

Network Working Group
Request for Comments: 1667
Category: Informational

S. Symington
MITRE Corporation
D. Wood
MITRE Corporation
M. Pullen
George Mason University
August 1994

Modeling and Simulation Requirements for IPng

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document was submitted to the IETF IPng area in response to RFC 1550. Publication of this document does not imply acceptance by the IPng area of any ideas expressed within. Comments should be submitted to the big-internet@munnari.oz.au mailing list.

Executive Summary

The Defense Modeling and Simulation community is a major user of packet networks and as such has a stake in the definition of IPng. This white paper summarizes the Distributed Interactive Simulation environment that is under development, with regard to its real-time nature, scope and magnitude of networking requirements. The requirements for real-time response, multicasting, and resource reservation are set forth, based on our best current understanding of the future of Defense Modeling and Simulation.

1. Introduction

The Internet Engineering Task Force (IETF) is now in the process of designing the Next Generation Internet Protocol (IPng). IPng is expected to be a driving force in the future of commercial off-the-shelf (COTS) networking technology. It will have a major impact on what future networking technologies are widely available, cost effective, and multi-vendor interoperable. Applications that have all of their network-layer requirements met by the standard features of IPng will be at a great advantage, whereas those that don't will have to rely on less-widely available and more costly protocols that may have limited interoperability with the ubiquitous IPng-based COTS products.

This paper is intended to serve as input to the IPng design effort by specifying the network-layer requirements of Defense Modeling and Simulation (M&S) applications. It is important that the M&S community make its unique requirements clear to IPng designers so that mechanisms for meeting these requirements can be considered as standard features for IPng. The intention is to make IPng's benefits of wide COTS availability, multi-vendor interoperability, and cost-effectiveness fully available to the M&S community.

2. Background: Overview of Distributed Interactive Simulation

The Defense Modeling and Simulation community requires an integrated, wide-area, wideband internetwork to perform Distributed Interactive Simulation (DIS) exercises among remote, dissimilar simulation devices located at worldwide sites. The network topology used in current M&S exercises is typically that of a high-speed cross-country and trans-oceanic backbone running between wideband packet switches, with tail circuits running from these packet switches to various nearby sites. At any given site involved in an exercise, there may be several internetworked local area networks on which numerous simulation entity hosts are running. Some of these hosts may be executing computer-generated semi-automated forces, while others may be manned simulators. The entire system must accommodate delays and delay variance compatible with human interaction times in order to preserve an accurate order of events and provide a realistic combat simulation. While the sites themselves may be geographically distant from one another, the simulation entities running at different sites may themselves be operating and interacting as though they are in close proximity to one another in the battlefield. Our goal is that all of this can take place in a common network that supports all Defense modeling and simulation needs, and hopefully is also shared with other Defense applications.

In a typical DIS exercise, distributed simulators exchange information over an internetwork in the form of standardized protocol data units (PDUs). The DIS protocols and PDU formats are currently under development. The first generation has been standardized as IEEE 1278.1 and used for small exercises (around 100 hosts), and development of a second generation is underway. The current Communications Architecture for DIS specifies use of Internet protocols.

The amount, type, and sensitivity level of information that must be exchanged during a typical DIS exercise drives the communications requirements for that exercise, and depends on the number and type of participating entities and the nature and level of interaction among those entities. Future DIS exercises now in planning extend to hundreds of sites and tens of thousands of simulation platforms

worldwide. For example, an exercise may consist of semi-automated and individual manned tank, aircraft, and surface ship simulators interacting on pre-defined geographic terrain. The actual locations of these simulation entities may be distributed among sites located in Virginia, Kansas, Massachusetts, Germany, and Korea. The PDUs that are exchanged among simulation entities running at these sites must carry all of the information necessary to inform each site regarding everything relevant that occurs with regard to all other sites that have the potential to affect it within the simulation. Such information could include the location of each entity, its direction and speed, the orientation of its weapons systems, if any, and the frequency on which it is transmitting and receiving radio messages. If an entity launches a weapon, such as a missile, a new entity representing this missile will be created within the simulation and it will begin transmitting PDUs containing relevant information about its state, such as its location, and speed.

A typical moving entity would generate between one and two PDUs per second, with typical PDU sizes of 220 bytes and a maximum size of 1400 bytes, although rates of 15 PDUs/second and higher are possible. Stationary entities must generate some traffic to refresh receiving simulators; under the current standard this can be as little as 0.2 PDUs per second. Compression techniques reducing PDUs size by 50% or more are being investigated but are not included in the current DIS standard.

With so much information being exchanged among simulation entities at numerous locations, multicasting is required to minimize network bandwidth used and to reduce input to individual simulation entities so that each entity receives only those PDUs that are of interest to it. For example, a given entity need only receive information regarding the location, speed and direction of other entities that are close enough to it within the geography of the simulation that it could be affected by those entities. Similarly, an entity need not receive PDUs containing the contents of radio transmissions that are sent on a frequency other than that on which the entity is listening.

Resource reservation mechanisms are also essential to guarantee performance requirements of DIS exercises: reliability and real-time transmission are necessary to accommodate the manned simulators participating in an exercise.

M&S exercises that include humans in the loop and are executed in real-time require rapid network response times in order to provide realistic combat simulations. For DIS, latency requirements between the output of a PDU at the application level of a simulator and input of that PDU at the application level of any other simulator in that exercise have been defined as:

- 100 milliseconds for exercises containing simulated units whose interactions are tightly coupled
- 300 milliseconds for exercises whose interactions are not tightly coupled [2].

The reliability of the best-effort datagram delivery service supporting DIS should be such that 98% of all datagrams are delivered to all intended destination sites, with missing datagrams randomly distributed [3].

While these numbers may be refined for some classes of simulation data in the future, latency requirements are expected to remain under a few hundred milliseconds in all cases. It is also required that delay variance (jitter) be low enough that smoothing by buffering the data stream at the receiving simulator does not cause the stated latency specifications to be exceeded.

There are currently several architectures under consideration for the M&S network of the future. Under fully distributed models, all simulation entities rely directly on the network protocols for multicasting and are therefore endowed with much flexibility with regard to their ability to join and leave multicast groups dynamically, in large numbers.

In some cases, the M&S exercises will involve the transmission of classified data over the network. For example, messages may contain sensitive data regarding warfare tactics and weapons systems characteristics, or an exercise itself may be a rehearsal of an imminent military operation. This means the data communications used for these exercises must meet security constraints defined by the National Security Agency (NSA). Some such requirements can be met in current systems by use of end-to-end packet encryption (E3) systems. E3 systems provide adequate protection from disclosure and tampering, while allowing multiple security partitions to use the same network simultaneously.

Currently the M&S community is using the experimental Internet Stream protocol version 2 (ST2) to provide resource reservation and multicast. There is much interest in converting to IPv4 multicast as it becomes available across the COTS base, but this cannot happen until IPv4 has a resource reservation capability. The RSVP work ongoing in the IETF is being watched in expectation that it will provide such a capability. Also some tests have been made of IPv4 multicast without resource reservation; results have been positive, now larger tests are required to confirm the expected scalability of IPv4 multicast. But issues remain: for security reasons, some M&S exercises will require sender-initiated joining of members to

multicast groups. In addition, it is not clear that IPv4 multicast will be able to make use of link-layer multicast available in ATM systems, which the M&S community expects to use to achieve the performance necessary for large exercises.

3. M&S Requirements for IPng

The identified network-layer service requirements for M&S applications are set forth below in three major categories: real-time response, multicast capability, and resource reservation capability. All of these capabilities are considered to be absolute requirements for supporting DIS as currently understood by the M&S community, except those specifically identified as highly desirable. By desirable we mean that the capabilities are not essential, but they will enable more direct or cost-effective networking solutions.

It is recognized that some of the capabilities described below may be provided not from IPng but from companion protocols, e.g. RSVP and IGMP. The M&S requirement is for a compatible suite of protocols that are available in commercial products.

a. Real-time Response

DIS will continue to have requirements to communicate real-time data, therefore the extent to which IPng lends itself to implementing real-time networks will be a measure of its utility for M&S networking. The system-level specifications for the DIS real-time environment are stated in Section 2 above.

b. Multicasting

M&S requires a multicasting capability and a capability for managing multicast group membership. These multicasting capabilities must meet the following requirements:

- Scalable to hundreds of sites and, potentially, to tens of thousands of simulation platforms.
- It is highly desirable that the network-layer multicasting protocol be able to use the multicasting capabilities of link-level technologies, such as broadcast LANs, Frame Relay, and ATM.
- The group management mechanics must have the characteristics that thousands of multicast groups consisting of tens of thousands of members each can be supported on a given network and that a host should be able to belong to hundreds of multicast groups simultaneously.

- Multicast group members must be able to be added to or removed from groups dynamically, in less than one second, at rates of hundreds of membership changes per second. It is not possible to predict what special cases may develop, thus this requirement is for all members of all groups.
- The network layer must support options for both sender- and receiver-initiated joining of multicast groups.

c. Resource Reservation

The M&S community requires performance guarantees in supporting networks. This implies that IPng must be compatible with a capability to reserve bandwidth and other necessary allocations in a multicast environment, in order to guarantee network capacity from simulator-to-simulator across a shared network for the duration of the user's interaction with the network. Such a resource reservation capability is essential to optimizing the use of limited network resources, increasing reliability, and decreasing delay and delay variance of priority traffic, especially in cases in which network resources are heavily used. The resource reservations should be accomplished in such a way that traffic without performance guarantees will be re-routed, dropped, or blocked before reserved bandwidth traffic is affected.

In addition, it would be highly desirable for the resource reservation capability to provide mechanisms for:

- Invoking additional network resources (on-demand capacity) when needed.
- The network to feed back its loading status to the applications to enable graceful degradation of performance.

4. References

- [1] Cohen, D., "DSI Requirements", December 13, 1993.
- [2] Final Draft Communication Architecture for Distributed Interactive Simulation (CADIS), Institute for Simulation and Training, Orlando, Florida, June 28, 1993.
- [3] Miller, D., "Distributed Interactive Simulation Networking Issues", briefing presented to the ST/IP Peer Review Panel, MIT Lincoln Laboratory, December 15, 1993.
- [4] Pate, L., Curtis, K., and K. Shah, "Communication Service Requirements for the M&S Community", September 1992.

- [5] Pullen, M., "Multicast Network Architecture for DIS, briefing presented to the Networks Infrastructure Task Force", George Mason University, C3I Center/Computer Science, November 10, 1993, revised November 11, 1993.

5. Authors' Addresses

Susan Symington
MITRE Corporation
7525 Colshire Drive
McLean, VA 22101-3481

Phone: 703-883-7209
EMail: susan@gateway.mitre.org

David Wood
MITRE Corporation
7525 Colshire Drive
McLean, VA 22101-3481

Phone: 703-883-6394
EMail: wood@mitre.org

J. Mark Pullen
Computer Science
George Mason University
Fairfax, VA 22030

Phone: 703-993-1538
EMail: mpullen@cs.gmu.edu

