LABORATORY FOR COMPUTER SCIENCE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Neutral Agents

A Novel Mechanism for Interactive Simulation

Computation Structures Group Memo 380
April 2, 1996

Arvind
Alejandro Caro

To be presented by Arvind as the keynote address in the session on High Performance Computing at the 20th Army Science Conference, 24-27 June 1996, Norfolk, Virginia.

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Neutral Agents
A Novel Mechanism for Interactive Simulation

Arvind and Alejandro Caro
Laboratory for Computer Science
MIT
{arvind,acaro}@abp.lcs.mit.edu

Abstract

Interactive simulation has the potential to effect fundamental improvements in the Army's war-fighting capabilities. Nevertheless, problems with *scalability*, *fidelity*, and *security* preclude current simulation architectures from satisfying future Army requirements. We introduce the concept of *Neutral Agents* as a novel mechanism to address some of these issues. A simulation framework based on this mechanism has the potential to be highly scalable since Neutral Agents act as natural control points for managing valuable simulation resources and for implementing security strategies. Neutral Agents also act as "referees" by providing simulated entities with an impartial interface to the state of the global simulation environment. We also discuss implementation concerns, in particular the use of high-performance open systems.

1 Introduction

We are experts in neither interactive simulation nor military affairs. We are computer architects with long experience in building complex systems out technology that evolves at a frantic pace. Nevertheless, we recently applied our skills and techniques to the problem of improving the Army's distributed simulation capabilities. For a computer architect, distributed simulation is a great problem. The amount of raw computation needed to provide high-fidelity, large-scale simulation is truly amazing. But there is really much more to distributed simulation than the need for fast computers. From a *system-level* perspective, distributed simulation pushes the frontiers of many fields including computation, communication, coordination, graphics, and extensible software design.

An analysis of current simulation architectures reveals several basic problems making them unsuitable for future Army applications. First, these architectures cannot scale to support interactive simulations involving hundreds of thousands of entities, as needed by the Army to conduct realistic exercises. Second, the reliability and fidelity of current simulation environments is suspect due to the way in which the state of a simulation — the "ground truth" of the simulated environment — is distributed to and manipulated by each participant. There are no facilities in place to ensure that every participant sees a single consistent version of the ground truth, and there is no way to prevent participants from misusing simulation state. For instance, a simulated entity can easily "peek" at information that would only be available to its opponent in the real world. Finally, security in current simulation architectures needs to be improved. There exists the possibility that a real enemy can "tap into" a simulated exercise, thus gaining valuable information about the capabilities, plans, and tactics of real forces.

We developed the idea of Neutral Agents to address these problems. Neutral Agents provide natural control points for managing the use of valuable simulation resources — such as computation and communications bandwidth — and for implementing security strategies. Neutral Agents also provide fair, impartial interfaces to the state of a simulation, thus improving consistency and fidelity.

The rest of this paper is organized as follows. Section 2 shows why simulation is a compelling tool for the Army and discusses current simulation technology. Section 3 introduces Neutral Agents and discusses how they are used in a new simulation architecture. Section 4 points out some of the benefits of the Neutral Agents approach. Section 5 emphasizes the importance of open systems in the implementation of future simulation infrastructures, and finally, Section 6 touches on some of the many open problems in Neutral Agents.

2 Simulation: An Emerging Force Multiplier

Since the end of the cold war, the U.S. Army has started to transform itself from "a massive strategically deployed Army with primarily a regional mission" to "a smaller continental U.S.-based force with a broad, global mission" – the force projection Army.[1] To meet this challenge the Army is leveraging the rapidly evolving capabilities of information technology to change the way it operates. Information technology is a force multiplier for today's Army, and its importance will increase in the future.

The goal is to create a digitized force where soldiers and their commanders have unprecedented situational awareness allowing them to maximize their combat effectiveness. Still, commanders must transform this additional knowledge into orders and actions, and these orders must be based on the likely evolution of the situation. Simulation is a new information subsystem that complements traditional command, control, communications, computers, and intelligence (C4I) by reducing uncertainty about the future; simulation is literally "computing the future". New simulation technologies will enable the Army to maximize its effectiveness by predicting and preparing for future contingencies at all levels. Simulation should evolve to become the sixth component of the Army's information infrastructure, so that it becomes appropriate to speak of the Army's Command, Control, Communications, Computers, Intelligence, and Simulation system – **C4IS**.

The are several basic requirements for future simulation environments. First, large-scale simulations consisting of at least 100,000 entities of all kinds — including manned and unmanned simulators, computer generated forces, and actual forces suitably instrumented — must be supported. Second, the interactions between these entities must occur with high-fidelity. This means physical laws, communication conventions, etc... must be obeyed by the simulation entities. Finally, the simulation environment must be highly secure. An infrastructure such as this is likely to play an important role in at least four domains of military activity:

- **Training and Readiness:** these are areas where distributed simulation is successful today. Present simulation architectures, however, limit the utility of simulation in this domain to small exercises.
- **Rehearsal:** rehearsal of operations before deployment, or between engagements, will extend the training function of simulation to a wartime capability. This will require effective integration of simulation into the operating C4I system so that simulation deploys with the force.
- **Planning:** from tactical operations to strategic force shaping, simulation can improve the quality of planning tasks. To achieve this goal, it must be possible to populate the simulated environment with a complete range of systems and forces (especially the opposing force) in an affordable manner. In addition, these systems and forces must, at each planning level, display credible, verifiable behavior.
- **Evaluation and Test:** simulation offers major long-term leverage by ensuring that scarce defense dollars are applied to the right future weapons and that those weapons are developed in a faster, more efficient process.

The past decade saw the development of the **Distributed Interactive Simulation (DIS)** system. This system permits collections of manned simulators where participants, regardless of their physical location, interact in a virtual battlefield through a shared, distributed representation capable of simulating a wide range of behaviors. The standardization of the DIS protocols (proposed IEEE Standard 1278) has enabled important exercises like the Synthetic Theater of War-Europe (STOW-E); the success of this exercise and others like it has given initial validation to the concept that *everything is simulation except combat.*

3 Neutral Agents: A Novel Approach to Simulation

The DIS infrastructure has proven quite useful for staging moderately sized virtual battles. Still, Army requirements call for larger simulations with improved fidelity and security. Unfortunately, the DIS protocols

[1] Salomon, Leon E. "Providing the Technological Edge". *ARMY*, February 1995.

are fundamentally based on broadcasting/multicasting *all* simulation information to *all* entities. As discussed below, we believe this property of DIS introduces a number of problems which make it unsuitable as the basis for future simulation infrastructures. We offer Neutral Agents as a solution to some of the problems.

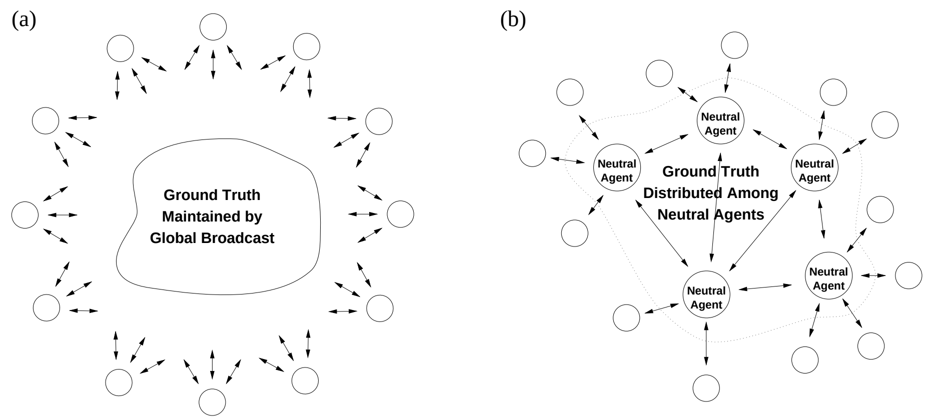

Figure 1: (a) All entities must broadcast to all other entities and filter all information from ground truth; this simulation architecture is *not scalable*. (b) We propose *Neutral Agents* which can act as mediators between end-point entities, reducing communications and computation costs.

Our investigation began by looking at the way the state of the simulation — the "ground truth" of the virtual battlefield — was distributed to participating entities. Figure 1(a) shows the current high-level architecture of a DIS simulation: every entity broadcasts updated information to and receives information from every other entity. This architecture is *not scalable* because communication costs grow as the square of the number of entities, directly proportional to the number of *interactions* between entities. Furthermore, entities must filter information packets that do not concern them from the communications stream. Entities must also perform a significant amount of redundant computation.

Figure 1(b) shows a high-level logical view of the architecture we envision to allow scalable distributed simulation. Rather than having every entity multicast updated information to every other entity, we propose a hierarchical architecture, where Neutral Agents act as mediators between endpoint entities for certain types of communication. In our new architecture, entities communicate mostly with Neutral Agents; direct entity to entity communication is likely to be both uncommon and highly localized.

Neutral Agents can manage many different kinds of information. Indeed, a separate virtual network of Neutral Agents might exist for each type of interaction between entities. For example, one virtual network of Neutral Agents might handle geographic proximity, whereas another virtual network might model communications links. Both virtual networks are likely to be implemented on top of the same *physical* network, and the Neutral Agents themselves may also be implemented on the same physical computers. Yet, the mapping between Neutral Agents and entities may be different for different virtual networks of Neutral Agents.

The mapping between entities and Neutral Agents is dynamic, but slowly changing. Although we expect to update information from Neutral Agents to entities fairly often, we expect to transfer the assignment of one entity from one Neutral Agent to another infrequently. The mapping between entities and Neutral Agents will be dependent upon several factors, including:

- **Physical Geography:** entities which are in physical proximity in the simulation should be handled by the same Neutral Agent.
- **Virtual Communications Links:** entities which must communicate with each other, even when not in close physical proximity, should be handled by the same Neutral Agent.
- **Security Requirements:** entities which communicate at the same level of security within the simulation should be handled by the same Neutral Agent.

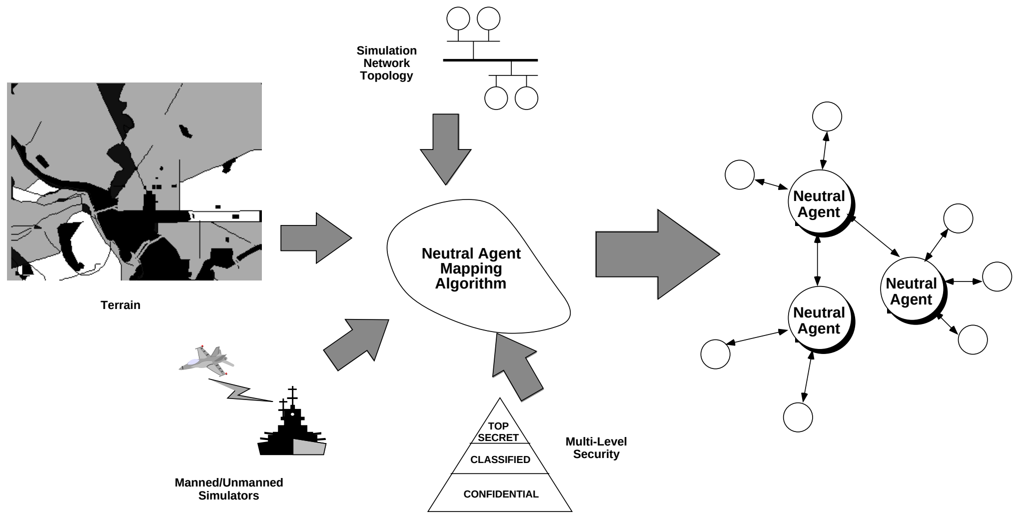

Figure 2: Mapping of entities and Neutral Agents to simulation topology involves a partitioning of work between Neutral Agents. The partitioning must take into account physical, communications, and security interactions between entities, as well as other concerns.

- **Simulation Network Topology:** entities handled by the same Neutral Agent should be in close proximity on the simulation network.

It should be clear that mapping an entity to a Neutral Agent is not a straightforward optimization problem since some of the mapping criteria are at odds with each other. The mapping process also must be dynamic, in order to handle load balancing, movement of entities across the virtual battlefield, and changing terrain and communications links.

4 Benefits of Neutral Agents

Better Use of Available Communications Bandwidth
Current broadcast or multicast techniques will not scale to simulations of 100,000 entities or more – such activity would overload even the fastest projected networks. However, by placing Neutral Agents at crucial places in the simulation infrastructure, packet traffic may be reduced dramatically. Consider the Neutral Agent not only as an arbiter of information, but also as a concentrator and router of information. The Neutral Agent is assigned to monitor a portion of the virtual battle space, and it controls most interactions between the entities populating this portion of the space (and, via neighboring Neutral Agents, all inter-subspace interactions). When an entity moves among battle subspaces, it is registered with the appropriate Neutral Agent.

Since the Neutral Agent determines what information is distributed to entities, it can exploit locality in the physical implementation of the entities. For example, if all the entities in a particular region of the battle space are implemented on the same LAN (Local Area Network), little if any simulation traffic need flow out of the local network onto the WAN (Wide Area Network) that comprises the rest of the simulation infrastructure. The exception, of course, occurs when a local entity takes an action which may affect a different subspace. While such locality will not always be present, the Neutral Agent can still reduce traffic by avoiding multicasts to all entities in the simulation. If information needs to flow out of a local LAN, the Neutral Agent can transmit it to an appropriate neighboring Neutral Agent, who will redistribute it on a need-to-know basis.

We expect that most interactions between entities will not require information broadcast. Neutral Agents take advantage of this. In addition, they provide a general way to control the use of communication bandwidth, which is perhaps the most important factor in making the simulation substrate scalable. Indeed, Neutral Agents embody two of the most important principles in the design of high-performance systems: *leave room for expansion,* and *make the common case fast!*

Less Redundant Computation Performed by Entities

Certain computations currently occur at each endpoint entity, but these could easily be shared by the entities assigned to the same Neutral Agent. *Inter-visibility analysis,* for instance, is an example of the type of computation that would benefit greatly by being centralized. A Neutral Agent could perform this analysis for all its assigned entities, and could then distribute the results to each one. This is much more efficient than having each entity perform the analysis on its own.

Efficient Distribution of Terrain Database

In a realistic physical simulation, terrain data can consist of gigabytes or even terabytes of information. Ideally, we would like to have our endpoint systems be low-cost, commodity computers which gain their power through software and networking. Even today, such machines can perform fairly good simulations of vehicles, given the appropriate software. However, it is impossible today, and unlikely in the near future that any low-cost commodity computer will be able to handle the large databases required to render realistic terrain. Thus, as a simulator moves through the virtual world, it must receive an updated area of the terrain database. This is a natural function for a Neutral Agent to perform since it knows the terrain of its simulation subspace, as well as the location and direction of movement of all entities within the subspace.

In this way, Neutral Agents exploit *temporal* and *spatial* locality in the traversal of the terrain database. Discovering and exploiting locality is a central theme in the design of efficient computer systems, and it is an investment which pays handsome dividends in improved performance.

Provide A Single "Ground Truth"

Neutral Agents can act as "referees" for interactions between entities. Currently, interactions occur without mediation – entities decide for themselves whether any action affects them, to what degree, and in what way. Neutral Agents provide a mediation facility, whereby actions by one entity can be handled by a Neutral Agent, which can report the results to any entities affected. In addition to reducing computational requirements at the endpoints, this function of Neutral Agents increases the fidelity and reliability of the simulation.

In this way, Neutral Agents provide entities with a well-defined interface to the state of the simulated battlefield. In the current simulation infrastructure, each entity is responsible for implementing this interface. This can lead different entities to have different views of "ground truth". Neutral Agents guarantee a uniform view of a single, coherent simulation state.

Ability to Incorporate Multi-Level Security

A broadcast/multicast infrastructure like DIS is vulnerable to attacks of its physical security. A real enemy can "tap into" the simulation infrastructure, and by virtue of the broadcast/multicast protocols, he will eventually gain access to all the information about an exercise. It can be argued that the communications packets can be encrypted to improve security, but in a broadcast/multicast architecture, this offers little relief. All endpoints in such an architecture would need to be aware of the encryption scheme, making the scheme itself difficult to keep secret.

Neutral Agents serve as control points for information flow. The number of connections between Neutral Agents is much smaller than the number of connections between entities in a DIS-style simulation, making the Neutral Agents connections easier to secure. This allows Neutral Agents to incorporate multi-level security naturally. For instance, the scheme used to encode entity to Neutral Agent communications can be simpler (less secure, but less computationally intensive) than the scheme used to protect inter-Neutral Agent communication. An attacker can crack the first scheme, but that will only yield information about a small piece of the battle space. Breaking the second level of security is much harder since the physical communication links are more tightly controlled and the encryption scheme itself is more complex.

Designed for Open Systems

Neutral Agents do not rely on proprietary hardware or protocols. They are designed to be implemented as software on generic machines. This topic is discussed in more detail in the following section.

5 The Importance of Open Systems

Based on our experience in computer architecture, we are convinced that open systems will be vitally important to the success of future simulation infrastructures. Economic pressures dictate that systems built using standard commercial technology will exhibit, at any point in time, unbeatable performance for a given price. Furthermore, over a period time, standard commercial technology will evolve faster than specialized proprietary technology, leading to drastically improved performance from generation to generation. Simulation systems built from commercial hardware can also take advantage of new developments in software technology. Software developers amortize their costs by concentrating on large markets, so it benefits implementors to use standard hardware.

We envision future simulation environments utilizing the entire spectrum of standard computing technologies. The communications fabric would consist of a hierarchy of networks based on FDDI, ATM, emerging gigabyte optical-fiber, or wireless technology. Like today's Internet, this network would provide redundant connections between endpoints, but with guarantees on the quality of service, perhaps honoring even real-time constraints. Endpoint systems for basic simulators should use commodity personal computers. More complex simulators and Neutral Agent implementations should utilize standard symmetric multiprocessors (SMPs) or if necessary, tightly-coupled clusters of these. In short, an open systems-based simulation infrastructure will offer excellent performance and the most flexible path to new technology upgrades.

6 Research Issues in Neutral Agents

As computer architects, we are always reminded to "understand the workload." It should be evident that the potential effectiveness of Neutral Agents is heavily dependent upon the patterns of communication among entities. We can discern these patterns using traces of actual simulations as recorded by packet logging facilities. By understanding the nature of the simulation workload, we can go a long way towards answering a number of the fundamental questions in the design of Neutral Agents, such as:

- How many Neutral Agents are necessary for a given size simulation?
- What kinds of interactions would benefit the most from a Neutral Agent approach?
- Given a simulation network topology, what would be the most effective placement of Neutral Agents?
- What computations should be moved from the entities to Neutral Agents?
- Can the computing and communications infrastructure built to support Neutral Agents be used for additional purposes during war-time operations?
- Can components of existing simulation systems be integrated with Neutral Agents?
- Can Neutral Agents be phased in gradually into current simulation architectures?

Beyond these basic questions lies fertile ground for future research. We expect these investigations will yield the high-fidelity, large-scale simulation capabilities for the Army of the next century.

7 Acknowledgments

Neutral Agents were developed within the context of a proposal in response to the ARL Federated Labs BAA-95. The ideas for this technical area resulted from a close collaboration between the authors, Prof. Seth Teller and Mr. Andrew Shaw of MIT, Drs. Andy Ceranowicz and Se-Hung Kwak of Loral Inc., and Prof. Mark Pullen of George Mason University. The exposition of these ideas in this paper is derived in part from the text of the proposal.