

Internet and Multicast Service Issues for XMSF

J. Mark Pullen and Dennis Moen
Department of Computer Science and C³I Center
George Mason University, Fairfax VA 22030
{mpullen,dmoen}@netlab.gmu.edu

Keywords:

Web services, Internet standards, Quality of Service, multicast

ABSTRACT: *The new extended modeling and simulation framework (XMSF) seeks to facilitate expanded distributed interoperation both among modeling and simulation systems and between them and other systems. The facilities and performance provided by underlying networks represent an important constraint on any distributed simulation. In XMSF, their importance is increased because the whole approach rests on capabilities that are fundamentally integrated with the Internet. Because XMSF seeks to exploit technologies developed for the commercial sector, it cannot achieve its objectives if it demands capabilities not available in the commodity Internet of the near future. This paper reviews the network services likely to be needed by XMSF and their projected availability in the near-term Internet, where any-to-any connection cannot be assumed and applications must be able to function over available transport. Network issues include quality of service, multicast, security, and overall performance. Corresponding capabilities needed in XMSF software are reviewed, showing that it will be possible to apply available network capabilities under XMSF to meet modeling and simulation goals. The analysis includes a review of recent new work in overlay (end system) multicast that can provide for efficient many-to-many communication and also provides an approach to providing quality of service on a prioritized basis, within the capabilities of the near-term Internet.*

1. Introduction

The Extensible Modeling and Simulation Framework (XMSF), as defined in [1], is a composable set of standards, profiles and recommended practices for Web-based modeling & simulation (M&S). XML-based markup languages, Internet technologies and Web Services will enable a new generation of distributed M&S applications to emerge, develop and interoperate. XMSF is not a replacement for existing standards such as HLA [2] and DIS [3], but rather an extension that allows Web technologies to be used in extending systems interoperability more quickly and at lower cost. It integrates several high-level requirements derived from years of experience with M&S frameworks, and with the challenges of their effective deployment across diverse networks and systems. XMSF will enable simulations to interact directly and scalably over a highly distributed network, achieved through compatibility between a web framework and networking technologies. It will be equally usable by human and software agents. Clearly XMSF must support composable, reusable model components. Its use will not be constrained by proprietary technology or legally encumbering patents, since such barriers discourage the free, open, *ad hoc* development of interconnected tactical models and simulations.

In [1], the precepts of XMSF are defined as:

- Web-based technologies applied within an extensible framework will enable a new generation of modeling & simulation (M&S) applications to emerge, develop and interoperate.
- Support for operational tactical systems is a missing but essential requirement for such M&S applications frameworks.
- The framework of Extensible Markup Language (XML)-based languages can provide a bridge between forthcoming M&S requirements and open/commercial web standards, while continuing to support existing M&S technologies.
- Compatible and complementary technical approaches are now possible for model definition, simulation execution, network-based education, network scalability, and 2D/3D graphics views.
- The Web approach for technology, software tools, content production and broad use provides best business cases from an enterprise-wide (i.e. world wide) perspective.

The above definitions came from an intensive study, including an invitational workshop and public symposium, whereby the potential structure of XMSF was carefully considered, with inputs from a broad cross section of the stakeholder technical communities:

- Modeling and simulation (M&S)
- Web services
- Internet and networking
- Command, control, communications, computing, and intelligence (C4I)

This paper addresses findings of the XMSF study in the area of supporting network technologies. The requirements are driven not only by the needs of interoperable M&S and C4I, but also by the need for XMSF to be supportable in the commercial Internet. Commercial industry today relies heavily on the Web and Internet for communications. Without Internet supportability, the continued commercial-sector technology development on which XMSF is predicated will not happen. Accordingly, we review XMSF network requirements, near-term commercial Internet capabilities, and technologies that can bridge gaps between the two.

2. XMSF Network Requirements

XMSF will support a scaled list of capabilities to support users with a wide range of network capacities, from wireless access and modems through megabit-per-second Digital Subscriber Line (DSL) and upward through gigabit networks. Over-the-network message formats for corresponding network protocols will be defined unambiguously. Such definitions can be represented flexibly in XML to allow rapid definition of application-specific data streaming formats that include run-time extensibility, portability and semantic interoperability.

2.1 Basic Assumptions

These core assumptions provide a baseline to consider the network services required to support XMSF objectives. The XMSF environment will not be confined to individual networks. Key objectives of XMSF are to expand the customer base, enable a new generation of modeling and simulation applications, and jointly accrue benefits with commercial industry. The implication is that XMSF must reach beyond individual private networks or individual Internet Service Provider (ISP) networks. XMSF must be able to run across the public Internet. Otherwise, it will not provide the benefits needed by commercial industry, upon which we plan to capitalize for Defense purposes.

1. In keeping with the layered abstraction approach, XMSF applications should not be network media-aware. Web Services are designed to be extremely flexible and are most effective when independent of supporting services such as network media. Therefore, XMSF applications should not be dependent on specific network media.

2. Scalability and resilience are essential in XMSF. XMSF applications, middleware and networks must not only be scalable and resilient in the sense of supporting a large number of users, but also from the perspective of

being responsive to unpredictable demands from various interoperating processes as a result of the unpredictable nature of the simulations involved. Without the ability to adapt to fluctuating demands and network services, the goals of XMSF will not be achievable. This implies support mechanisms for fault tolerance, but not direct implementation of fault tolerance, which is application specific.

3. XMSF is not itself a new standard, but rather a set of profiles for use of existing standards. If new standards are required, they will be developed through the existing commercial-sector standards process. For networking standards, this means the Internet Engineering Task Force (IETF).

2.2 Network Service Requirements

While the basic functional requirements for network services have been defined, the Modeling and Simulation community needs to characterize network requirements in a way that can be measured and understood. This includes a shared application-level understanding of the impact if the requirements cannot be met. Key network service requirements that warrant special consideration include:

Network Quality of Service (QoS). QoS must meet a specified or negotiated standard for end-to-end capacity, latency, jitter, and packet loss in a statistical sense. If the approach is a negotiated solution, then a mechanism for negotiation is required, with possibly different solutions for global and local negotiation. Today, QoS can be specified or negotiated within private networks or individual ISP networks, but not across the Internet. For Internet-wide QoS negotiation, no known strategy exists, nor is one expected in the next decade. QoS requirements include the consistency needs of applications and translation to network capabilities. For example, does the application need to know the order of message sending? Achieving certain QoS objectives also implies tradeoff. For example, two very important parameters to XMSF are reliability and latency. Unfortunately, these parameters work inversely as increased reliability implies greater latency.

Multicast. XMSF requires many-to-many multicast (group communication) among instances of distributed applications. The current trend is away from providing this as a network layer service because the business model for the Internet doesn't support the service. One-to-many multicast may become available from individual service providers under the IETF's source-specific multicast (SSM) protocol, but probably will not be available end-to-end across the Internet. This implies that many-to-many multicast must be provided in an overlay/middleware

solution, using a non-multicast network layer. Implicit in the approach is a requirement for an ability to identify and respond to congestion, because multicast networks are very susceptible to congestion.

Reliable Multicast Transport. Internet unicast achieves ordered, reliable transmission using the Transport Control Protocol (TCP). Reliability is also a major concern for multicast networks, as it is impractical in a large multicast group to have fully reliable, order-preserving real-time multicast, comparable to TCP in unicast. The IETF is developing reliable bulk-transfer multicast protocols, but these by their nature will not support real-time requirements. An application that needs to know the order of sending must deal with the problem itself by including sequence numbers in its messages. This implies a need to specify XMSF requirements for reliability in the format of a selectively reliable/real-time and fully reliable/non-real-time capability.

Graceful startup. In the Internet, transmissions are expected to be “TCP friendly” in the sense that they ramp up their sending rate using a ‘slow start.’ The IETF currently is considering proposals that will allow TCP to scale better at higher data rates; however the requirement for graceful startup will remain, thus instant startup will be available only in private networks.

End-to-end network status and performance monitoring. A mechanism must be defined and implemented to provide real-time end-to-end network status and performance monitoring. This information is necessary to the application of middleware for use in adapting to changing network conditions, specifically capacity availability, information loss, and congestion.

Management of policy-based filtering technology. When considering communications across multiple management domains or Autonomous Systems, routing policies, firewalls, and Network Address Translation (NAT) generally prevent straightforward any-to-any communications. Therefore, a mechanism for dealing with policy-based filtering technology that will be encountered in the Internet is required.

Security. XMSF network security requirements must be defined to include authentication, denial of service protection, confidentiality, auditing, and integrity. See section 6.3 of [1] for more information.

Multi-sensor systems. XMSF must support multi-sensor systems, and thus needs the ability to manage streaming data with low buffering latency and also the ability to coordinate groups of sources.

Middleware requirements. There are critical middleware functions that must be included above the network. These include a real-time object request broker, authentication/authorization services, real-time directory services, group coordination/synchronization, and session coordination provided by the Session Initiation Protocol (SIP) or a similar protocol, with addition of an automated setup/teardown capability. XML requires network transfer mechanisms such as XML-RPC or SOAP.

Network time service. Network Time Protocol (NTP) and /or Global Positioning System (GPS) are required to provide synchronized network time for XMSF. GPS is more accurate and can be used to synchronize a local NTP master service.

Over-the-net protocols. Standardizing on over-the-net protocols is a key requirement for success. Riding over standard Internet protocols is a proven basis for enabling interoperability.

Test environment. A dedicated and monitorable test environment will accelerate development of an XMSF community. This can be accomplished using Next Generation Internet (NGI) networks such as Abilene and the Defense Research and Engineering Network (DREN). An initial test environment is being established in 2003 at the Naval Postgraduate School (NPS), George Mason University (GMU), and Old Dominion University (ODU).

3. Near-Term Internet Capabilities

In this section, we review capabilities known to be available in the commercial Internet today, and additional capabilities that can be expected in the near term across the Internet or within ISP networks or private networks.

3.1 Network Services Available Today

QoS and multicast services can be provided today on private networks, including the NGI. In some cases, individual ISPs are capable and willing to provide guaranteed levels of QoS, but make no guarantees for traffic that leaves their managed network. Performance that might be expected in this environment includes:

- Individual path flows to around 100 Mbps.
- Latency under 100 ms round-trip in North America
- Jitter is manageable by buffering but has a latency penalty of 10% or more; with fiber networks overall latency generally can be held to 100 ms round-trip.
- Packet loss guarantees in a private or single source network are easily attainable at less than one percent
- High performance end-to-end with instant startup is practical as long as reliable delivery is not needed. Reliable delivery via TCP is available up to about

100 megabits per second; TCP flow control does not scale well to wide-area flows above this.

- Instant startup of high-performance flows is available on private networks only.
- Good global time synchronization can be made available with NTP/GPS. Review of requirements for secure NTP for special applications is still required.

3.2 Network Services Achievable in 3 to 5 Years

The Internet is a continuously changing environment with a number of on-going initiatives that will offer new services and improved guarantees of QoS. In addition, there are alternative approaches that can be considered by the XMSF that will provide an improved XMSF network environment. Below is a brief summary of key capabilities that the network technology group at the August 2002 XMSF workshop believed will be achieved, based on the current state of technology development:

- QoS on a multi-network basis seems likely, though not Internet wide. This is not a technology issue, rather one of a business case for ISPs. Individual ISPs are likely to form agreements that will allow QoS policy transfer across network interfaces.
- Multicast can be accomplished through applying overlay networks. This can be done using Virtual Private Networks (VPN) or through a middleware that provides application-transparent multicast. Both of these approaches require adequate capacity, but available capacity is not expected to be a problem.
- Significant advances in dynamic caching, based on products that are available today, can significantly enhance performance for digital libraries. It is possible to provide individual data flows up to about one gigabit per second by localizing access. This approach doesn't apply to dynamic data exchanged by simulations, but has the effect of relieving competition for network services demanded by the large file transfers associated with the digital libraries.
- Reliable multicast will be available for non-real-time bulk data transfer.
- Guidelines will be developed for effective management of distributed XMSF simulations including session initiation, network monitoring, and authentication procedures for participants, key distribution, software version control, etc.

4. Technology Development for XMSF

During the process of reviewing XMSF network requirement and near-term commercial Internet capabilities, it became apparent that work is needed to refine XMSF networking requirements and to adapt for

XMSF some techniques that have been demonstrated in networking research but are not available commercially.

4.1 Refined XMSF Networking Requirements

The XMSF report identifies specific recommendations for work necessary to fill out the XMSF profiles in networking. This relatively modest set of fundamental initiatives is required to lay the foundation for providing a successful networking environment for XMSF. These early initiatives include definitional work and development of mechanisms supporting unique XMSF requirements. The following definitional work is needed to develop a clear characterization of M&S network requirements, with well-defined metrics:

- Develop workable definitions of the consistency needs of applications. This is a key aspect that will allow translation from applications requirements to network capabilities.
- Define acceptable tradeoff between reliability and latency in a parameterized form. There must be agreement between M&S and networking on definition and measurement of acceptable reliability and latency.
- Define requirements for reliability in multicast (group communication), such as selectively reliable/real-time, fully reliable/non-real-time.

4.2 Network Technology Enablers for XMSF

The work on mechanisms for effective network support of XMSF includes:

- Develop mechanisms that will allow XMSF to function in the presence of policy based filtering technologies, such as firewalls, NATs, and policy-based routing.
- Provide mechanisms for application or middleware processes to be aware of available network capacity.
- Develop a strategy/capability to support M&S needs for networked group communications over a non-multicast network layer as an overlay network.
- Develop mechanisms for end-to-end network status and performance monitoring

These networking efforts should not be pursued in isolation, but rather in cooperation with applications development for XMSF exemplars. As stated previously, any new networking standardization should take place under the auspices of the IETF.

4.3 Overlay Multicast for XMSF

Two areas where near-term Internet capabilities clearly fall short of XMSF requirements are QoS and multicasting. Here we provide a summary of techniques demonstrated in recent research that can enable these key

services within the context of near-term Internet capabilities.

End System multicast provides an attractive alternative to network based multicast enabling the distributed application to become independent of network based multicast. The case for this approach follows the historical end-to-end argument [6] that any functionality should be (a) pushed to higher layers if possible, (b) unless implementing it at the lower layer can achieve large performance benefit that outweighs the cost of additional complexity at the lower layer. Recent work in overlay networks [7-11] indicates that QoS and multicasting can be realized effectively above the network layer, with the constraint that the overlay network can never achieve better performance than the best that the underlying network is capable of if the use of its resources is optimized.

Figures 1 to 5 show how overlay multicast works. Figure 1 shows the original Internet Multicast approach, using the Distance Vector Multicast Routing Protocol (DVMRP). Data is delivered from the source to the receivers using an IP multicast tree of the shortest paths from each receiver to the source. Figure 1 presents an IP multicast tree constructed using DVMRP. Routers A and B receive a single copy of each packet and forward it along multiple interfaces. However, only one copy of the packet is sent over any physical link. In addition, packet delay is the same to all receivers as though packets were sent directly by unicast.

We can expand the representation of the multicast function from this simple example to a complete graph (figure 2) where every pair of nodes is the set of ends of an edge representing the virtual flow of packets. DVMRP then becomes nothing more than solving the minimum spanning tree problem for this complete graph. If we now remove the multicast function from the routers A and B in the network and build a new tree based on the end systems, then we get a spanning tree as presented in figure 3. The underlying physical path of the packets is now represented as in figure 4. The result is that we are now using some of the physical paths twice and we have added a small amount of delay as a result. The End System multicast concept can be generalized to include nodes at the edge of a network where proxies can use the LAN multicast functionality [7]. These proxies can also act as a router and forward packets on behalf of other nodes (Figure 5).

This tree approach to multicast overlay is simple and efficient and results in a single path between any pair of nodes in the overlay network. BTP, HMTP, and Yoid are examples [8] of tree-based multicast overlay protocols. While the simplicity of these protocols is desirable, they are vulnerable to single link failure potentially resulting in

the partitioning of group membership [9]. Thus, the tree class of protocols requires mechanisms to protect from loss of a link.

For highly dynamic environments with QoS requirements such as that of RT-DVS, it may be more desirable to consider mesh-based overlay multicast protocols. These protocols, such as NARADA and Hypercast, [9] provide mechanisms for redundant connections between group members in the overlay, therefore, offering a protection against single link failure. These protocols, however, do require an approach to protect against routing loops between group members [9]. But they offer the promise of overlay networks that use all available network capacity among group members. To meet Internet requirements, the protocols used will need to be "TCP friendly", that is, to reduce their sending rate in a way similar to TCP when network congestion is encountered, so as not to take unfair advantage of TCP users who generate the majority of today's Internet traffic. An associated capability will be the ability for an XMSF group to regulate the traffic offered by its members on a priority basis. This will ensure that the most essential information is transmitted first, so that XMSF systems can make most effective use of the underlying network.

4.4 QoS for XMSF

For the overlay multicast approach to successfully meet the requirements of XMSF, the client application must be able to specify or request its QoS requirements across the overlay network. The basic QoS parameter generally specified today is bandwidth. This is inadequate to meet XMSF performance objectives. As indicated in paragraph 2.2 above, XMSF will require bandwidth, end-to-end delay, packet loss and jitter as specified constraints to meet performance objectives.

Protocol implementation of multidimensional QoS constraints across a network is very complex. While it clearly is feasible to make available via the overlay network nearly all of the connectivity of which the underlying network is capable in a static situation where long setup times are possible, the capability to maintain QoS in shared networks under dynamic traffic loads has not been demonstrated. The best that can be said is that the overlay network will be able to cooperate with the distributed application to optimize QoS within the priorities of the application.

Most QoS protocol approaches today focus on a single metric such as bandwidth [12] and have problems scaling with two metrics. At least two strategies are worthy of consideration to address this requirement. One approach would be to consider adding to the example overlay protocols described above features similar to those

proposed in the Quality of Service Multipath Multicast Protocol (QoSM²P) [12]. The QoSM²P strategy is to allow any client to specify its QoS requirements using as many constraints as necessary. This approach implies that every overlay node in the network is able to maintain information about QoS parameters on incoming paths and there is a mechanism in place to perform required measurements. Another strategy might be to use a proxy for dynamic QoS control similar to that proposed by Menasce for use with work load modeling of server sites [13]. In this strategy, a QoS proxy would continuously monitor the QoS of the paths in the overlay network for the QoS parameters of interest. A metric would be defined that reflects the waited sum of the parameters of interest, thus providing a single metric that can be used in the overlay protocol for establishing and maintaining multicast paths. This single metric approach would result in a reduction of the complexity of the overlay protocol.

The multicast overlay protocol can assist in providing QoS in the overlay by providing a congestion control mechanism. This will provide a means to manage and improve throughput performance of priority data as the number of users increases. Urvoy-Keller and Biersack [14] have proposed a new TCP-friendly Multicast Congestion Control (MCC) model for overlay networks, using a limited buffer for packet storage at the application level. Instead of performing an end-to-end congestion control between the source and every receiver, overlay MCC uses unicast congestion control between adjacent overlay nodes. This approach is consistent with the many-to-many multicast requirements for XMSF. It also is consistent with the general “network awareness” characteristic required of applications in the XMSF environment.

The optimum condition for providing QoS in an overlay would be that the underlying network itself included a premium IP QoS service across multiple domains. As indicated earlier in this paper, there is no business incentive to provide a many-to-many multicast service across independent domains, but there is a business incentive for end-to-end point-to-point QoS service across multiple domains. Approaches to support this type of service have been demonstrated in the Pan-European GEANT network interconnecting European NRENs [15]. Experiments introduced into the GEANT network have demonstrated the ability to provide QoS including guaranteeing bandwidth, delay, jitter, and packet loss based on DiffServ. Such an approach for the Internet could provide the best environment for success with QoS in overlay multicast.

5. Conclusions and Future Work

XMSF holds great promise for improving interoperability of M&S and C4I systems by taking advantage of emerging commercial standards. Web services and networking standards will join the legacy IEEE simulations standards 1278 and 1516 to increase the scope of interoperability and reduce time and cost to achieve it. The work outlined in section 4 above is necessary to provide the Internet component of XMSF. Some of this work is already underway. We encourage the M&S community to complete the required tasks so that the full benefits of XMSF can be achieved.

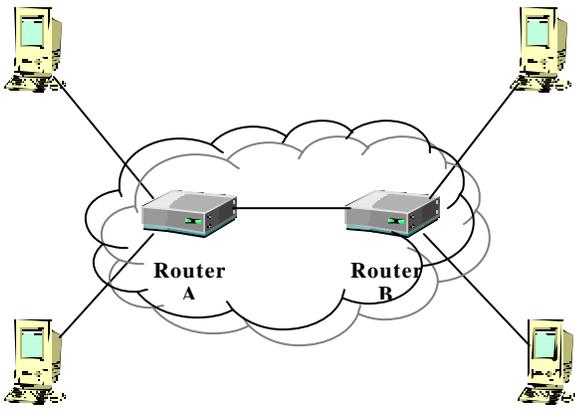


Figure 1. IP Multicast tree using DVMRP

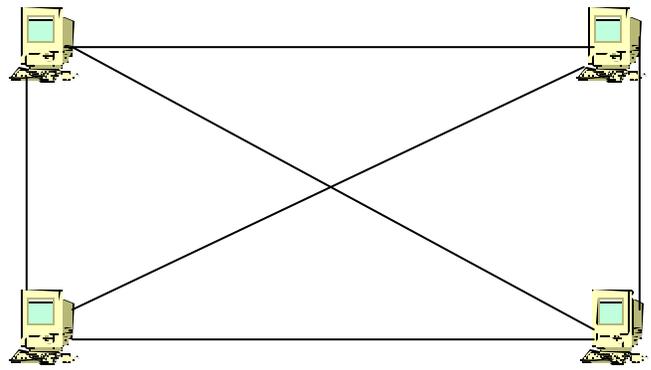


Figure 2. Complete Graph



Figure 3. Spanning Tree Graph

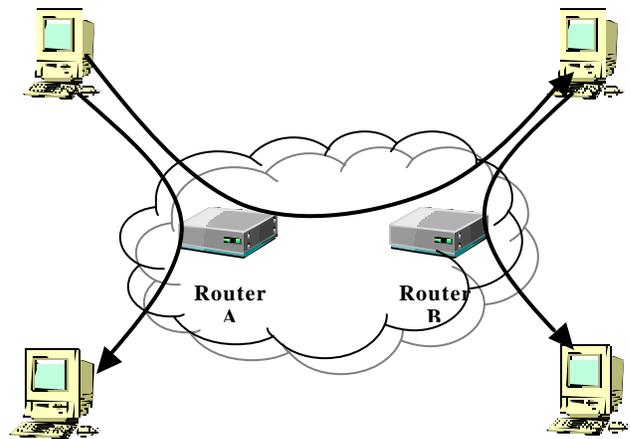


Figure 4. Physical path of the packets

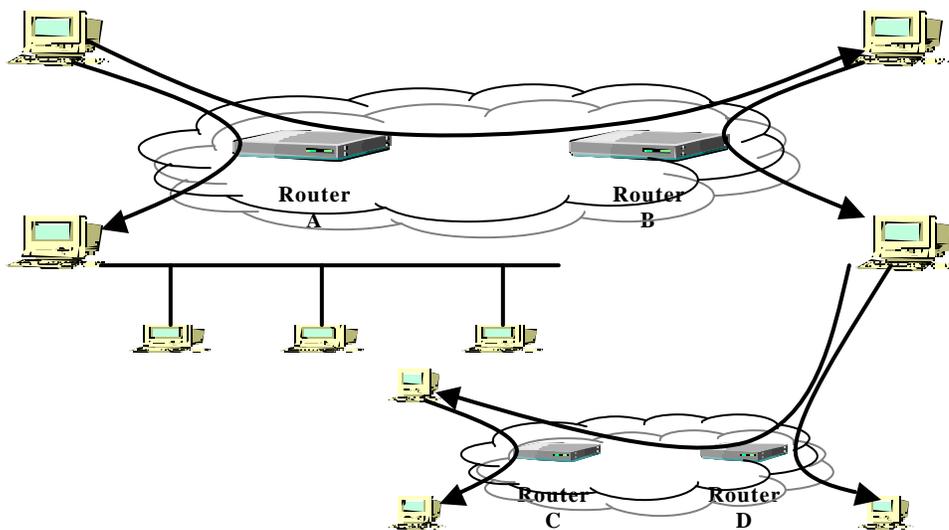


Figure 5. End system proxy nodes

Acknowledgements

This work was supported in part by the Defense Modeling through US Navy contract N00244-02-P-2860.

References

- [1] Brutzman, D., M. Zyda, M., J.M. Pullen, and K.L. Morse, *Extensible Modeling and Simulation Framework (XMSF): Challenges for Web-Based Modeling and Simulation*, US Naval Postgraduate School, 2002
- [2] IEEE Standard 1516, Modeling and Simulation High Level Architecture, 2001
- [3] IEEE Standard 1278, Distributed Interactive Simulation, 1995
- [4] Pullen, J.M., and Simon. R., "Selectively Reliable Multicast for the HLA," Simulation Interoperability Workshop, Orlando, FL, September 2002
- [5] Moen, D. and J.M. Pullen, "A Performance Measurement Approach for the Selectively Reliable Multicast Protocol for Distributed Simulation," *Proceedings of the IEEE Distributed Simulation and Real Time Applications Workshop (DS-RT'01)*, Cincinnati, OH, August 2001
- [6] Saltzer, J., D. Reed, and D Clark. "End-to-end arguments in system design," *ACM Transactions on Computer Systems*, 2(4): p277-288, 1984.
- [7] Chu, Yang-hua, S. G. Rao, S. Seshan, and H. Zhang, "A Case for End System Multicast," *Proceedings of the ACM Sigmetrics*, June 2000.
- [8] Chu, Yang-hua, S. G. Rao, S. Seshan, and H. Zhang, "Enabling Conferencing Applications on the Internet using an Overlay Multicast Architecture," *Proceedings of ACM SIGCOMM2001*, August 2001.
- [9] Wang, W., D. Helder, S. Jamin, and L. Zhang, "Overlay Optimizations for End-host Multicast," *Proceedings of the ACM NGC 2002*, p154-161.
- [10] Garcia-Luna-Aceves, J., D. Balasubramanian, and R. Balakrishnan, "A Scalable and Fault-Tolerant Architecture for Internet Multicasting Using Meshes," *Proceedings of the ACM NGC 2002*, p87-94.
- [11] Savage, S., A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The End-to-End Effects of Internet Path Selection," *Proceedings of ACM SIGCOMM 1999*, p. 289-299.
- [12] Baccichet, Pierpaola, Elena Pagani, and Gian Paolo Rossi, "Quality of Service Multipath Multicast Protocol," *Proceedings of ACM NGC '02*, p. 123-129.
- [13] Menasce, Daniel A., "Automatic QoS Control," *IEEE Internet Computing, January/February 2003*, p. 92-95.
- [14] Urvoy-Keller, Guillaume and Ernst W. Biersack, "A Congestion Control Model for Multicast Overlay Networks and its Performance," *Proceedings of ACM NGC '02*, p. 141-147.
- [15] Roth, Rudolf, Mauro Campanella, Simon Leinen, Roberto Sabatino, Nicolas Simar, Michal Przybylski, Szymon Trocha, Athanassios Liakopoulos, and Afrodite Sevasti, "IP QoS Across Multiple Management Domains: Practical Experiences from Pan-European Experiments," *IEEE Communications Magazine, January 2003*, p. 62-69.

Author Biographies

J. MARK PULLEN is Professor of Computer Science at George Mason University, where he heads the Networking and Simulation Laboratory in the C³I Center. He holds BSEE and MSEE degrees from West Virginia University, and the Doctor of Science in Computer Science from the George Washington University. He is a licensed Professional Engineer, Fellow of the IEEE, and Fellow of the ACM. Dr. Pullen teaches courses in computer networking and has active research in networking for distributed virtual simulation and networked multimedia tools for distance education.

DENNIS MOEN received his M.S. degree in electrical engineering from the University of Arizona. He is a licensed Professional Engineer, and is pursuing a Ph.D in Information Technology at George Mason University. He joined Lockheed Martin Corporation in 2002 as a senior staff member in survivable network architectures and design. His research interests are network performance modeling and adaptive Internet protocols.