

ON A GLOBAL INFORMATION GRID SIMULATION PLATFORM FOR INVESTIGATIONS OF END-TO-END PERFORMANCE

R.G. Cole*, L. Benmohamed, A. DeSimone and B. Doshi
Johns Hopkins University Applied Physics Laboratory
Laurel, Maryland

ABSTRACT

We have embarked upon the development of a reusable Global Information Grid (GIG) simulation platform in order to support GIG architecture, design and engineering. We followed a classical systems engineering approach to the design of the GIG simulation platform. We have identified a set of useful performance studies to perform related to GIG architecture and design. We present our approach to the design of our GIG simulation platform and the development of a GIG topology model allowing for automated simulation configurations. As an example of the platform's utility, we briefly discuss our initial studies of BGP performance in GIG-like environments. We conclude with a discussion of follow on work on our GIG simulation platform.

1. INTRODUCTION

The Department of Defense (DoD) is moving toward IP convergence as it develops the Global Information Grid (GIG) IP network. This IP convergence will permeate all aspects and components of the GIG, its backbone transport networks, its satellite networks and its tactical Mobile Ad-Hoc Networks (MANETS). The GIG will comprise a vast array of network technologies, requires support for Military Unique Features, and needs to function under harsh environments. Critical to the design and deployment of the GIG, is the development of a detailed understanding of its end-to-end performance. This includes its con-

trol and management capabilities, as well as its end-to-end packet transport.

For these reasons, we have embarked upon a set of investigations, relying on simulation models, intended to aid decision makers in designing, configuring and deploying the GIG. Our initial studies focus on improvements to the GIG's control plane functions, i.e., the inter-domain routing layer. These studies include Border Gateway Protocol (BGP) [11] performance of satellite links [5], Quality of Service (QoS) enhancements to BGP [2], and eventually scaling and reliability enhancements. In order to perform these studies, we are developing a large scale simulation platform of the GIG. This work builds upon our efforts this year and last to promote the transition of the tools developed under the Defense Advanced Research Projects Agency (DARPA) Information Processing Technology Office (IPTO) Network Modeling and Simulation (NMS) Program to the greater DoD community. Our scalable GIG simulation platform is built upon these tools and our models of the GIG topology were developed as part of the DARPA NMS sponsored project.

In this paper, we present our efforts, progress and initial results in developing a large scale simulation model of the GIG network for Modeling and Simulation (M&S) purposes. In Section 2, we discuss our approach to developing a large scale simulation model of the GIG. In Section 3, we briefly present results on BGP performance from some of our initial investigations. In Section 4, we describe our plans for future investigations and necessary modifications to our GIG simulation platform. The final section of this paper presents our conclusions.

*Corresponding Author. This material is based upon work supported by the Defense Advanced Research Projects Agency Information Processing Technology Office Network Modeling and Simulation Program. Issued by DARPA/CMO under Contract MDA 972-01-D-0005, Task Order Number 0037

2. SIMULATION PLATFORM

In this section, we describe our approach to building up a large scale simulation platform of the GIG. This work was performed under the auspices of the DARPA Network Modeling and Simulation (NMS) program over the last year. Our objective was to encourage transition of NMS tools and techniques to the DoD. The approach we took was to identify a DoD system under development where NMS tools would be of value. We then used this target system as a platform to demonstrate the value of the NMS tools. Toward this goal, we followed a classical, systems engineering, top-down approach resulting in the following steps to developing our GIG simulation platform:

- *Performance Study Definition* - we defined a set of performance studies, that add value to the activities of the GIG Working Groups and the DoD organizations and services in designing the GIG.
- *Requirements* - we derived a set of requirements for tool selection for the GIG simulation platform. The requirements were derived from the needs of the identified performance studies.
- *NMS Tools Assessment* - the tools and techniques from the DARPA NMS program were assessed with respect to their utility in supporting the desired performance studies, i.e., the Requirements.
- *GIG Characterization* - a consistent and realistic model of the GIG was developed, which would lend itself to large scale simulation efforts. This model is consistent with GIG plans across the various branches of the DoD.
- *Test Case* - it was useful to work through a Test Case to validate the methodology. Toward this end, a set of initial studies on BGP performance was performed in order to validate the approach chosen to build our GIG simulation platform. These cases also provided for initial Verification and Validation (V&V) of the current platform tools. Further, they served to refine our understanding of future work to improve the simulation platform.

We expand on our methodology in the following subsections.

2.1. PERFORMANCE STUDIES

Here we discuss the performance studies, which we identified as a useful set of M&S studies for the GIG designers and architects. The intent of defining these studies was to develop a meaningful set of requirements for the tools selection, driven by the end-users' needs, i.e., the DoD. The following performance studies were proposed:

BGP Stability and Scale – The GIG is relying upon BGP to meet various routing demands in a multi-domain and large scale network. This set of simulation studies focuses on the scalability of the alternative inter-domain routing architectures under consideration by the GIG designers. This work also includes convergence time studies of BGP routing behavior under various failure scenarios and under various network mobility assumptions. We proposed using the Network Simulator 2 (NS2) [10], the Parallel and Distributed NS2 (PDNS) [8], and the BGP++ [6] [7] tools for the execution of these studies. We discovered that some integration work is required to support these studies.

QoS Enhancements – The GIG will comprise of various and diverse technologies supporting a broad range of applications with diverse network performance requirements. In addition, the GIG may pioneer the implementation of Precedence and Preemption (P&P) capabilities. In order to accomplish this, enhancements of networking protocols are required, including QoS enhancements to routing protocols, resource reservations through signaling, source routing, tunnel establishments, etc. These will require supporting M&S efforts. A specific example of studies of this type is the investigation of modifications to BGP to carry QoS related metrics [2]. This study would require better integration of BGP++, NS2 and the Integrated Fluid Flow Models (IFFM) [13] from the DARPA NMS project. This would also require modification of the BGP++ code.

End-to-End Application Modeling – Due to the rather broad range of network technologies supporting the overall connectivity within the GIG, this study focuses on the range of end-user perception of performance expected at various application attachment points to the GIG. Potential applications include Voice over IP (VoIP), Situational Awareness (SA), Mobility Alert (MA), Sensor-Fusion-Shooter, and Target Assessment (TA) applications. NMS tools

related to application and network traffic modeling would be useful for these types of study.

Information Assurance – Our GIG simulation platform, and tool sets from the NMS program, could be applied to Information Assurance (IA) investigations of the GIG. Overlay security systems and applications could be run over our GIG simulation platform in order to better understand the performance and effectiveness of various security protocols and systems prior to deployment in the GIG. These studies need to be fleshed out in terms of the nature of the attacks and mitigation technologies under study for potential GIG deployments. Example studies could include the performance of High Assurance IP Encryptor (HAIPE) discovery and routing protocols in GIG deployments, worm infection spreading and mitigation technology effectiveness and prevention, and attacks against the routing infrastructure.

2.2. TOOLS SELECTION REQUIREMENTS

The selection of DARPA NMS tools for our GIG simulation platform was based upon the following criteria, which are driven by our set of identified performance studies:

- *Reusability* - a goal of our work is the development of a reusable GIG simulation tool of utility to the greater DoD community. This imposes requirements on the nature of the tools, their ease of configuration and inter-operation. It also implies that manual methods are not reasonable for the final, integrated simulation environment.
- *Applicability to Studies* - the studies require building a baseline GIG simulation supporting both wired and wireless links, e.g., satellite links, MANETS, transport capabilities including Multi-Protocol Label Swapping (MPLS) and Differentiated Services (DiffSrv), inter and intra-domain routing protocols, and a vast range in link speeds ranging of OC-48 links down to low-speed wireless links on the order of 100 Kbps. Further, various application models and traffic models are required.
- *Scale* - Given the size of the GIG, parallel implementations of the simulation models will be required. Some of the studies address vastly different time scales, e.g., the investigation of BGP

performance over a high-speed network there packet transmissions are sub-microsecond while BGP time scales are on the order of several minutes. This will require hybrid packet-level and flow-level modeling. At a minimum, the simulation tools need to support tens to hundreds of thousands of nodes. More likely, they will be required to simulate several million nodes.

- *Interoperability* - interoperability of the tools is extremely important for ease of use and for developing true end-to-end modeling. It is not feasible to build a reusable GIG simulation platform that requires manual methods for patching together separate simulations of various aspects of the network. Further, running separate, non-interacting simulations may neglect dynamic interactions between the various pieces of the GIG network being simulated separately.
- *DoD Open Source* - the tools used should be generally available to the DoD community. Tools developed under the DARPA NMS program fall within this category. For integration and further development, access to source code is clearly advantageous.

2.3. SIMULATION TOOLS AND PLATFORM

There is no one single tool (or collection of tools) that would currently support the modeling of the diversity of topologies, technologies, and complexities found in the GIG and perform the above identified performance studies. In this section, we describe the selected set of NMS tools and capabilities, which we deem suitable for the execution of the above identified performance studies. We lay out a progression of tools and integration work necessary to incrementally build up a single GIG simulation platform necessary to execute the progressively more complex GIG performance studies.

We decided to base the GIG simulation platform upon the NS2 [10] simulation tool. When necessary, we would then migrate to the PDNS [8] tool. NS2 was chosen as the base simulation engine due to the broad support and breadth of contributions and capabilities made by the Research and Development community over the last fifteen years. NS2 has the broadest protocol support among existing simulation tools; including wired-based and wireless communications, exten-

sive transport layer protocol support, IP and MPLS modeling, extensive DiffSrv models, contributed routing protocol models for intra and inter-domain routing, application models, etc. PDNS represents an enhancement to NS2 for the express purpose of running simulation models in a distributed computing cluster. Both of these tools are available in open source and hence allow for code verification and modification. In the DARPA NMS Program, the PDNS toolkit demonstrated a several million node simulation capability running on a distributed supercomputer at the Penn State University. Basing the GIG simulation facility on NS2 and PDNS affords a growth path to extremely large scale modeling of the GIG. Other selected tools and capabilities, are the BGP++ [6] code for modeling the BGP routing protocol, the Integrated Fluid Flow Models (IFFM) [13], the high fidelity MANET models incorporated into the GloMoSim tools [14], and the background traffic models based upon the work in [3] and [12]. As previously mentioned, these tools will require levels of additional integration coding in order to fully exploit their capabilities in GIG performance studies.

We built the GIG simulation model for our initial studies on the NS2 and BGP++ tools. The NS2 and PDNS tools rely on the TCL scripting language to describe the simulation configuration and their core code is written in C++. BGP++ is also written in C++ and relies on standard BGP daemon router configuration files for the simulation configuration. We developed a Practical Extraction and Report Language (PERL) script, to automatically generate the NS2 and BGP++ configuration code based upon a set of GIG input parameters, e.g., number of core nodes in the GIG-BE, etc. The PERL script generates configuration code for a broad set of parameter definitions and limits. The PERL code automatically generates the GIG topology models described below and allows for the definition of extremely large scale configurations.

We built a core hardware facility to host the GIG simulation platform. Specifically, we deployed Apple Xserve Head and Cluster nodes for our distributed computing complex. These servers dual boot both MAC OSX [1] and Yellow Dog HPC Linux [15]. Distributed simulations will be supported with the Apple XGrid [1] and the Linux Mosix [9] distributed computing tools. The parallel network simulations will be supported via PDNS.

2.4. GIG TOPOLOGY MODELS

We now discuss our work to characterize the topology and scale of the GIG. Our goal was to characterize the GIG as an end-to-end system for large scale M&S purposes. Hence, an objective was to parameterize the topology so that automated generation of simulation configurations were possible. This placed some restrictions on the level of detail placed in the models.

The GIG network is currently under design and development. Hence, the topology decisions made within our simulation models include assumptions on a future network topology. The intent of the topology described herein was to provide a reasonable representation of the future GIG with enough complexity to capture the performance details under investigation. We are not trying to create a detailed model of constituent programs. Instead, we looked to abstract the capabilities of constituent programs to develop an appropriate model for the analysis of end-to-end application and inter-domain routing performance; treating the GIG as an overall system.

Two views of the GIG architecture were addressed: one corresponding to near-term evolution (prior to 2010) and another long-term (after 2012 and the deployment of the Transformational Satellite (TSAT)). The deployment of the TSAT system represents a dramatic change in the routing architecture within the overall GIG network. Hence, we chose the deployment of the TSAT system as the demarcation between our near-term and long-term architectures.

A key architectural construct for the GIG is the *Black Core*. In the black core, all traffic has been encrypted and can be treated as unclassified. In particular, classified IP addresses are not visible in the black core. Red networks - networks that use classified IP addresses - connect to the black core through HAIPE devices that carry the classified traffic in Encrypted Secure Protocol (ESP) tunnels. From a routing perspective, we assumed that classified networks are stub networks and that all routing information is aggregated into route-able prefixes at the point of attachment to the GIG.

Figure 1 shows an example high level, near-term topology chosen as our baseline for our simulation studies. The GIG is composed of various networking domains managed by various organizations within the DoD. These include the GIG Bandwidth

Expansion (GIG-BE) and GIG End User (GIG-U) networks, Teleport, the Transformational Satellite (TSAT) network, the Automated Digital Network System (ADNS), ADNS Users (ADNS-U), the Air Force’s Combat Information Transport System (CITS) network and the Army’s War-fighter Information Network - Tactical (WIN-T). The solid lines in the figure are meant to indicate wired-lines; while the dashed lines indicate wireless connectivity, e.g., satellite links. The core networks, i.e., GIG-BE, ADNS, CITS, and Teleport, are comprised of very high speed communications facilities. Additionally, GIG-BE plans on supporting IP and Multi-Protocol Label Swapping (MPLS) based services for transport. Other GIG networks may provide MPLS transport as well, in addition to IP transport. Our current GIG models capture the dynamic nature of the satellite connectivity down to the tactical environments, e.g., ADNS-U (battleship groups) and WIN-T MANETS (Army Unit-of-Action deployments). Our development of this GIG topology model for M&S purposes is more fully described in [4] and in our related PERL configuration generation script.




Figure 1: The high level topology of the near-term GIG model for simulation studies.

Additional GIG topology work is required to complete the characterization of the GIG for M&S purposes. This additional work includes characterizing the scale of the number of routing prefixes within the GIG for intra-domain routing purposes, characterizing the nature of network mobility, e.g., frequency of deployment of expeditionary forces from battleship groups to WIN-T Unit-of-Employment networks, and further specifics on the anticipated internal architec-

ture of the various GIG networks identified in Figure 1.

3. BGP INVESTIGATIONS

In this section, we briefly discuss our initial performance investigations of BGP performance in the GIG. This represents an initial Test Case of our methodology. BGP is being considered as the routing protocol between various networking domains comprising the GIG. Operational experience on BGP scaling, performance and robustness is derived from the Internet. However, the GIG is a vastly different networking environment than the Internet. Hence, it is extremely important to evaluate BGP capabilities within a GIG deployment scenario.

One large difference between the GIG and the Internet is the GIG’s heavy reliance on satellite communications. In [5], we reported on BGP performance over intermediate satellite links using our GIG simulation platform. Figure 2 shows representative results of BGP traffic overhead under various satellite channel models.

Figure 3 shows BGP example convergence results on a GIG-like topology model consisting of GIG-BE, ADNS, GIG-U, ADNS-U, WIN-T and Teleport Autonomous Systems (ASes). This example topology includes various satellite channels (and models) connecting ADNS-U ASes to ADNS and WIN-T ASes to the Teleport AS. The particular simulation run in Figure 3 shows a satellite channel fading on a WIN-T link to Teleport causing the loss of routing connectivity during the network convergence.

Other studies currently in progress include the impact of QoS enhancements to BGP, see [2]. We will continue development of the GIG simulation platform With better fidelity in our GIG topology models and simulation engine, we expect improved results and utility from this reusable GIG simulation platform.

4. FUTURE WORK

Much work is necessary to achieve our goal of developing a reusable GIG simulation platform based upon the DARPA NMS tools and capabilities for use within the DoD community. In this section, we identify several key, future work items:




Figure 2: The BGP traffic transmitted over a random satellite link as a function of the mean cycle time of the channel model.

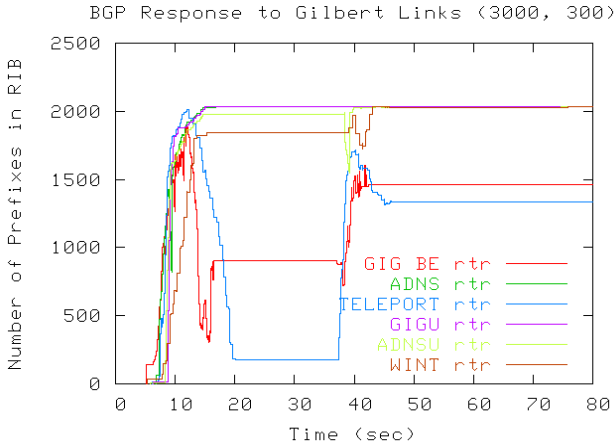


Figure 3: The short-term BGP response in a 500 router implementation.

- *Verification and Validation (V&V)* - extensive V&V is necessary to test the simulation tool sets at each stage of the development. A possibility is to leverage the work of the GIG-Evaluation Facility (GIG-EF) Working Group to accomplish this.
- *Application Traffic Models* - there is a need to develop a set of DoD application models and turn these into useful, integrated traffic generation tools for the simulation studies of the GIG.
- *Improved GIG Topology Models* - work to refine and improve the GIG topology models for the simulation platform is necessary. One example area of future work is incorporating a better, more detailed view of the intra and inter-domain routing architectures, once these are developed by the relevant GIG Working Groups.
- *BGP Routing Integration* - currently the BGP++ simulation tool models only the control messages passed between the BGP routers and the BGP processing within each router. An interface with the routing modules within the NS2 and PDNS simulation packages needs to be developed.
- *Integrated Fluid Flow Models* - the IFFM models do not currently coordinate the flow path with dynamic routing decisions in the NS2 simulation tool. This will require integration of the IFFM with the dynamic routing within the NS2 tool.
- *Integrated Background Traffic Models* - there were several work projects within the DARPA NMS program investigating the characteristics of Internet traffic. These models would be useful in the development of background traffic models for simulation platforms. This would require taking the results from these efforts and developing a background traffic module in NS2.
- *Improved Simulation of Wireless Physical Layer* - the work at UCLA to develop high-fidelity, physical layer wireless models into their simulation [14] is extremely valuable. These models need to be ported to the NS2 simulation tool in order to enhance the GIG simulation model.

5. CONCLUSIONS

We have embarked upon a program to develop a reusable, GIG simulation platform to investigate per-

formance issues within current and future GIG architectures, designs and deployments. This work was initiated during our involvement in the DARPA NMS Program. We have based our GIG simulation platform development on tools extensively developed under DARPA projects, specifically the DARPA NMS Program.

We have initiated a set of performance studies investigating issues with the use of BGP4 as the inter-domain routing protocol for the GIG. We discussed some of our initial studies and results on BGP performance with GIG-like scenarios. More studies are underway. We have concluded this paper with a discussion of future work items necessary to more fully develop our GIG simulation platform and hence improve the range of utility of this platform in addressing the needs of the GIG architects, designers and engineers.

ACKNOWLEDGMENTS

We would like to thank Phil Chimento from JHU/APL on numerous and insightful discussions on our platform development and on the GIG routing plans. Further, we wish to acknowledge our discussions with Eric Naber at JHU/APL on the BGP++ tool, its design, capabilities and on ways to enhance its code. We wish to thank Xenofontos Dimitropoulos at Georgia Tech University for answering our numerous questions on the BGP++ tool. Finally, we wish to thank the reviewers for their insightful comments.

References

- [1] Apple, Corporation, *The Apple Web Site*, <http://www.apple.com/>, January 2005.
- [2] Benmohamed, L., B. Doshi, DeSimone, A. and R.G. Cole, *QoS Enhancements to BGP in the Global Information Grid*, abstract submitted to MILCOM05, February 2005.
- [3] Cleveland, W., Purdue University, *Internet Traffic Models*, <http://www.stat.purdue.edu/~wsc/>, February 2005.
- [4] Cole, R.G., Benmohamed, L., Chimento, P., DeSimone, A. and B. Doshi, *A Topology Model of the Global Information Grid for M²S*, JHU/APL Internal Technical Report, February 2005.
- [5] Cole, R.G., Benmohamed, L., DeSimone, A. and B. Doshi, *Border Gateway Protocol 4 (BGP4) Performance over Intermittent Satellite Links*, invited talk at IEEE PacRim Conference on Communications, August 2005.
- [6] Dimitropoulos, X.C., and G. Riley, *Simulation Tool for Border Gateway Protocol Studies*, <http://www.ece.gatech.edu/research/labs/MANIACS/BGP++/>, February 2005.
- [7] Dimitropoulos, X.C. and G. Riley, *Simulation Tool for Border Gateway Protocol Studies*, Internal Georgia Tech University Report, November 2004.
- [8] Fujimoto, R., Georgia Tech University, *The Parallel and Distributed NS2*, <http://www.cc.gatech.edu/compass/pdns/>, January 2005.
- [9] MOSIX, *MOSIX; Cluster and Grid Management*, <http://www.mosix.org/>, January 2005.
- [10] The Network Simulator 2, <http://www.isi.edu/ns2/nam/>, March 2003.
- [11] Rekhter, Y. and T. Li, *The Border Gateway Protocol - 4 (BGP4)*, Internet Engineering Task Force (IETF), RFC 1777, November 1995.
- [12] Riedi, R., Rice University, *Internet Traffic Models*, <http://www.stat.rice.edu/~riedi/> February 2005.
- [13] Towsley, D., University of Massachusetts - Amherst, *Integrated Fluid Flow Models*, <http://www-net.cs.umass.edu/>, January 2005.
- [14] GloMoSim, *GloMoSim; Global Mobile Information System Simulation Library*, <http://pcl.cs.ucla.edu/projects/glomosim/>, January 2005.
- [15] Yellow Dog Linux, Corporation, *The Yellow Dog Linux Web Site*, <http://www.yellowdoglinux.com/>, January 2005.