

Extreme Simulation of C4ISR Communications and Networking

Donald M. Leskiw¹

The Ultra Corporation
PO Box 50, University Station
Syracuse, NY 13210
leskiwd@ultracorp.com

Donald R. Pflug

AFRL/IFGC

525 Brooks Road
Rome NY 13441-4505
pflugd/sistia@rl.af.mil

Alex F. Sisti

AFRL/IFSB

ABSTRACT

The Internet is viewed by many to be the network of choice for worldwide communications. Indeed, as the Internet proliferates, one will have a global base from which to choose technologies and methods for building intra- and inter- network systems. This emerging information infrastructure comprises wire and wireless networks of vast complexity. A fundamental issue for military users, however, is how to describe the Internet so as to predict its behavior under various conflict scenarios. An unpredictable Internet would certainly contribute to the "fog of uncertainty" that inevitably accompanies command and control operations. In this paper we describe ongoing research that is focused upon modeling the Internet and simulating its behavior: Extreme Simulation of C4ISR Communications and Networking project. Its objective is to demonstrate the feasibility of producing a Global Network Simulator (GNS) for Air Force C4ISR operations, focusing on the Internet for Air Force Operations within the Global Grid. Our approach employs: hierarchical modeling abstraction techniques using object-oriented design and implementation; a new parallel, discrete event simulation framework for large, scalable simulations; and a database schema to specify input and system constructs at runtime.

1. Background

New communication network simulations are needed in C4ISR to demonstrate ensuing advances in global operations and distributed resource management. These include full-scale simulation of networking, protocols, information distribution, and display (including visualization of the global grid itself) to lead to integrated demonstrations of full-scale land-, air-, and space-based internets

¹ This research is being supported in part by: the Air Force Research Laboratory AFRL/IFG under contract F30602-99-C-0089; DARPA/ITO, under Navy contract N6601-96-C-8530 to Rutgers University's Center for Discrete Mathematics & Theoretical Computer Science (a National Science Foundation Science and Technology Center); and by the Ballistic Missile Defense Organization, under Army contract DASG60-92-C-0155. The opinions presented here, expressed or implied, are those of the authors and are not necessarily those of the United States government or its agencies.

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 1999	2. REPORT TYPE	3. DATES COVERED 00-00-1999 to 00-00-1999			
4. TITLE AND SUBTITLE Extreme Simulation of C4ISR Communications and Networking		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Lab/IFGC,525 Brooks Road,Rome ,NY,13441		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 15	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

having explicit modeling of network dynamics, links, antennas, protocols, physical (fading wideband) channels, and traffic. A network simulation having detailed models of information sources and sinks (such as sensor to shooter) together with scenarios and concept-of-operations (CONOPS) descriptions would support C2 concept and technology evaluation, device emulation, and drive interfaces for hardware-in-the-loop demonstrations for Internet-based operations. Beyond scaling to millions of objects, a "big internet" simulation for C4ISR must be capable of reaching down into the lower layers of the protocol stack, i.e., link and physical. For example, verification of an IP-security (IPSEC) mechanism designed to protect wireless frequency bands of a theater intranet may require simulation of a global distributed attack scenario with models of antenna beam forming, link signaling, and power management algorithms, and include multi-path and interference elements of the theater environment.

The reliance upon Dual-Use internets poses inherent challenges for global next generation C4ISR. For example, contrast the traditional military procedure of operating controlled, dedicated theater networks with having to share a cheap, satellite-based multi-access Internet that is also available to anyone from any point on Earth. Currently, IP attacks can affect routing over wire systems and result in "denial-of-service" at the target node. In a wireless system, spectrum saturation within an entire band could result. In a few years the commercial viability of low-cost space-based Internet services will be demonstrated. Currently, low earth orbit (LEO) systems can provide global voice or low-bandwidth data services, i.e., positioning and navigation, while geosynchronous earth orbit (GEO) systems provide wide-area video and also low-bandwidth services such as paging. In the mid-term we see competition among multiple systems offering integrated services. In the longer-term we expect space-based networks to be fully interoperable — internets of LEO systems for multiple access and regional services with mid-earth orbit (MEO), GEO and high earth (or highly elliptical) orbit (HEO) systems providing wide area asymmetric bandwidth services. The emerging global information infrastructure comprises wire-based and wireless networks of vast complexity that will provide internetworking services using Next Generation Internet (NGI) technologies: revolutionary applications, gigabit capacity wireless networking, internetworking satellites, and deeply networked systems.

1.1. Modeling Internet Protocols

Traditional approaches for describing network traffic has relied on statistical representation of aggregate traffic using exponential (Poisson) type models. From August 1989 to February 1992 millions of high quality data were collected on Ethernet LAN's and Leland and Wilson's analysis of those measurements showed the behavior of aggregate traffic to be "self-similar" or fractal-like [Leland and Wilson, 1991]. They demonstrated that a bursty model is more appropriate for LAN Ethernet (TCP) traffic modeling than the assumptions on smoothness that are implied by the Poisson approach [Leland et al, 1994]. These findings have since been verified for wide-area networks [Paxson and Floyd, 1995]. Since then the statistical properties of self-similar network traffic have been characterized using models of long-lasting transient conditions having heavy tailed workloads (distributions whose tails obey a power law) [Crovella and Lipsky, 1997].

These basic changes to assumptions regarding LAN/WAN traffic, from Poisson to fractal characterization, affects how we model the Transmission Control Protocol (TCP) portion of the

Internet. Modeling the routing, the Internet Protocol (IP) is another matter. Here too, empirical studies have shown that an important distinction should be made between routing protocols and actual routing behavior [Paxson, 1996]. Some basic questions are: "What sort of pathologies and failures occur in Internet routing?"; "Do routes remain stable over time or change frequently?"; and "Do routes from *A* to *B* tend to be symmetric (the same in reverse) as routes from *B* to *A*?" Paxson analyzed forty thousand end-to-end Internet route measurements involving a set of diverse Internet sites. He observed and classified routing pathologies in the data, including routing loops and failures deep inside the network. Regarding routing stability it was concluded that "Internet paths are heavily dominated by a single prevalent route, but that the time periods over which routes persist show wide variation" (ranging from seconds to days). Throughout the data, different site-to-site pairs exhibited disparate routing characteristics and no "typical" Internet routing pattern seemed to exist.

The Internet comprises many independent administrative domains, each having its own routing policies. These regions are called autonomous systems (AS). Within each AS routing information consisting of all known routing tables is periodically broadcast using interior gateway protocols. Between autonomous administrative domains, routing information is exchanged at the boundaries through exchange points (EP), typically using the Border Gateway Exchange (BGP) protocol. BGP is an incremental update mechanism. Information on changes inside one domain takes time to reach others; inter-domain information may be outdated, erroneous. It is not uncommon for a router to employ route caching, which works well under stable network conditions. The use of such high-speed mechanisms can create phenomena called "route flap storms". Under BGP, the global routing structure is supposed to converge, however, this has been observed to be not always the case [Labovitz et al, 1997]

1.2. Simulating the Internet

The problem of simulating the Internet is vexed by its very nature as an immense juggernaut. The virtue of IP is that it allows different higher protocols, e.g., http, and network administration policies to seamlessly interoperate. This also makes modeling IP behavior more difficult — IP allows, encourages user diversity. The Internet is a moving target and has undergone dramatic changes in its evolution: in 1992 the Mbone did not exist and the Web was little known. Such fundamental "phase" changes are expected to continue: consider the implications of using the Internet as a "local loop" for video-on-demand services, replacing the ubiquitous POTS (plain old telephone service). At present, the pathologies are being identified, but no unified theory for understanding Internet behavior appears to be on the horizon. New simulation approaches are needed for describing the Internet (wire-based, cellular, and space-based) and to provide a "big picture" which is sufficiently detailed to understand such pathological behavior as well make predictions as the Internet undergoes exponential growth and dramatic paradigm shifts in user applications.

The most encouraging recent development affecting our ability to simulate the Internet is the advent of scalable parallel discrete event simulation (PDES) techniques. One class of such simulation engines is the scalable simulation framework (SSF) being developed by Cooperating

Systems Corporation.² Two versions of SSF are available: a Java implementation from Cooperating Systems Corporation and a C++ implementation from Dartmouth College called DaSSF. SSF provides mechanisms for constructing PDES that can scale to millions of objects using a basic programming interface: SSF.Config, SSF.Entity, SSF.Process, SSF.Event, SSF.inChannel, and SSF.outChannel. Using these one can hope to build large high performance simulators that may begin to reproduce in the laboratory what has been experienced by practitioners [Cowie et al, 1999].

1.3. The Challenge of Extreme Simulation

The objective of the Extreme Simulation of C4ISR Communications and Networking project is to demonstrate the feasibility of producing a Global Network Simulator (GNS) for Air Force C4ISR operations, focusing on the Internet for Air Force Operations within the Global Grid. Figure 1 below illustrates the networking modeling space for this system that will enable the user to specify the scope and detail of selected network segments and global communications scenarios. The system will initially execute on shared memory multi-processor machines. The eventual target is distributed-shared memory processing architectures.

The overall goals of GNS are to:

- Accomplish a High Fidelity Model of the Global Internet for Ground and Space Systems
- Provide Visualization of Global Grid
- Support Demonstration and Evaluation of:
 - New Technologies: Protocols, Network Management, Security, Programmable Radios, Sensor to Shooter
 - Quality of Service Requirements for: C2 Decisions and Planning
- Interface and Operate with
 - External Models, Detailed Tabular Data (ACN Antenna/ Waveform Descriptions)
 - Hardware in the Loop

2. Big Concerns

GNS is intended to uncover big network pathologies, support analysis of large-scale real world situations, and give insight into network operational robustness and susceptibility to attack. We concentrate on four areas of concern that a big network simulation should be capable of addressing. We dub these "big concerns".

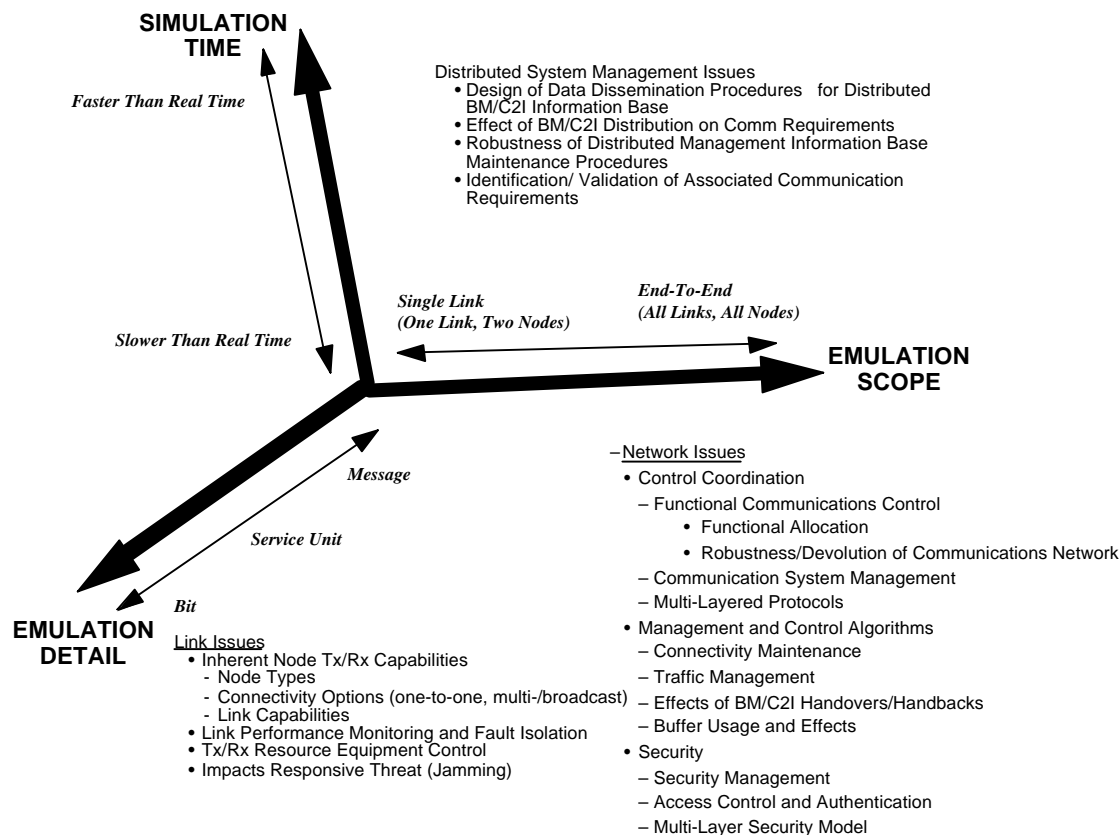


Figure 1 Boundaries of Communication Modeling Space

2.1. Global Extent

The next twenty years will witness dramatic growth in Internet capability and utilization. It is apparent that a second Internet is emerging — above the atmosphere, a vertical, multi-provider, dual use infrastructure of multiple LEO satellite constellations that ultimately will be linked to a backbone of MEO, GEO, and HEO gateways by microwave and laser links. Presently, we have a plethora of disparate space systems providing a variety of services to users. The near-term is viewed as the arena where competing "point-designs" will validate the cost-effectiveness of space-based wireless systems for commercial services. In various stages of planning and deployment, there are, for example: at LEO Iridium, Orbcomm, Skybridge, Celestri, Globalstar, Teledesic; at MEO Orblink, ICO (formerly Odyssey); and in the Geostationary orbital ring at HEO Spaceway, Cyberstar, Astrolink, Immarsat, Intelsat, Aster, KaStar, etc. Services they provide range from video bentpipe/broadcast, to voice/FAX, to low bandwidth signaling/paging. Exponential growth in user demand for services is expected to be the dominant force in shaping a highly volatile market. For example, when the Asian economic recovery starts, interactive services will be a major demand. Although only a fraction of the 1.2 billion people in China are expected to log on in the next few years, the available Ku-based satellite capacity will be swamped, even assuming asymmetric Internet service with the U.S. The Asia-Pacific regional DTH (direct to home)

market is expected to triple by year 2006 to forty million households. VSAT (Very Small Aperture Terminal) based services are likely to be the greatest driver in generating the revolutionary applications which always follow jumps (in speed, in coverage, and in price) in Internet access technologies. This is best summarized in the motto "access is the killer app". Proliferation of cheap satellite access from any point on Earth is the coming, but little researched frontier. The dominant service model undoubtedly will evolve to multiple access, rather than traditional circuit-oriented, in order to provide universal global service across broad spectrum of price/performance. However, the benefits of cheap satellite multi-access to the Internet from any point on Earth will accrue only after addressing the potential for strong detrimental interactions among layered protocols. Attempts to characterize and manage these interactions have so far been hobbled by the lack of scalable, multilayer protocol design and analysis tools.

2.2. Complex Multi-Layer Protocol Models

Standards for Quality of Service (QoS) generally focus on user bandwidth and latency requirements and differ depending on the application characteristics: voice, FAX, data, imagery, video, signaling (transponder paging). For the systems described above, however, the near-term protocols and switching technologies are *not* optimized for high-performance integrated services. Existing Internet protocols implicitly assume a wired infrastructure with symmetric low-loss links, allowing the strict separation of protocol layers above and below IP within the end-to-end paradigm. TCP assumes the underlying network is reliable and that most packet losses are due to congestion. In contrast, for wireless networks packet losses are often due to link loss/fading. Satellite links have the additional problem of latency due to line-of-sight propagation distances (for GEO systems this is at least 250 milliseconds, and with framing, queuing, on-board switching delays can reach as high as 400 milliseconds latency). Standard TCP mechanisms for mitigating packet losses generally make the situation worse when applied to wireless systems. For example, TCP flow control currently assumes the maximum *window size* is 64K bytes (practical window sizes are often 32K bytes). A satellite link with round-trip delay of 0.8 seconds and bandwidth of 1Mbps would require an optimal (theoretical) window size of 100K bytes. TCP performance problems over long delay links can possibly be mitigated by using: large windows (TCP-LW) with bandwidth adaptation; selective acknowledgement (TCP-SACK) and slow start techniques; congestion avoidance (e.g., using random early detection or RED gateways); and transaction TCP or T/TCP. Methods for performance improvement over long delay links have been proposed that include TCP extensions explicitly optimized for space applications, using middleware such as split-TCP and TCP spoofing, or by giving the problem to the application. This last area considers application protocols such as http and includes techniques such as persistent TCP, caching, and application-specific proxies. Although, GEO satellites pose vexing obstacles to Internet applications, Multicast is one that can benefit greatly from the characteristics of GEO platforms. (Hughes Research Laboratories, for example, has demonstrated this using the NASA ACTS, satellite-based MBONE operations.)

Wireless networks (both terrestrial- and space-based) are further challenged by the multiple access problem in which a large number of users transmit to a network server. Techniques such as FDMA, TDMA, and CDMA are employed to manage this chaos with variations for spread-spectrum and wideband applications. Designs for multiple access digital networks may assume

fixed capacity assignments, demand access assignments, or random access capacity allocation. The ALOHA random access protocol developed by Abramson has been shown to be successful with both spread and high bandwidth multiple access.

Rutgers University's Andrew Ogielski has proposed that the next generation of scalable Internet modeling tools must incorporate "design for change" methodologies, supporting a production pipeline that allows a large set of protocol variants to move through simulation, emulation, and deployment without substantial rewriting of protocol code. Design of such a production process is made significantly more difficult by the observation that in satellite Internet, end-to-end application performance is heavily dependent on interactions between multiple layers in the protocol stack: a lesson first painfully learned in terrestrial wireless networking. In the context of terrestrial radio Internet, improving the quality of service seen by a TCP application requires integrated treatment of link level, transport level, and application level concerns. Furthermore, mitigation techniques such as proxying, caching, multilevel compression, and link access scheduling strategies must work in concert to deliver good end-to-end performance to all users. This poses extreme challenges for protocol designers: predicting, inter-activity and in advance, the effects of changes in any protocol layer on all the others, under various alternative traffic patterns characterizing future generations of integrated services networks. For good reason, then, the design and analysis of network protocols has lagged far behind the pace of change and diversification in link-layer technologies, and much of the blame lies in the historical shortage of scalable modeling tools. Without them, achieving an adequate understanding of multiprotocol behavior requires wide deployment in real networks. At this pace, because the Internet is changing rapidly in both size and nature, quantitative characterizations of its dynamics are stale long before they hit the printed page.

2.3. Include the Physical Layer

The physical characteristics of wire-based networks are typically static. Once the physical layer problem is solved for a given medium (copper, optical fiber) and a QoS specification, its properties may be assumed constant and accurately predicted in all ensuing applications. In contrast, the physical layer for wireless networks is more dynamic and susceptible to adverse affects from the environment. Wireless LANs in office buildings can suffer from multipath problems as employees move about office furniture and partitions. These degrade the overall network due to ensuing TCP retransmission requests. Cellular radio transmission in cities can experience severe degradation due to multipath effects from buildings. With military C4ISR applications, sources of (friendly) interference and hostile jammers, will likewise degrade wireless LANs and theater Intranets. At the physical layer transmitter power management and frequency management are problems having solutions that can be very sensitive to terrain composition, antenna elevation, and weather.

The problem of power and frequency management together with antenna beam forming and nulling creates a requirement for an adaptive form of space-time processing. This must take into consideration the locations of its friendly sources and users of transmissions, locations of eavesdroppers and jammers, and the geo-spatial electromagnetic characteristics of terrain and environment. A simple example of the increasing importance of the physical layer in

communications is antenna design for military applications of GPS. The greater accuracy provided by differential GPS using reception from multiple GPS satellites can be easily degraded by jamming. An omni-directional GPS receiver antenna is required, but it needs to selectively null in the direction of an interfering source(s). When multipath is present the problem is far more difficult to mitigate. Certain battlefield applications require high bandwidth, such as the transmission of imagery or video (as in the hand-off between radar or laser ATP sensor and a laser guided weapon system). A UAV tasked with loitering for target selection or BDA mission will also require low-probability of intercept of its (video) transmissions.

Boeing has developed a new phased array antenna that can rapidly acquire and track communication satellites, operate at SHF and EHF, and provide anti-jamming through beam nulling. An IBM wireless ATM system together with the Boeing antenna technology promises battlefield communication bandwidths of one gigabyte. Portions of this system were demonstrated during the Air Force's Expeditionary Force Experiment EFX '98. Participating aircraft included TS-3, AWACS, C-135 avionics testbed, KC135R transport, a B-1B, and a JSTARS. Wireless system researchers are also considering the use of smart materials for self-adjusting satellite antennas. These could change their shape according to signal quality conditions or to provide a limited field of view (LFOV) scanning capability for large aperture satellite antennas.

As systems become increasing complex, steps are being taken to achieve standardized physical layer description languages. An example of this is the Open Systems Architecture for Integrated RF (or photonic) circuits being funded by the Open Systems Joint Task Force (OS-JTF). An eight-layer model is proposed to relate functions from antenna aperture, to signal processing through mission processing interface. The model allows functions for antenna pointing, waveform selection, and processing thresholds to be dynamically threaded and scheduled according to a given mission profile. In summery, the QoS of a wireless network is highly dependent on the reliability and dynamic management of its physical characteristics. Future systems will employ links using antenna systems having adaptive space-time processing and beam and frequency management.

2.4. Scale to Millions of Objects

Recent empirical studies (Ogielski of Rutgers) have established that very large networks such as the Internet exhibit unpredicted pathological behaviors such as "fractal traffic patters" and "congestion storms". Ogielski observes that the lack of a "big picture" of the Internet makes it difficult to analyze empirical data. The example of "congestion storms" is described as "non-local phenomena [that] are anecdotally associated with sudden routing changes, random patterns of aggregation of flows, and other events that may be impossible to correlate using limited, localized measurements.

The capability of GNS to address big network issues implies that GNS itself is big. The only project that has attempted to go beyond the modeling of small (by comparison) networks is the DARPA funded effort described below in the section Related Work. Constructing models of "big Internet" faces three major challenges [Cowie et al, 1999]

- The very large numbers of heterogeneous network elements to be modeled;
- The very long time scales required to model phenomena with long-range correlations, and
- The required level of detail, including the use of protocols and empirical traffic sources.

Network simulations of 100,000 objects, scalable to millions, are unprecedented in scale. Cowie et al have shown that the computational challenge may be met using parallel simulator kernels that can respectively configure and simulate 100,000 network nodes at rates of up to 10^6 network events per second. They claim simulations that generate 10^{10} network events are within practical reach today, and by year 2002 be capable of running in real time. Their claims are based on the Scalable Simulation Framework (SSF) pioneered by Cowie of Cooperating Systems Corporation. Two versions of SSF currently exist: Java SSF for prototyping moderately large network models in Java and DaSSF for C++ (by Nicol at Dartmouth), which offers the extreme high-performance. Java SSF is presently available on Linux/x86 and SUN/Solaris platforms. Additional Unix ports are planned. DaSSF presently runs on SGI IRIX, SUN Solaris, DEC/OSF, Linux, and Windows. Ultra is currently running communications networking models for ASAT scenarios using Java SSF.

3. Approach to the Problem

Our efforts to achieve extreme simulation of C4ISR communications and networking are based on recent developments in three interrelated areas:

- dynamic, hierarchical modeling techniques with runtime binding for coarsening-refinement transitions among alternative subsystem model abstractions that sense and adjust according to activity thresholds;
- runtime support for self-organizing simulations, a common runtime environment that enables objects to communicate indirectly and allows independent dynamic insertion of objects into an ongoing simulation; and a
- system description language that addresses the challenge of building large telecommunications simulations having millions of entities, and executing on a variety of networked heterogeneous platforms.

Our approach to model abstraction of the physical layer is based upon the Kalman methodology for dynamic stochastic processes. A Kalman system model is used to abstract the physics and a Kalman measurement model provides interoperability between disparate state-space representations. The Kalman error covariance equations are used to determine dynamic model accuracy for managing object aggregation and deaggregation. These stochastic objects are embedded into the SSF object-based simulation and provide behavior for sensors, satellites, antenna pointing, propagation, etc.

3.1. Near-Term Technical Objectives

Our primary objective for the near-term is to demonstrate the feasibility of building the Global Network Simulator (GNS) and characterizing its performance (speed and scalability). The global Internet is evolving into a vastly complex system that almost defies description. The two key elements here are the use of a hierarchical-relational database schema (the Graphical Modeling Language, GML) and a highly efficient runtime system for parallel discrete event simulation (the Scalable Simulation Framework, SSF). We plan to rely initially on GML and SSF for building prototype software.

Accordingly, the specific near-term technical objectives are:

- Verify GML is a suitable choice as a hierarchical-relational database schema for describing global Internet and C4ISR constructs
- Verify GML can be used to provide descriptions of networks that scale to millions of objects, still be *manageable*, and support analysis of network communication behavior.
- Demonstrate the use of GML by constructing models of "Global Grid" and describing C4ISR configurations and scenarios;
- Develop and demonstrate (using SSF) multi-level model abstractions (highly- to low-detailed) for GNS ;
- Develop a preliminary design for Interoperating GNS with external protocol, networking, and communication device test-articles, including hardware-in-the-loop; and
- Show how models for physical layer phenomenology can be incorporated into GNS.

3.2. Technical Approach

The Global Network Simulator must employ high performance computing and networking technology to achieve scalable runtime performance. An object-oriented software design approach is needed to build models. Figure 2 provides a schematic overview of the top-level object design for GNS.

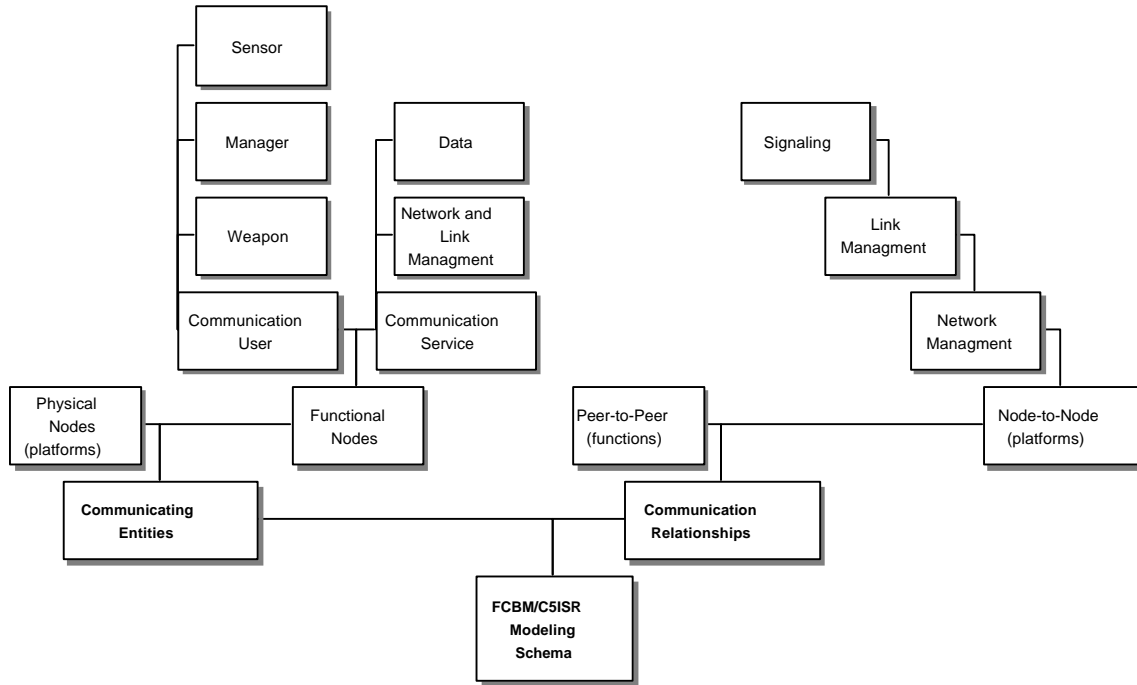


Figure 2 GNS Object Class Architecture

The scope of flexibility illustrated in figure 1 earlier requires that we also use hierarchical multilevel multi-level models to abstract functionality of C4ISR entities and employ runtime object aggregation/deaggregation techniques during run time. Figure 3 illustrates the entity modeling approach for the multi-level abstraction approach.

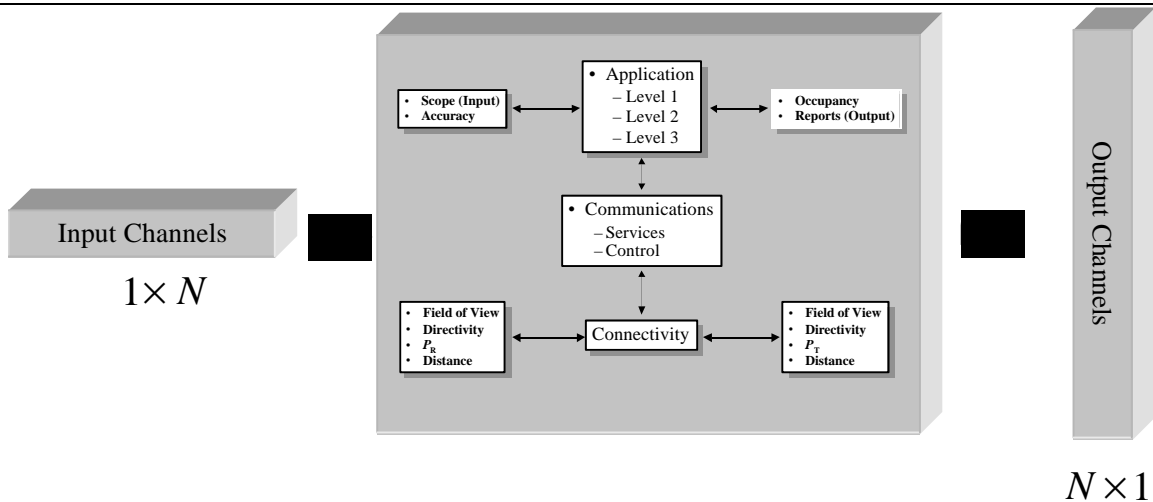


Figure 3 Entity Modeling Abstraction Schematic

A critical step in achieving the goals of GNS is to provide a mechanism for specifying configurations of network architectures, C4ISR architectures, and characteristics of decision makers, sensors, and shooters (the sources and sinks of communications messages for the network). Toward this end two technologies are of interest: the Telecommunication description Language (TeD) currently in use at Georgia Institute of Technology with their version of Time

Warp as a parallel discrete event simulation engine or GML (Graph Modeling Language), currently in use at Rutgers DIMACS. The GML approach is currently preferred as it is a generic hierarchical-relational database schema. GML description fragments can be stored, accessed and manipulated as files, URLs, in a DBMS, and within a program. GML is an extensible, portable data description format (first used in the computational graph theory community for graph drawing and operations on graphs). In GNS a set of reusable GML configuration files would describe hierarchically the attributes (and subattributes) listed below in Figure 4.

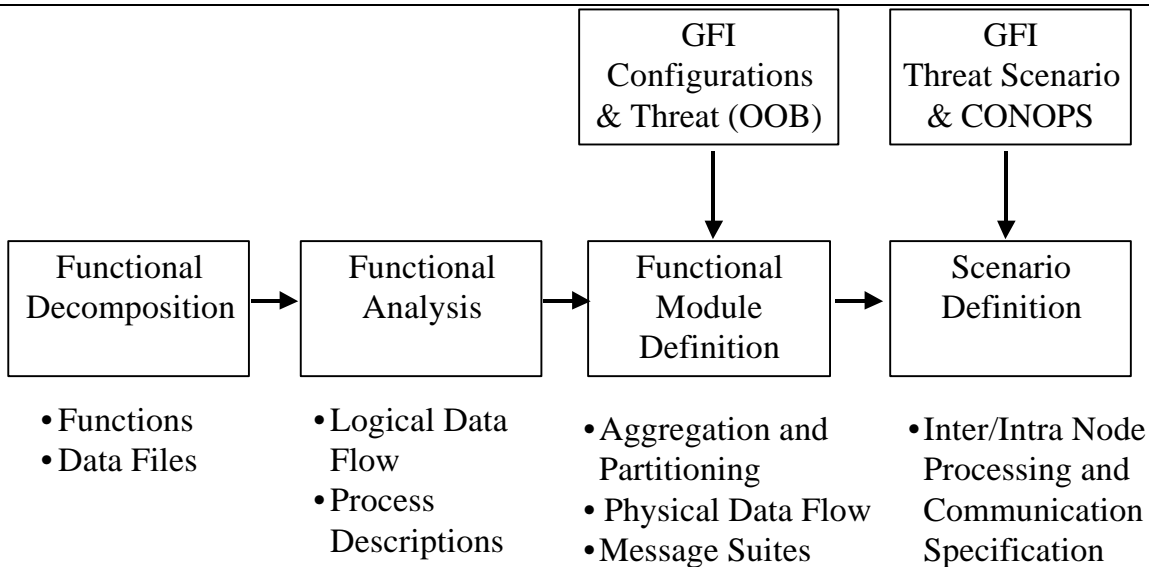


Figure 4 Derivation of GNS Configuration Files

Since GML is both a hierarchical and relational database schema, extremely complex descriptions can be generated and processed to specify C4ISR functions, sensor and shooter characteristics, CONOPS (e.g., rules of engagement), and threat scenarios (including attacks on the network itself). An example of selected portions of a C4ISR runtime configuration is illustrated below in Figure 5, where a set of increasingly complex constructs for sensor, C4ISR, and interceptor combinations for a global scenario involving an ASAT launch (e.g., a crude attack by North Korea on the Soviet space station). The overall scenario involves DPS providing missile launch detection, the global Space Surveillance Network performing surveillance, detection, and tracking of space objects, the SPADOC performing space attack response selection engagement decisions. A generic direct ascent US ASAT is modeled to provided interception of the hostile ASAT. Underlying the sensor, C4ISR, and shooter operations is the communication system, which carries the messages according to satellite link networking.

Scenario	Nodes (Inter/Intra)	Functions	Communications & Networking	Mission Data
ASAT1 FBCCCCISR	• SSN • GBI	• SSN • GBI	• Virtual S00P1000C000	• SV+COV TIVA
ASAT2 FBCCCCISR	• Space Operations Defense Center	• Space Attack Warning/ Attack Assessment Display	• Virtual S00P1000C000	• SV+COV TIVA
ASAT3 FBCCCCISR	• SSN6	• Bandwidth • Time-Delay	• QOS S10P1100C100	• SV TIVA
ASAT4 FBCCCCISR	• GBI2-D (Land- and Ship-Based)	• GBI2	• Virtual S00P1000C000	• SV TIVA
ASAT5 FBCCCCISR	• SSN6	• Bandwidth • Time-Delay	• QOS S10P1100C100	• SV TIVA
ASAT6 FBCCCCISR	• GBI5-D	• GBI5	• QOS S10P1100C100	• SC+COV • Image TIVA
ASAT7 FBCCCCISR	• ASAT6	• Kill Assessment	• QOS S10P1100C100	• SC+COV • Image TIVA

FBCCCCISR: **F**ire Control/ **B**attle Management/ **C**ommand&Control/ **C**ommunications/ **C**omputing/
Collaboration/ **I**ntelligence/ **S**urveillance/ **R**econnaissance
S11P1111C111 **S**ervice Interface Protocols (Levels 1 & 2)/ **P**eer-to-Peer Protocols (Level 1, 2, 3,4)/ **C**ontrol
Protocols (Levels 1, 2, &3)
TIVA **T**ext/ **I**mage/ **V**ideo/ **A**udio

Figure 5 Top-Level C4ISR Scenario Construct Descriptions

4. Progress to Date

A prototype of these scenarios has been built by Ultra and is called SMSC. SMSC was hosted on both HLA and the Java SSF at Ultra and run on the 16-node Solaris machine at Rutgers University. Figure 6 provides an overview of the software engineering builds for SMSC, an early prototype of GMS.

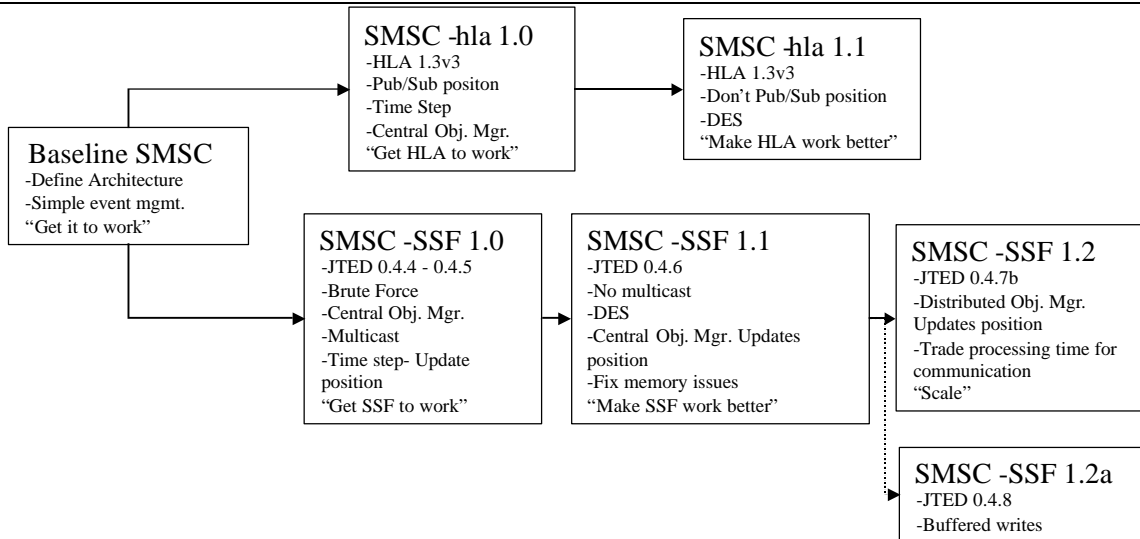


Figure 6 GMS Pre-Prototype Build Configurations

Using the Java SSF-based implementation of SMSC, performance measures were obtained for various loading distributions on a 16-processor Sun system running Solaris. These early results are provided in Figure 7 below. Case 1 involved 50,000 simulation events while Case 2 involved 500,000 events. The results are normalized and show initial scaling for a small number of nodes and that the behavior is consistent for both the modest and large sized problems, but further analysis and tuning is needed to improve performance for greater degrees of parallelization.




Figure 7 Preliminary Runtime Results

5. References

- Leland et al, 1994 W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Internet traffic (extended version)" in IEEE/ACM Transactions on Networking, Vol. 2, No. 1, February 1994, pp. 1-15.
- Leland and Wilson, 1991 W. E. Leland and D. V. Wilson, "High time-resolution measurement and analysis of LAN traffic: Implications for LAN interconnection," in Proc. IEEE INFOCOM '91, Bal Harbour, FL, 1991, pp. 1360-1366.
- Paxson and Floyd, 1995 V. Paxson and S. Floyd, "Wide-area traffic: the failure of Poisson modeling," IEEE/ACM Transactions on Networking, Vol. 3, No. 3, 1995, pp. 226-244.
- Crovella and Lipsky 1997 M. E. Crovella and L. Lipsky, "Long-lasting transient conditions in simulations with heavy-tailed workloads," Proceedings of the 1997 Winter Simulation Conference, 1997.
- Paxson, 1996 V. Paxson, "End-to-end routing behavior in the Internet," Proceedings of GIGCOMM '96, 1996.
- Labovitz, 1997 C. Labovitz, G. R. Malan, and F. Jahanian, "Internet routing instability", ACM SIGCOMM, 1997
- Cowie et al, 1999 J. H. Cowie, D. M. Nicol, and A. T. Ogielski, "Modeling the global Internet," Computing in Science and Engineering, Vol. 1, No. 1, 1999, pp. 42-50.