Strategic Directions in Distributed Simulation

Volume 2 in the *Simulation 2000* Series

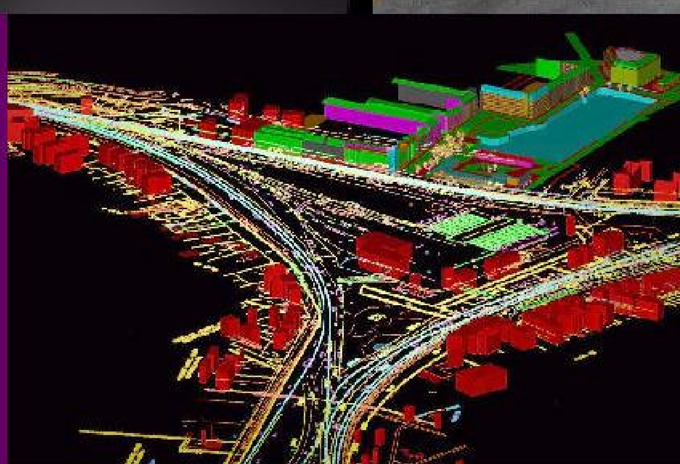

Roger Smith

Title:	**Strategic Directions for Distributed Simulation**
Author:	**Roger D. Smith** Chief Scientist, ModelBenders LLC http://www.modelbenders.com/

Summary:	This paper presents strategic directions and research challenges in distributed simulation. The strategic directions are areas in which simulation can be applied immediately, but where we have not taken full advantage of the technology that is available. These include systems management, real-time decision making, persistent virtual worlds, and virtual verisimilitude. The research challenges are those technologies that are essential for the progress of distributed simulation, but for which we need a deeper understanding before we can proceed. These include modeling human behavior, creating simulation domain architectures, providing abundant information bandwidth, and developing practical event management techniques.
	
	Volume 2 in the ***Simulation 2000*** series.
	
	Intended Audience:
	• News Reporters and Story Researchers
	• High School and College Students
	• Classroom Teachers
	• Visionaries and Creators of Virtual Worlds
	
	Outline of Material:
	• Distributed Simulation
	• Strategic Directions
	  ▪ Systems Management
	  ▪ Real-time Decision Making
	  ▪ Persistent Virtual Worlds
	  ▪ Virtual Verisimilitude
	• Research Challenges
	  ▪ Human Behavior Modeling
	  ▪ Simulation Domain Architectures
	  ▪ Abundant Network Bandwidth
	  ▪ Practical Event Management
	• Conclusion
	• References
	• Acknowledgement

1 DISTRIBUTED SIMULATION

This paper presents strategic directions and research challenges in distributed simulation. In searching for these technologies the author polled several prominent members of the simulation community and reviewed recent publications that characterized many areas of simulation:

- *Proceedings of the 1998 Winter Simulation Conference* (Medeiros, Watson, Carson, & Manivannan, 1998),
- *Proceedings of the Twelfth Workshop on Parallel and Distributed Simulation* (Unger and Ferscha, 1998),
- *Proceedings of the 1999 Game Developers Conference* (Yu, 1999), and
- *Digital Illusions: Entertaining the Future with High Technology* (Dodsworth, 1998).

Distributed simulations are those applications that span multiple computer devices, executables, or geographic areas. These include what are often referred to as parallel and distributed simulations (PADS) and distributed interactive simulations (DIS) (Fujimoto, 1998). These communities vary widely in their techniques for implementing a distributed simulation, but they both fall under the general category of distributed simulation.

Distributed simulation is widely applied in military training systems in which computers and executables have been joined together through techniques like the Distributed Interactive Simulation (DIS) protocol, Aggregate Level Simulation Protocol (ALSP), and the High Level Architecture (HLA). It is also used in analytical models in which networked and parallel computers divide a problem into smaller pieces that can be solved more rapidly. The entertainment community has applied distributed simulation ideas in attractions like the Battle Tech Entertainment Center and the Internet-based Virtual Worlds environment. Most computer games also contain a distributed simulation mode that allows them to interoperate with other people playing the game on the Internet. Games like Quake II, Rainbow Six, Command & Conquer, and the entire Star Wars line are well know and well sold for this capability.

2 STRATEGIC DIRECTIONS

The strategic directions are areas in which simulation can be applied immediately, but where we have not taken full advantage of the technology that is available. These include:

- Systems management,
- Real-time decision making,
- Persistent virtual worlds, and
- Virtual verisimilitude.

2.1 Systems Management

It is possible to embed simulation modes in the operating systems of computer systems. These systems can feed data about their operations into a data store that is accessible to simulation processes. Periodic execution of these would evaluate this performance data and identify the operational trends in the data. This can then be used to optimize the system for its most characteristic applications.

The PC is a general-purpose computer that is put to specific tasks once it is in the hands of the user. If the operating system contained a simulation kernel it would be able to evaluate the uses to which each machine was being put and optimize that machine for those applications. The simulation would need a database of application characteristics such as word processing, accounting, databases, graphic art, sound editing, games, web serving, web surfing, telephone management, and hundreds of others.

The advantage of simulation-based adaptation is that the user need not be an expert in configuring the machine and the simulation can re-optimize the machine when it is applied to a different function. Since most systems are used for more than one application, the simulation would also be able to adjust the configuration to best satisfy two or three applications - a task beyond the abilities of most PC users.

2.2 Real-time Decision Making

The world is filled with opportunities to apply computer simulations to assist in real-time decision making. Any place that information is available in a digital form and humans are evaluating that information to making decisions based on that information, there is an opportunity to support the human with a simulation.

These opportunities occur in thousands of fields, only a few of which will be described here.

Combat Consultant. Large military organizations are migrating their communication and decision-making tools to computer systems. This provides combat information in a form that can be accessed and evaluated by a simulation. We are lucky to live in an age in which our citizens are not faced with life-and-death combat decisions on a daily basis. As a result, soldiers that encounter this kind of event are relatively inexperienced at dealing with it. Military organizations mitigate this through extensive training activities (some of which also involve simulations), but there is no substitute for experience. A Combat Consultant is a simulation mode embedded inside of the command, control, communications, computers, and intelligence systems being used by the soldiers (these are commonly referred to as C4I Systems). The simulation is equipped with the expertise of previous commanders and the real-time expertise of other commanders currently using similar systems on the network. The Combat Consultant can monitor the information in the system and suggest alternatives that may be successful under the current situation.

Though this may begin as an expert system, it also includes the real-time experience of other commanders solving similar problems at this moment in time and a simulation engine to project this situation into the future. The system searches for the best strategies for handling each combat situation in real time.

The term C4I evolved from C2 over the last two decades to more accurately describe the operations performed by commanders and their decision support systems. It is time to add simulation to the acronym - C4IS.

Aircraft Navigation. The Federal Aviation Administration is planning to change the mechanism for controlling commercial air traffic across the country. Under the new method, entitled "Free Flight", pilots will have unprecedented decision making authority in selecting their flight paths and adjusting them throughout the trip. Simulations can assist these pilots by evaluating environmental data, aircraft status, data received from sensors, and data transmitted from the ground. The simulation can constantly study the current situation, looking for the optimum solution for reaching a destination. Perhaps more importantly, the simulation can also generate customized plans for use in an emergency. When the unexpected happens, a plan is ready and the flight consultant is there to support a pilot who is confused, scared, and unable to make decisions.

Crowd Management. All large cities face the problem of managing the flow of people trying to accomplish their own objectives. These people may be driving in rush hour traffic, searching the mall for a sale, or rushing to the best rides in a theme park. In all of these cases, we could optimize operations by directing this traffic. Using traffic flow sensors we can measure the location and density of people in the system, feed this information to a simulation, and look for solutions in real time.

In the case of the theme park, entertainment events could be scheduled by the simulation in patterns that push and pull the guests in specific directions. Good theme parks are designed to direct the flow of traffic from the time you enter the main gates until you finish your tour of the attractions. These designs would be assisted in real-time by simulations that recognize overcrowding in one area and schedule activities to pull part of the crowd to another area. The "pull" mechanism may be the appearance of a costumed character down a side path; beginning a computerized entertainer directly behind the accumulating mass; or the sounds of a roaring dinosaur in a different direction. These tactics are designed to redirect the crowd in a manner that is non-intrusive and that appears to be of the guest's own volition. Events may also be scheduled to direct the guest's attention away from the fact that they are waiting in line.

Market Prediction. Banks and financial institutions are already using simulation and gaming techniques to analyze past performance and predict future activities. These simulations influence commodity trades, stock speculation, and currency exchanges. They provide an edge over competitors that can result in millions of dollars in additional profits. Simulations of this type can be embedded into many forms of stock selection and advice software, including those used by your stock broker, internet stock trading web page, and personal asset management package (e.g. MS Excel, Quicken, MS Money). These are also useful tools for teaching a novice how markets work and what to watch for in future investments.

2.3 Persistent Virtual Worlds

The networked world is a natural host for a persistent virtual world that is accessible to all users. We need to create virtual environments that are persistent over many years and that form the foundation for specific studies, training, and entertainment that will be conducted within them. The gaming community has already accomplished this with online persistent virtual worlds like Ultima Online, Diablo, and Everquest. These provide persistent fantasy worlds that evolve as the users interact with them.

Similar virtual worlds need to be created by high level sponsors of studies and training events. These would be the seeds from which scenarios are drawn and the environment in which distributed interactions occur. Organization like the Office of the Secretary of Defense, the Defense Modeling and Simulation Office, the Central Intelligence Agency, the National Air and Space Administration, and others need to become the hosts for persistent virtual worlds that support the needs of their entire customer base. It should be possible for globally distributed customers to enter these worlds at any time and explore solutions to current problems.

Commercial versions of this can be used to track the activities of specific individuals in the population. The popularization of cell phones and pagers has placed electronic tracking devices on the belts of a demographic of people that we are most interested in tracking. These tagged people can serve as a sample of the general population, allowing us to see customer density in airports, malls, highways, and large entertainment events.

This could be used to identify potential witnesses to crimes based on their presence in the area and predictions of the path they were likely to have followed while in the area. It may even be possible to identify the perpetrator of the crime using this technique.

2.4 Virtual Verisimilitude

In the simulation business we strive to create virtual worlds that accurately represent the real world. This always involves a high degree of abstraction to help us experiment with systems that are far too detailed to fit into any model. However, we have been so conditioned by our lack of computational power and seduced by our skills at abstraction that we sometimes avoid extending our simulations when we have the tools to do so.

There are few simulations that portray a really convincing virtual world. With all of the computational power now available and the increasing maturity of software tools to build models and virtual worlds, we need to explore a new level of representation. It is time for the next big advance in modeling detail and the richness of virtual environments

Statistically accurate simulations are excellent for many applications, but we need to begin equipping ourselves with models that accurately represent individual objects, events, and interactions without relying on actuarial effects to make them correct.

3 RESEARCH CHALLENGES

The research challenges are those technologies that are essential for the progress of the field of distributed simulation. Though there are many areas of valuable research, the four listed here are broad enough and essential enough to be listed as research challenges. These are:

- Human behavior modeling,
- Simulation domain architectures,
- Abundant network bandwidth, and
- Practical event management techniques.

3.1 Human Behavior Modeling

Many simulations are driven by statistical distributions that characterize the average behavior of a system, but do not claim accuracy for individual events or small time intervals. These distributions represent the activities of machinery, the population growth rates of animals, and human performance of specific tasks. However, they do not model instantaneous behaviors of intelligent or reactive beings in the virtual world. We are in dire need of techniques for inserting intelligent, reactive, unique human behavior in the virtual world.

Military training simulations and computer games require interactions between human operators and automated virtual humans. In the past, this has been accomplished through techniques like Finite State Machines that encode specific behaviors and define the transition conditions from one behavior to another. However, we are discovering the limitations imposed by this technique. These are very difficult systems to create and maintain. Human operators that interact with them regularly identify their limitations and take advantage of them. The entities controlled by these techniques do not exhibit realistic behaviors, rather they exhibit correct behavior - "by-the-book", robotic actions.

We need to discover and create techniques for representing the behavior of human leaders, followers, and groups that give them the ability to appear "live" or "real" to the humans interacting with them. Both the military and the gaming communities are augmenting their robotic methods with "softer" models that include human emotion, training, and fatigue. These result in objects that are all slightly different in spite of being driven by the same software.

The distributed simulation community needs a set of behavior libraries that can be linked into a simulation in the same way we currently link in statistical distributions. This will require the definition of a set of categories of behavior and API's that are necessary to stimulate those categories.

3.2 Domain Architectures

Within the Department of Defense we have been developing standard protocols for joining multiple, previously independent, simulations. These methods have included the

DIS protocol, the Generic Data System (GDS), ALSP, and most recently, the HLA. With HLA we have begun to identify simulation functionality that is generally necessary for all systems and which should be included in an infrastructure rather than within specific simulation models. This approach allows a simulation development team to reuse some of the essential functionality that is included in the general infrastructure. It has also encouraged us to question the uniqueness of every simulation system. We recognize that simulations fall into domain areas in which the degree of commonality is much higher than it is across all simulation systems. We begin to imagine a layered view of simulation uniqueness similar to network protocol layers. Higher layers become more specific until they narrow to a single application.

It should be possible to develop an architecture that supports an entire domain of simulation systems, providing a large common pool of functionality. These architectures may include a general interoperability standard like the HLA, but would go further by defining a set of domain tools for operating the simulations, common interfaces to connect to external systems, and object base-classes from which to extend unique object instances.

3.3 Information Bandwidth

Distributed simulations can not exist without sufficient reliable communications bandwidth for delivering events and synchronizing execution of the entire system. This bandwidth is currently one of the limiting factors on the size of a distributed simulation. Luckily, bandwidth is also a limiting factor for all applications using the Internet. This has attracted millions of commercial research and development dollars to the problem. That work can and will be applied directly to simulation applications. The global communications industry will discover methods for providing abundant information transfer. These will include methods for configuring the physical medium of delivery and efficient protocols for transferring data. We may productively put our efforts into simulation-specific communications protocols that are not addressed by other communities.

3.4 PDES Management

For twenty years we have been involved in research to discover techniques for practical and efficient synchronization of distributed simulation processes. This has resulted in some very clever and powerful ideas for addressing this problem. However, these ideas have been embraced by few industrial and government applications. The constructive wargaming community has adopted Conservative Time Management, but Optimistic Time Management is still searching for an ideal application.

We must identify applications that are well served by the different methods of event management. To justify further study, our research in this area needs to find a practical and valuable home in commercial, government, or military simulation systems. By 2010 we should be able to apply the appropriate synchronization technique to a distributed simulation by analyzing the problem, setting configuration variables, and attaching the event management engine to our simulation. Trial-and-error and fine-tuning of the engine must become standardized such that a simulation professional can perform these operations, rather than a PDES specialist.

4 CONCLUSION

The strategic directions and research challenges presented in this paper emphasize two different aspects of the future of distributed simulation. The first is the need for additional development and imagination in applying the technologies we already have. The second is the need for additional research and innovation in areas that will allow us to advance the state-of-the-art. . It is the author's opinion that, while the research challenges provide stimulating problems, the strategic directions for applications are much more urgent at this time. The world is in the middle of an information, communication, and computational explosion. Thousands of advanced applications are being fielded every year and many of these could be improved through the inclusion of existing simulation technologies. However, these opportunities are being lost or the technology reinvented by others because of the lack of communication, marketing, and proselytization by members of the "core" simulation community.

REFERENCES

Dodsworth, Clark. 1998. *Digital Illusion: Entertaining the Future with High Technology*. ACM Press. New York, NY.

Ferren, Bran. 1999. "Some Brief Observations on the Future of Army Simulation". *Army RD&A Magazine*. May-June 1999.

Fujimoto, Richard M. 1998. "Parallel and Distributed Simulation". *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*. John Wiley & Sons Inc. New York, NY.

Medeiros, Watson, Carson, & Manivannan. Editors. 1998. *1998 Winter Simulation Conference Proceedings*. Washington D.C.

Unger, Brian & Ferscha, Alois. Editors. 1998. *Proceedings of the Twelfth Workshop on Parallel and Distributed Simulation*. Banff, Alberta.

Yu, Alan. Editor. 1999. *Proceedings of the 1999 Game Developers Conference*. San Jose, California.

ACKNOWLEDGEMENT

The author would like to thank Dr. Tony Valle, Chief Architect of the Joint Simulation System, and Dr. David Pratt, Chief Scientist for SAIC ASSET, for their support in identifying the Research Challenges listed in this paper.