

Establishing the Limits of Networked Simulation: A Grand Challenge for Modeling and Simulation

Thomas L. Clarke
Institute for Simulation and Training
University of Central Florida
3280 Progress Drive, Orlando, FL 32826
tclarke@ist.ucf.edu (407)882-1327

Michael D. Myjak
The Virtual Workshop, Inc
P. O. Box 98 Titusville, FL 32781
mmyjak@virtualworkshop.com

Keywords: interoperability, DIS, HLA, VV&A, simulation

ABSTRACT

Computer science began with by establishing limits; Turing showed that the halting problems and many others are not computable. A rigorous analysis to establish the limits, if any, of what may be simulated within a network of simulators needs to be done. Conducting such an analysis at a sufficient level of abstraction to cover all possible network topologies and protocols is a grandly challenging problem

Since Turing's time other restrictive aspects of computing have been found and analyzed. These include computational complexity, instability, and chaotic sensitivity to initial conditions, among others. Emerging techniques such as strategic sampling should provide tools to analyze and establish limits for networks where even small differences in systems can mean dramatic changes in throughput. Tools for establishing the limits set by instability and chaos are provided by modern non-linear analysis. These and other techniques, supplemented with new research, will provide a sound scientific underpinning to determining the limits and behaviors of ultra-large networked simulations.

Research under this grand challenge problem will insure the execution scalability of ultra-large scale networked simulations. It will help to avoid hardware/software congestion within simulation networks and it will help keep network designers from becoming trapped in computational cul-de-sacs. This problem thus has significant impact on the economics of large networked simulations.

INTRODUCTION

Thorpus and Skepticus last debated at 1996 I/ITSEC meeting in Orlando. Five years will have passed until they again debate at the First International Conference on Grand Challenges for Modeling and Simulation at the 2002 SCS Western Multiconference. What follows is an update of

their original dialog that takes into account developments of the past five years. It is expected that interactions Thorpus and Skepticus have with GCMS attendees will do much to help refine their ideas.

Participants: Thorpus, Skepticus

Thorpus is a legionary centurion who is the chief advocate of training soldiers with distributed networks of interactive simulators. Under his guidance, a networking protocol known as DIS (Distributed Interactive Simulation) had been developed through a series of DIS workshops that bring together representatives of all major simulation manufacturers. At these workshops problems with the DIS protocol are identified and solutions are suggested. Periodically the workshop then issues a new standard version of DIS that codifies DIS protocol elements that have been accepted by consensus at the workshops.

Since Thorpus' time the DIS workshops have become the SIW (Simulation Interoperability Workshops) and DIS has been absorbed into the High Level Architecture (HLA).

Skepticus is a natural philosopher and computer scientist who is skeptical of the ability to reach the goal of a communications protocol, such as DIS, that can encompass the intercommunication needs of all simulators. While Skepticus has been amazed by the progress provided by Moore's law since his last debate with Skepticus in 1996, he still thinks there are fundamental limits that cannot be breached by sheer computing horsepower.

Thorpus: I have a vision for the future of simulation that has yet to be fully articulated. Thousands perhaps even millions of individual simulators networked together to implement simulations beyond the capability of individual simulators. These simulators will be distributed across both real and virtual spaces but will continue to interact; for short I called this vision DIS, Distributed Interactive Simulation. Others have now taken up the torch of my vision and have produced the High Level Architecture or HLA, that by virtue of its abstraction and generality promises to interconnect not only a large number of simulators, but the full range of types of simulator.

Skepticus: I have serious doubts about this. The work of Turing shows that any universal computer can do anything any computing device whatsoever can. Your network of simulators, each of which I assume contains a universal computer, would be just a large universal parallel computer. As such it would fall under the limitations of Amdahl's law; communication inefficiencies would make the networked simulators less effective than a single simulation computer of equal cost.

Th: You are assuming a naive simple model of inter-simulator communication. I have developed schemes and protocols that greatly reduce the network traffic. Indeed the recent popularity of Beowulf clusters and other network-of-workstation computational schemes show that human ingenuity is capable of barriers such as Amdahl's so-called law.

Sk: For the moment then let us concede that communication overhead can be reduced, but I would like to return to this point later. I foresee other problems with your network. I assume that you will want to have different makes and models of simulators on your network. How will you insure that they can all work together? How do you keep it from becoming a Network of Babel?

Th: That is easily taken care of by the establishment of network protocol standards.

Sk: We will have to discuss these standards later in the context of language theory, but more problems with HLA continue to occur to me. Assuming that both communications overhead and network standards issues can be dealt with, what will be the performance of the network? If your simulators are very far apart, there will be significant propagation delays that will slow your HLA network.

Th: I assure you there are means to deal with the issue of propagation delay. By properly assigning computation tasks to the networked simulators involved in the HLA exercise, the deleterious effects of propagation delays can be eliminated.

Sk: Again you will have to give me the details later, but now I remember that you mentioned virtual space. What does this mean, do you include simulators in space, near the moon? Will this not cause an accuracy problem for your simulators? What about continuous fields, like electromagnetism and sound? What about intelligence?

Th: (sputtering.) Hold on! Let me explain my concept which has grown into HLA from the beginning. I think I can meet all your objections. If I can't, perhaps you can recommend ways that I can improve HLA.

HLA is an outgrowth of DIS which is in turn an outgrowth of an earlier project that I directed, SIMNET. Under my sponsorship, Bacchus, Barthalomew, and Nestor developed the first network of chariot training simulators.

Certain communication problems became evident in developing SIMNET. Communication delays, limited network bandwidth, etc, were problems that had to be overcome.

BBN's brilliant solution was to use dead reckoning. Each simulator maintains a model of every other simulator on the network. In the absence of communications from the other simulators, this model simulator continues to move from the last known position at the last known velocity, so that network bandwidth is not used up constantly sending position information ...

Sk: (Interrupting) Yes! Very clever. Then you would send true position information only every so often to maintain accuracy. Are these true positions sent at fixed time intervals?

Th: In the absence of other reasons, true positions, or updates, are sent at a minimum rate that is set by the protocol. HLA is divided into Federations, and Federations that have real-time requirements typically follow the old DIS ideas. Typically one update or status packet is sent every second so that other simulators are constantly reminded of the sender's existence and do not drop the simulator from consideration. The usual terminology within DIS was to refer to a packet of data a protocol data unit (PDU), now within HLA the OO terminology of message is used.

But to maintain highest accuracy, a given simulator not only runs models of other simulators, but also contains a model of itself. The dead-reckoned position of the self-model is constantly compared with the true position of the simulator. Whenever the self dead-reckoned position and the true position differ by more than a protocol determined threshold, a new update PDU or message is broadcast to all other simulators. Since the self dead-reckoning is in error, it is known that all other dead-reckonings will be in error as well and will require correction.

This scheme for using dead-reckoning models, sometimes referred to as ghosts, also alleviates many problems associated with propagation delays, transmission errors etc. No simulator can remain in error for very long since it is constantly receiving update PDUs/messages from other simulators.

If a simulator joins the networked simulation exercise in progress, within a short time it will have picked up all the other simulators on the network via their minimal stay-alive PDU/message transmissions. If a PDU/message is missed due to transmission error, the dead-reckoning will maintain a good approximation until a good PDU/message is obtained.

Sk: All right, I can see how this scheme with its dead reckoning or ghosts will support a network of training simulators, at least in the first approximation. However, I

see many problems that present severe difficulties for implementing any such network.

Let me explain some of these objections in detail.

A LONG and ONGOING DEBATE

Skepticus and Thorpus continue their debate touching on:

- Communications Overhead
- Network Standards
- General Impossibilities
- Validation
- Limitations On Physical Simulation

Due to limitations on conference preceding paper size, the details are omitted.

CONCLUSION

Th: Well this sounds very good. If I get some physicists laid off from that Super Collider project to work on the HLA team, you make it sound like physical phenomena can be modeled without going outside the HLA message passing dead-reckoning paradigm.

Sk: The success of quantum mechanics makes me think that in principle the HLA paradigm can handle physics, but I have some doubts about the practicality. In the microscopic realm, the world really seems to work by particles exchanging information via virtual particles. For the scale of physics concerned to the battlefield, however, the use of the Feynman formalism is more of a mathematical/computational convenience. The continuous distribution of water vapor etc. that constitutes the cloud or whatever is conceptually Fourier transformed so that the cloud is decomposed into a series of modes. Each mode is then treated as a particle which is instantiated as a HLA entity.

What concerns me is the number of entities or modes that would be needed to represent detailed features of the environment like clouds. If every mode requires an entity this could lead to the sorts of problems we discussed under the size. Granularity becomes an issue as well. The phenomena of chaos wherein systems evolving according to non-linear dynamics become extremely sensitivity to initial conditions could place very high demands on the precision of representation. Modern techniques for decomposing into modes such as the wavelet transform can help to reduce network traffic demands, but this is clearly an area that requires much study.

Th: You are thinking a bit in terms of the non-server architecture imposed by the old DIS architecture. HLA now has the run time interface layer (RTI) that can perform some of the functions that a server would have in DIS. So what if server-like entities are allowed?

Sk: If you had asked me that question at the beginning I would have told you that servers are clearly the way to go. A server can incorporate the latest thought on modeling environmental phenomena from the relevant scientific community, and act as the interconnection media between the HLA federates. For the EM case the radar emits to a server that simulates the propagation through the environment including effects such as inversion layers, and the server then in turn communicates with the target. All the nasty details of trying to model the environment strictly within server-less DIS-like paradigm are avoided.

However, now that I see the possibility of building a server-less HLA federation that can model the environment by designing the federation in analogy to the techniques of quantum mechanics and relativity I am not so sure. There is a certain attractiveness to the pure server-less DIS-style, although it loses some of its simplicity if parallel federations with different speeds of influence are required. The big advantage of server-less HLA might be forward compatibility with future versions. If servers are introduced there will be a big discontinuity with past versions so that legacy systems will be orphaned since they will be incompatible with the servers.

That's a decision I will have to leave to the SIW.

REFERENCES

- Clark, Myjak, and Lake (1999) "RTI Interoperability Study Group Final Report" *1999 Fall Simulation Interoperability Workshop*, Vol. 1, pp.1 – 28.
- Clarke, Long, and Prasad (1997) "Non-Intrusive HLA Interoperability Testing" *1997 Fall Simulation Interoperability Workshop*, Vol. 1, pp.203-209.
- Clarke, T. L. (1998) "Simulation's Ultimate Challenge" *IITSEC 1998 Proceedings*.
- Clarke, T.L. (Ed) (1995) "Distributed Interactive Simulation Systems for the Simulation and Training in the Aerospace Environment" *SPIE*, Bellingham, WA. pp.327
- Clarke, T.L. and Myjak, M. (1996) "A Dialog on the Limits of DIS." *IITSEC Conference Proceedings*.
- Schiavone, G.A., Sureshchandran, S., and Hardis, K. (1997) "A Multifaceted Approach to Interoperability Issues in Training with Distributed Interactive Simulation using Dissimilar Simulators" *ACM Transactions on Modeling and Simulation (TOMACS)*, Vol. 7, No. 3, July, 1997, pp.332-36