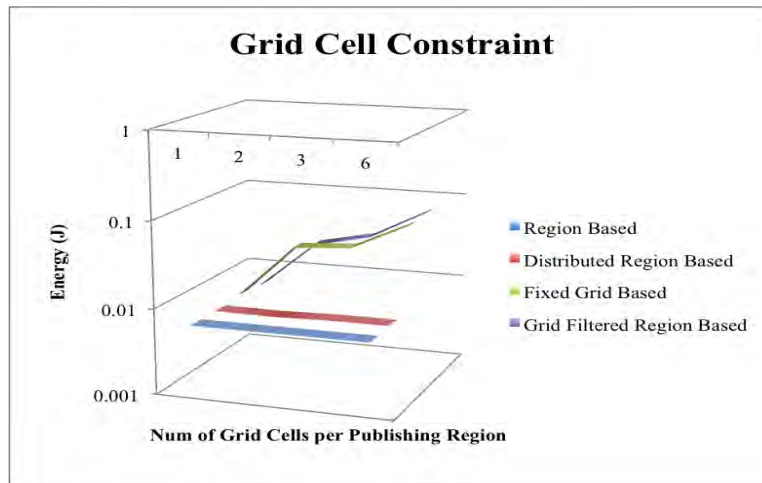


Figure 6: Publication region Constrained to N Grid Cells

Figure 6 shows that as we increase the number of grid cells encompassed by the publication region we increase the number of duplicate messages that a subscribing federate will receive. Since the publication region in this experiment is constrained to the cell boundaries there are no irrelevant messages. This constraint still allows duplicate messages to occur under the grid filtered region based approach because regions will always surpass the threshold. Utilizing this constraint makes grid filtered region based use just as many messages as fixed grid based would because they both will not incur irrelevant messages, but duplicate messages will still occur.

6.5 Communication: Varying Grid Cell Size

Grid based approaches are directly affected by the grid cell size. Smaller grid cells tend to lead to more duplicate messages and larger grid cells tend to lead to more irrelevant messages. In the following experiment we examine the effect that grid cell size has on energy consumption. We examine this effect while leaving the publication and subscription regions a static size for each grid cell size. The average energy consumption of ten independent runs were measured.

Figure 7 shows the impact that grid cell size has on grid based DDM approaches. We see that larger grid cell size greatly impacts the energy consumption of the fixed grid based approach. An interesting phenomenon is the small effect that grid cell size has on the grid filtered region based approach as the grid cell size increases. Here we see that a change in energy consumption for this approach does not occur until the grid cell size is at the largest size evaluated where we see a decrease in energy. This shows that utilizing a threshold of 0.6 to trigger direct matching between regions that fall below this threshold in the grid filtered region based approach gives a trend similar to a region based approach and close to the region based approach energy consumption when large grid cell sizes are used.

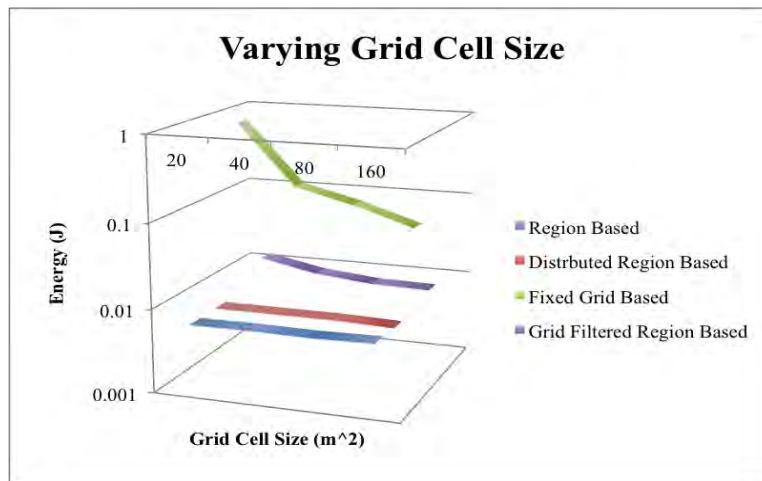


Figure 7: Varying Grid Cell Size

7 CONCLUSION

Using data distribution management in an energy constrained environment requires an understanding of the energy consumed by the elements of the DDM system. We analyze the energy required for computation and communication operations of four DDM approaches. Our experimental results quantify the high amount of energy consumption required in a centralized region based approach. We also conclude that that grid based DDM approaches can be costly in energy consumption when sending update messages due to irrelevant and duplicate messages. Restricting publication regions to one grid cell eliminates irrelevant and duplicate messages, but, utilizing this constraint may not always be suitable for the application. We observe that reducing the grid cell size reduces the number of irrelevant messages, but at the cost of increased duplicate messages. We also conclude that overall grid cell size with no constraints causes an increase in energy as grid cell size increases for grid based approaches.

An area of future work is to develop new DDM approaches created specifically to be used in energy constrained environments. Such an approach might combine the communication energy efficiency of the region based approach with the scalability and computation energy efficiency of grid based approaches.

ACKNOWLEDGMENTS

Funding from this research was provided by NSF/AFOSR Grant 1462503, AFOSR grant FA9550-17-1-022, and DoD SMART fellowship.

REFERENCES

- Bedard, D., A. Porterfield, R. Fowler, and M. Y. Lim. 2009. "PowerMon 2: Fine-Grained, Integrated Power Measurement". *Technical Report TR-09-04*, Renaissance Computing Institute, University of North Carolina, Raleigh, North Carolina.
- Boukerche, A., N. J. McGraw, C. Dzermajko, and K. Lu. 2005. "Grid-Filtered Region-Based Data Distribution Management in Large-Scale Distributed Simulation Systems." *Simulation Symposium, 2005. Proceedings. 38th Annual*, Institute of Electrical and Electronics Engineers, Inc, 259-266.
- Boukerche, A. and A. Roy. 2002. "Dynamic Grid-Based Approach to Data Distribution Management." *Journal of Parallel and Distributed Computing* 62(3):366-392.

- Chen, H., J. Wang, and L. Feng. 2012. "Research on the Dynamic Data-Driven Application System Architecture for Flight Delay Prediction." *Journal of Software* 7(2):263-268.
- Darema, F. 2004. "Dynamic Data Driven Applications Systems: A New Paradigm for Application Simulations and Measurements." *Computational Science-ICCS 2004, Springer*: 662-669.
- IEEE Std 1516.3-2010. 2010. *IEEE Standard for Modeling and Simulation (M&S) High Level Architecture (HLA) -- Interface Specification*. New York, NY: Institute of Electrical and Electronics Engineers, Inc.
- Liu, Y., H. Sun, W. Fan, and T. Xiao. 2015. "A Parallel Matching Algorithm Based on Order Relation for HLA Data Distribution Management." *International Journal of Modeling, Simulation, and Scientific Computing* 6(2):1540002.
- Long, Q. 2017. "A Framework for Data-driven Computational Experiments of Inter-Organizational Collaborations in Supply Chain Networks." *Information Sciences* 399: 43-63.
- Mandel, J., L. Bennethum, M. Chen, J. Coen, C. Douglas, L. Franca, C. Johns, M. Kim, A. Knyazev, and R. Kremens. 2005. "Towards a Dynamic Data Driven Application System for Wildfire Simulation." *Computational Science-International Conference on Computational Science* 197-227.
- Moreno, M., and G. D'Angelo. 2017. "Parallel Sort-Based Matching for Data Distribution Management on Shared-Memory Multiprocessors." In *Proceedings of the 21st IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications*. Rome, Italy: Institute of Electrical and Electronics Engineers, Inc.
- Pan, K., S. J. Turner, W. Cai, and Z. Li. 2007. "An Efficient Sort-Based DDM Matching Algorithm for HLA Applications with a Large Spatial Environment." In *Proceedings of the 21st International Workshop on Principles of Advanced and Distributed Simulation*, 70-82. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Pan, K., S. J. Turner, W. Cai, and Z. Li. 2011. "A Dynamic Sort-Based DDM Matching Algorithm for HLA Applications." *ACM Transactions of Modeling and Computer Simulation* 21(3):1-17.
- Raczy, C., G. Tan, and J. Yu. 2005. "A Sort-Based DDM Matching Algorithm for HLA." *ACM Transactions of Modeling and Computer Simulation* 15(1):14-38.
- Raczy, C., J. Yu, G. Tan, S. Tay and R. Ayani. 2002. "Adaptive Data Distribution Management for HLA RTI." In *Proceedings of 2002 European Simulation Interoperability Workshop*. Harrow, UK:
- Rak, S. J. and D. J. Van Hook. 1996. "Evaluation of Grid-Based Relevance Filtering for Multicast Group Assignment." In *Proceedings of 14th Designing Interactive Systems Workshop*, 739-747.
- Tan, G., Y. Zhang, and R. Ayani. 2000. "A Hybrid Approach to Data Distribution Management." In *Proceedings of the 4th IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications*, 55-61. San Francisco, California: Institute of Electrical and Electronics Engineers, Inc.
- Trepp Profiler. 2016. Trepp Power Profiler. <https://developer.qualcomm.com/software/trepp-power-profiler>. Accessed: April 2017.

AUTHOR BIOGRAPHIES

SABRA NEAL is a PhD student in the School of Computational Science and Engineering at the Georgia Institute of Technology. She received her B.S. in Computer Science from North Carolina Agricultural and Technical State University. Her email address is sneal6@gatech.edu.

RICHARD FUJIMOTO is a Regents' Professor in the School of Computational Science and Engineering at the Georgia Institute of Technology. He received a Ph.D. in Computer Science and Electrical Engineering from the University of California-Berkeley. His email address is fujimoto@cc.gatech.edu.