

An HLA-Based Approach to Quantify Achievable Performance for Tactical Edge Applications

Brian Rivera

Army Research Laboratory
2800 Powder Mill Road
Adelphi, MD
brivera@arl.army.mil

John A. Tufarolo

Pragmeering Corporation
3610 Beaver Ford Rd
Woodbridge, VA
John.Tufarolo@pragmeering.com

Gary Comparetto

Vinay Lakshminarayan
Mohammad Mirhakkak
Ernest H. Page
Nancy Schult
David Yoo

The MITRE Corporation
7525 Colshire Drive
McLean, VA
garycomp,vinay,mmirhakk,epage,ns
chult,dyoo@mitre.org

Keywords:

Mobile Ad Hoc Network Models, High Performance Computing, OPNET

ABSTRACT: *The DoD is pursuing an end-to-end, seamless, network-centric enterprise communications infrastructure to support a wide range of operating conditions and network topologies. Evaluating the achievable performance of this communications infrastructure, as it evolves, is essential to the user community in order to guide their ongoing requirements, design, and procurement activities. Tactical edge applications present significant challenges to network evaluation methods since they often include mobile ad-hoc networks (MANETs) that employ a wide range of platform types (ground-based, air-based, and satellite-based), traffic types (data, voice, video, and multimedia), delivery methods (unicast and multicast), offered traffic loads (kilobits/sec through megabits/sec), and numbers of nodes (from 10s to 1000s). The complexity exhibited by tactical edge applications typically demands the use of Modeling and Simulation (M&S) techniques, supported by high-fidelity models, to adequately quantify achievable performance on an end-to-end basis. However, these high-fidelity models often have very long runtimes, and restrictive limitations on scenario sizing.*

We investigate the application of the DoD High Level Architecture (HLA) and High Performance Computing (HPC) platforms to address the performance demands associated with analyzing tactical edge applications. A federation comprised of two Soldier Radio Waveform (SRW) federates and one Wireless Network after Next (WNaN) federate is developed and executed within an HPC environment at Aberdeen Proving Grounds (APG). High-fidelity OPNET models are used to represent the SRW and WNaN waveforms. Situational Awareness (SA) multicast traffic is delivered among the nodes represented within each of the three federates. Unicast traffic is exchanged between the SRW federates, in the presence of this SA background traffic, using the WNaN federate as a transit network. Performance metrics include: run time, memory allocation, and achievable throughput and latency as a function of background SA traffic load.

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

MAY 2011

2. REPORT TYPE

3. DATES COVERED

00-00-2011 to 00-00-2011

4. TITLE AND SUBTITLE

An HLA-Based Approach to Quantify Achievable Performance for Tactical Edge Applications

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)

MITRE Corporation, 7525 Colshire Drive, McLean, VA, 22102

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

14. ABSTRACT

The DoD is pursuing an end-to-end, seamless, network-centric enterprise communications infrastructure to support a wide range of operating conditions and network topologies. Evaluating the achievable performance of this communications infrastructure, as it evolves, is essential to the user community in order to guide their ongoing requirements, design, and procurement activities. Tactical edge applications present significant challenges to network evaluation methods since they often include mobile ad-hoc networks (MANETs) that employ a wide range of platform types (ground-based, air-based, and satellite-based), traffic types (data, voice, video, and multimedia), delivery methods (unicast and multicast), offered traffic loads (kilobits/sec through megabits/sec), and numbers of nodes (from 10s to 1000s). The complexity exhibited by tactical edge applications typically demands the use of Modeling and Simulation (M&S) techniques, supported by high-fidelity models, to adequately quantify achievable performance on an end-to-end basis. However, these high-fidelity models often have very long runtimes, and restrictive limitations on scenario sizing. We investigate the application of the DoD High Level Architecture (HLA) and High Performance Computing (HPC) platforms to address the performance demands associated with analyzing tactical edge applications. A federation comprised of two Soldier Radio Waveform (SRW) federates and one Wireless Network after Next (WNaN) federate is developed and executed within an HPC environment at Aberdeen Proving Grounds (APG). High-fidelity OPNET models are used to represent the SRW and WNaN waveforms. Situational Awareness (SA) multicast traffic is delivered among the nodes represented within each of the three federates. Unicast traffic is exchanged between the SRW federates, in the presence of this SA background traffic, using the WNaN federate as a transit network. Performance metrics include: run time, memory allocation, and achievable throughput and latency as a function of background SA traffic load.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

1. Introduction

The DoD has been pursuing a doctrine of *network-centric warfare (NCW)* for the past decade and a half [1]. And while the appeal of using information technology to enable the robust networking of well-informed, geographically dispersed forces is clear, designing and reasoning about the behavior of network-enabled forces has proved challenging. A particular challenge is reasoning about network behaviors in a wireless regime – tools that support the robust design and evaluation *mobile ad-hoc networks (MANET)* at very high fidelities and very large scales do not yet exist.

The U.S. Army Research Laboratory (ARL) Mobile Network Modeling Institute (MNMI) seeks to achieve a capability to design and test mobile ad-hoc networks at the levels of fidelity and scale necessary to understand the behaviors of NCW technologies in the full range of conditions in which they will be employed. The MNMI is pursuing the use of High Performance Computing (HPC) as a primary enabler for this capability.

In this paper we describe the development of an initial modeling and simulation (M&S) capability for the MNMI. The capability supports the integration of existing waveform simulations running in an HPC context. Section 2 presents a brief overview of MNMI and its near-term and long-term objectives. Section 3 describes the design principles for a near-term M&S framework. In Section 4, we discuss the development of an instance of the near-term framework using OPNET models of the Soldier Radio Waveform (SRW) and Wireless Network after Next (WNaN) waveforms. Conclusions and directions for future work are given in Section 5.

2. Mobile Network Modeling Institute (MNMI)

The current practice in reconfigurable mobile ad hoc networks is to use empirical models, simulations, emulations, and experimentations in a stovepipe fashion with minimal use of high performance computing (HPC) resources. These evaluations do not have the fidelity or the scalability to adequately predict how large-scale networks will perform in realistic environments. Furthermore, there is currently no efficient way to exploit the results of each of these approaches with a subsequent evaluation.

To promote the use of high performance computing resources for network modeling, the DoD High Performance Computing Modernization Program Office (HPCMOD) funded ARL to create the Mobile Network Modeling Institute in 2008. The MNMI includes researchers from ARL, Naval Research Laboratory, the Communications-Electronics Research, Development, and Engineering Center (CERDEC), MITRE Corporation, Program Executive Office Command Control Communications Tactical (PEO C3T), Rensselaer Polytechnic Institute, Kitware, Stanford University and the University of Minnesota.

The MNMI vision is to develop scalable software tools that transform the ways in which DoD models, simulates, emulates, and experiments with dynamic reconfigurable mobile warfighter networks.

The Institute seeks to exploit the power of HPC and scalable software to (1) develop the fundamental knowledge required to enable *a priori* prediction of the behaviors of diverse and dynamic networks; (2) understand the design trade-offs and impact of various technologies under a wide variety of dynamic adverse conditions; and (3) quantify the impact of network technologies both technically and operationally to make acquisition decisions.

To achieve this vision, the MNMI established several objectives:

- Develop and apply HPC software for the analysis of MANETs in complex environments.
- Develop an enabling interdisciplinary computing environment that links models throughout the Simulation, Emulation, and Experimentation (SEE) cycle.
- Leverage the powerful synergistic relationship between simulation, emulation, and experimentation.
- Expand DoD workforce that is cross-trained in computational software and network science skills.
- Deliver/support software and train the DoD HPC user community, and significantly extend it to key NCW transformation programs.

A key component of the MNMI is the linking of simulation, emulation & experimentation. This is being done through the development of the Network Interdisciplinary Computing Environment (NiCE) [2]. NiCE provides a common data model and format that links models throughout the simulation

(S), emulation (E) and experimentation (E) cycle as illustrated in Figure 1.

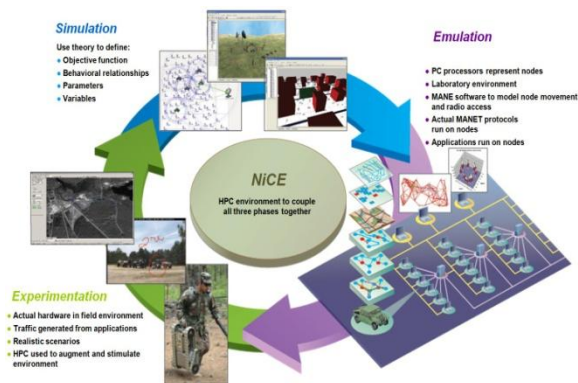


Figure 1: Simulation, Emulation Experimentation (SEE) Cycle.

At the core of NiCE is the cross-platform extensible discrete-event mobile Networking Data Model and Format (NetDMF) [3]. NetDMF is a superset of the eXtensible Data Model and Format (Xdmf). Xdmf is used worldwide in physics based simulations such as computational chemistry, structural mechanics and fluid mechanics. NetDMF extends Xdmf to include network and mobility features needed to support the modeling of mobile networks. Using NiCE, researchers can leverage visualization and analysis tools across all elements of the cycle.

The DoD has made a large investment in developing radio waveforms using existing discrete event simulators (DES) such as OPNET and Qualnet. Given the amount of effort invested in developing and validating these models, the Institute decided it would be infeasible and unrealistic to start from a clean slate. Instead, the Institute decided to break the modeling and simulation effort into near-term and long-term thrusts.

The near-term thrust seeks to demonstrate the value of the SEE cycle approach to Institute stakeholders by using existing legacy waveform models and DES tools with minimal changes and installing them in the HPC environment. While these tools and models cannot fully exploit the HPC environment, they can demonstrate that the vision is achievable and provide the stakeholders with a significantly enhanced analysis capability. A detailed discussion of the near-term thrust is found in Section 3.

The longer term thrust is to develop a massively parallelizable discrete event simulator that can utilize the HPC resources more effectively.

Figure 2 shows how the Institute’s M&S capabilities are evolving over time. The Institute initially started out with individual radio waveform models operating in a single DES running on a single processor. The Institute is currently running multiple radio waveform models on multiple processors still within a single DES. The next planned capability improvement is running multiple radio waveform models developed for use on different DES engines.

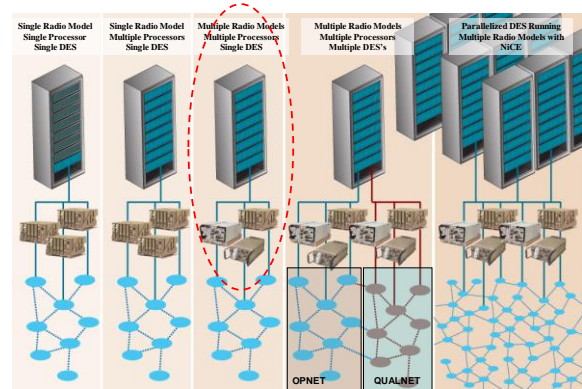


Figure 2: MNMI M&S migration path.

3. Near-Term M&S Framework

The *tyranny of scales* is well-known in distributed systems. Solutions that scale effectively to tens of processing elements (in today’s parlance: cores), often do not scale to hundreds. Similarly, solutions that scale effectively to hundreds of cores may not scale to thousands, nor thousands to tens of thousands, and so on. These limitations are a function of well-known issues in synchronization, communication and fault tolerance, and the availability of exploitable parallelism ala Amdahl’s Law [4] or of a sufficiently-sized workload ala Gufstafson’s Law [5].

While petascale and exascale computing on systems consisting of hundreds of thousands of cores is the scientific frontier [6], the figure of merit for the MNMI is the provision of tools that can effectively scale to use the resources available in today’s “commodity” HPC systems, consisting of tens of thousands of cores.

To date, only a few MANET modeling packages have demonstrated the ability to approach today’s HPC scales (see e.g., [7, 8]). These systems are based on well-studied algorithms in parallel discrete event simulation (PDES) [9] and often have been

finely tuned to match the architecture of the underlying HPC hardware [10, 11]. MNMI is evaluating extant PDES-based MANET modeling packages for their suitability to serve as its long-term M&S framework, with specific attention to two open-source frameworks: the Rensselaer Optimistic Simulation System (ROSS) [7], and ns-3 [12]. Once the long-term framework is selected, significant effort is expected to be required to populate the framework with the waveform models necessary to make it a useful analytical capability. Part of this effort will be undertaken within the auspices of the MNMI, but much of the effort will need to be accomplished through other programs.

In the interim, the objectives of the MNMI include the development of a near-term, HPC-based analytical capability based on waveform models that exist today.

Since a great many models of future waveforms have been developed in the OPNET simulation environment, and since OPNET provides – and actively supports – a High Level Architecture (HLA) interface for its product line, an obvious choice for the interim solution is “OPNET-over-HLA”. Such a solution potentially enables: (1) the analysis of heterogeneous (multi-waveform) networks by federating models of different waveforms, and (2) the scaling of single waveform models by splitting the model into multiple federates.

3.1 Near-term Federation Object Model (FOM)

Given the general approach of federating OPNET simulations, we need to answer the question, at what level(s) in the network protocol stack can runtime interoperation be supported and how does this choice constrain the kinds of scenarios and analyses that can be run?

We observe that for objective (1) above – the analysis of multi-waveform networks where each waveform is encapsulated within a single federate – interoperation at the network layer, specifically via the exchange of Internet Protocol (IP) packets between federates, is sufficient for many network design and analysis problems. In many operational contexts, and for most future waveforms, geographic and frequency separation limits the need to represent interactions at the physical and data link layers. On the other hand, this assumption does not hold for objective (2) above – scaling single waveform models by splitting them into multiple federates. Here, interactions at the physical and Media Access Control (MAC) layers, which are represented in the standalone OPNET model, must generally be preserved as the model is broken into federates.

In the spirit of “you have to start somewhere” and “crawl before you walk”, we adopt an initial approach based on IP packet interoperation (similar to the approaches presented in [13, 14]) and leave MAC and physical layer interoperation for future work.

3.2 HLA scaling

As we note previously, the long-term M&S Framework for the MNMI is anticipated to be in the form of a parallel discrete event simulation (PDES) engine. And it is the long-term M&S framework that must realize HPC scales. The scalability of the near-term HLA-based framework is of lesser importance than its analytical value. However, we are still interested in understanding, and maximizing, the general scalability of the near-term approach.

The performance of the commercially available HLA Runtime Infrastructures (RTIs) is generally well-studied (see, e.g., [15, 16]). These RTIs provide the full-range of HLA services and are generalized to operate in any Ethernet-based network. While the performance advantages of supporting special-purpose computing architectures and very high speed networking fabric have been demonstrated [17, 18, 19], the business case for their support in commercial contexts remains an open question [20, 21].

The target computing platform for the MNMI near-term M&S framework is Harold, an SGI Altix ICE 8200 containing 10,752 compute cores, 32 TB memory and a peak system performance of 120 Teraflops. The compute cores are arranged into 1,344 compute nodes each with dual quad-core Intel Xeon Nehalem-EP 2.8 GHz processors. The compute nodes are connected via 4 Gbps infiniband [22].

Based on the benchmarking approach suggested by Knight *et al.*, we conducted a series of latency and throughput evaluations of three commercially available RTIs running on Harold. We attempted to configure each RTI to achieve maximum performance for each benchmark. However, since our analysis was conducted without active detailed engagement from the RTI vendors, we avoid making any claims regarding the definitive performance of the RTIs here.

The results of a simple time advance benchmark are given in Figure 3 below.

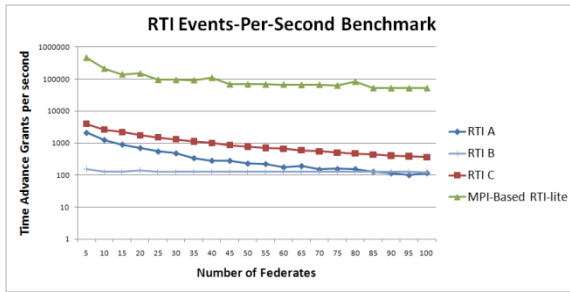


Figure 2: An RTI performance benchmark. Time Advance Grants (TAGs) per second as the number of federates in the federation increases.

In this test, we examine the number of time advance grants (TAGs) per second that an RTI can generate in the presence of some number of federates, each of which is attempting to advance time as rapidly as possible. Essentially, this benchmark evaluates the scalability of the underlying Lower Bound Time Stamp (LBTS) calculation. We scale from 5 to 100 federates and observe that two of the commercial RTIs demonstrate similar results, ranging from thousands of TAGs/sec for 5 federates to hundreds of TAGs/sec for 100 federates. The third exhibits a flat behavior, providing around 100 TAGs/sec for any sized federation.

For sake of comparison with the results reported for “high performance” RTIs (e.g., [17, 18, 19]) we adapted ROSS to provide *minimal* RTI capabilities, namely: (1) support for conservative time synchronization (see [9] for a discussion of adding conservative synchronization to an optimistic simulation engine), and (2) support for HLA services to create/destroy, publish/subscribe, send/receive, and advance time. Running the same benchmark, we observe that this PDES-based minimal RTI is able to generate 500,000 TAGs/sec for 5 federates, scaling to 52,000 TAGs/sec for 100 federates. We conjecture that primary reasons for the superior performance of the ROSS-based RTI are: (1) omission of the full range of HLA services; (2) pointer-based implementation using the high-performance Message Passing Interface (MPI) [23]; (3) tailorability to Harold architecture; and (4) highly-optimized LBTS algorithm (based on ROSS Global Virtual Time (GVT) algorithm).

Since a desire for the near-term M&S framework is to provide a usable, supportable, analytical capability, and since the performance demands required to support an IP-packet-based FOM are expected to be reasonably low, the MNMI is currently pursuing a commercial solution. Our efforts to date have been conducted using the RTI

NG Pro software from RaytheonVTC. As we begin to address MAC and physical layer interoperability—potentially resulting in vastly greater numbers of interactions among federates—and if, and as, the number of federates in our federations scales beyond the hundreds, the need to revisit PDES-based RTI support for the near-term framework may arise.

3.2 Adapting OPNET models for HLA

As noted above, we adopt a basic approach to federating wireless network models at the IP layer similar to those described in [13, 14]. Within each OPNET simulation, we simply add an “HLA node” which serves as a synchronizing agent and a gateway for sending and receiving IP packets to nodes in other federates. While adding this HLA node supports conservative time synchronization across the OPNET federates, with very little perturbation of the original model code, this basic modification does not explicitly transfer any data between models. Additional work is required in order to correctly route desired traffic from one network model to another. Unfortunately, there is no general solution for this; rather it requires understanding the existing model designs, and adding customized interfacing nodes in each model to intercede, encode, send, receive, decode, and utilize traffic appropriately. A case study is described in Section 4 below.

Once data arrives at an HLA routing node, the actual transmitting of data bytes (in this case a representation of an OPNET IP packet) from the OPNET HLA node to the HLA network is trivial compared to the other design steps. This data transmission requires the definition of a small HLA Federation Object Model (FOM) in order to send data across HLA. The OPNET mechanism maps HLA FOM elements to a customized OPNET packet definition. Neither OPNET nor HLA attempt to translate this data; part of the custom HLA routing code must interpret the data and utilize appropriately.

3.3 Executing OPNET/HLA in a batch environment

Many (most) HLA federations operate in an interactive mode, where an operator interacts with a federate or management application to run and pause the federate/federation. Our system is intended for batch submission on remote HPC assets, so relying on such interactive means for control is infeasible.

When operating single OPNET simulations with HLA enabled, the first OPNET federate to successfully join (unless designed otherwise) will immediately begin advancing time in accordance with OPNET’s internal event list.

In order to orchestrate our group of federates and to prevent the above condition, we developed and introduced a controlling “pacer” federate. This federate joins first and forms the federation; sets time management features such that it participates in the time advancement decisions (time-regulating in HLA parlance), then remains at time = t_0 until the expected number of OPNET federates have joined. Once this condition is achieved, the pacer federate removes the time-regulating setting, triggering the federation to proceed forward with all OPNET federates advancing in synchronized fashion based on the modeling activities.

4. Case Study – A Multi-Waveform Prototype

Two key waveforms under development by the DoD to significantly enhance warfighter communications at the tactical edge include the Soldier Radio Waveform (SRW) [24] and the Wireless Network after Next (WNaN) waveform [25].

SRW is being developed through the Joint Tactical radio System (JTRS) program to provide mobile ad-hoc network (MANET) voice and data for dismounted soldiers and small form factor applications. SRW is intended to form stub/leaf networks with limited transit capability to other SRW networks and will rely upon other waveforms for backbone transport service.

The Defense Advanced Research Projects Agency (DARPA) developed the Wireless Network after Next (WNaN) waveform to support inexpensive warfighter-portable multi-hop radios that support spectrum agility, dynamic spectrum access (DSA), and disruption tolerant networking (DTN).

For the purposes of this effort, the WNaN waveform is used to provide the backbone transport service for two separate SRW leaf networks. The SRW OPNET model used in the multi-waveform simulation was developed by the U.S. Army Communications Electronic Research & Development Center (CERDEC) and Materiel Systems Analysis Activity (AMSAA) while the WNaN OPNET model was developed by BBN Technologies.

4.1 A simple multi-waveform architecture

The top level architecture of the Multi-Waveform Tri-Federate Capability is depicted in Figure 4. Although unicast traffic may be sent and received within any federate, only the SRW federates are the source and destination of unicast traffic. Within each SRW federate, routing is handled by the DS Routing

module which is a component of SRW waveform and resides below the IP layer. If the destination of a traffic packet is in the same federate as the source of the packet, DS Routing module routes this packet to its destination. However, if the destination belongs to another SRW federate, the DS Routing module in each SRW node considers the destination “unknown” and routes the packet to the gateway as a default route.

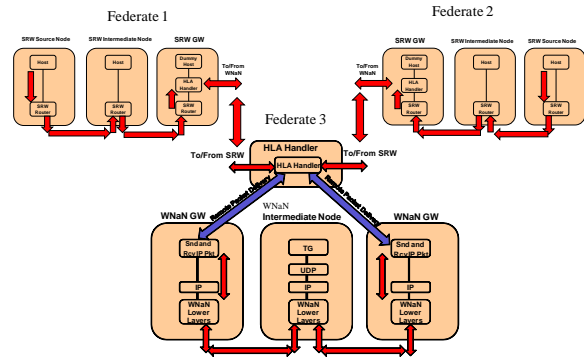


Figure 3: Top Level Architecture of the Multi-Waveform Tri-Federate Capability.

When the packet is received by the SRW radio located at the gateway node, it recovers the original IP packet and forwards it on its host interface. The subnet model representing the gateway has been modified to include an HLA/RTI Handler. A packet with a destination outside the federate within which it is generated is received by the HLA/RTI Handler for transmission to its intended federate. The HLA/RTI Handler receives each of these IP packets, serializes them for transmission through the HLA/RTI environment, and delivers the packet to the HLA/RTI environment. This causes all the HLA federates to receive a copy of the HLA packet. The serialized packet includes the source and destination federate identifiers (IDs) which are used to determine the intended receiver federate of the packet. If the source federate ID is either 1 or 2, the packet is received by federate 3 (the WNaN federate) which acts as a transit network between the two SRW federates.

In the WNaN federate, an HLA/RTI Handler has been added to the WNaN network model, to process packets from/to the HLA/RTI environment. For each external federate, an internal gateway node within the WNaN network is identified as the entry point for packets to/from that federate. The gateway node is an “enhanced” version of the WNaN node/radio

model, with an additional “external_traffic” module above IP to process the external IP traffic received from other federates. This “external_traffic” module takes an IP packet sent from another federate, determines the destination gateway within the WNaN network/subnet, based on a set of destination federate information, and sends the packet as the payload of a WNaN IP packet to the that gateway’s IP address. At the destination gateway node, the “external_traffic” module takes the original IP packet and sends it to the HLA/RTI Handler for processing and subsequently to the HLA/RTI environment.

When the WNaN federate acts as the transit network, the HLA/RTI Handler in the WNaN federate receives the serialized IP packet and reconstructs the original IP packet. Since each federate has a gateway associated with it, the packet is delivered to the appropriate gateway for transmission through the WNaN network to another gateway that is collocated with the destination SRW federate.

When a packet is received by the WNaN gateway that is collocated with destination SRW federate, it forwards the payload IP packet to the WNaN HLA Handler, which serializes the packet and delivers it to the HLA/RTI environment. For this transmission, the source federate ID is set to 3 (ID of WNaN federate), but the original federate ID is maintained. The source and destination node ID and federate ID of a node are derived from the IP address assigned to a node.

When a serialized IP packet is received by an SRW federate, it determines if it is the destination federate and if the source of the packet is the WNaN federate. If so, it routes the packet via its internal network DS routing module. Otherwise, it ignores the packet.

Situational Awareness (SA) traffic can be generated within any federate as multicast traffic. All multicast packets are routed to destinations within the federate. Multicast packets are not forwarded to the HLA/RTI Handlers for delivery to other federates.

4.2 Multi-waveform scenario

Two multi-waveform scenarios were developed as part of this effort and are referred to as the 9-9-9 and 42-75-42 configurations, respectively. Figure 5 provides a high-level portrayal of both configurations. The 9-9-9 configuration includes 9 nodes in each of the three federates whereas the 42-75-42 configuration includes 42 nodes in each of the SRW federates and 75 nodes in the WNaN federate.

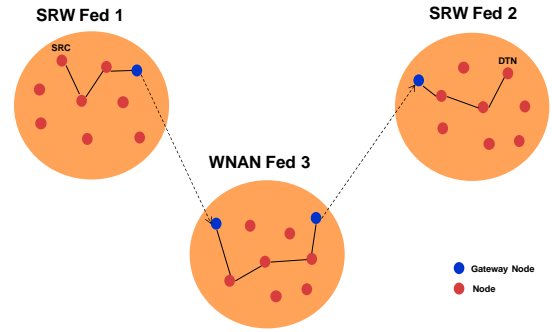


Figure 4: High-Level Multi-Federate Configuration.

Varying levels of multicast situational awareness (SA) traffic is generated within each federate for both configurations. In addition to this SA traffic, a unicast traffic flow is sent from SRW federate 1 to SRW federate 2, using the WNaN federate as a transit network. The complete set of offered traffic loads for both configurations is shown in Table 1.

Both scenario configurations were positioned in the Fort Dix area. The Terrain Integrated Rough Earth Model (TIREM) [26] was used to generate the path attenuation data between each pair of nodes within a federate – there was no connectivity between federates other than through the gateway nodes.

4.3 Execution in HPC environment

As noted Section 3.3, executing an HLA federation in an HPC environment offers a few challenges not typically found in “traditional” HLA applications. HPC execution is predominantly a non-interactive activity, typically initiated by a user submitting a scripted job for execution, and allowing the scripting mechanism to determine where and when to eventually execute the job. For our ARL host HPC system, Harold [22], this scripting feature is provided by the “Portable Batch System” PBS [27].

For this federation, we developed a submission script containing 7 jobs:

- 1 - main launching job, utilizing the PBS “job_array” feature to cause N iterations of the script to run, with increasing index numbers;
- 2 – RTI executive process; job start after job 1 starts
- 3 – Pacing federate; job start after job 2 starts
- 4 – SRW#1 model;
- 5 – SRW#2 model;
- 6 – WNaN model;

5 thru 7 depend on job 3 starting successfully; also made “co-dependent” so they only start when they all can start.

7 – shutdown process; only starts after completion of 3,4,5, and 6. This job includes post-cleanup and collection of output files.

We observe a few challenges to operating in this environment:

- *Little support for RTIexec in non-interactive mode.* There is no clean remote “shutdown” mechanism for the executive once it is started. We rely on PBS job queue delete to remove a running RTI executive process.
- *Wide variance in job execution time.* Often, one of seven jobs in a set started, but others remained queued for many hours. The allotted time for the entire job set expired before full set could complete. This required re-evaluation and adjustment of job submission dependences.
- *Job node selection and control use MPI feature.* PBS features to specifying node resources are primarily oriented towards MPI jobs, and are not easily applied in our context.

4.4 Results

We begin this section with an important disclaimer: the purpose of generating performance results for the two configurations under investigation is to demonstrate the utility of the tri-federate capability developed in this effort. The results presented should not be used to evaluate the capabilities of the SRW and/or WNAN waveforms.

We define 24 configurations for runs made using the Harold high performance computing (HPC) platform, as listed in Table 1. The unicast traffic flow period varies from 1 second to .01 seconds (i.e., packet rate of 1 packet/sec to 100 packets/sec), with a packet size of 1024 bytes. The SA traffic periods are chosen to stress the network performance across the range of unicast traffic levels investigated. Three levels of SA traffic are investigated for each configuration and are referred to as “low”, “medium” and “high”. The SRW SA packet size used is 30 bytes while the WNAN SA packet size is 56 bytes. The resulting SA traffic loads for the 9-9-9 and 42-75-42 configurations are shown in 6 and 7, respectively.

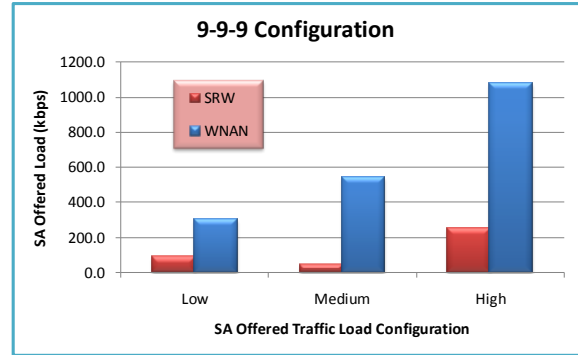


Figure 6: SA Offered Traffic Load for the 9-9-9 Configuration.

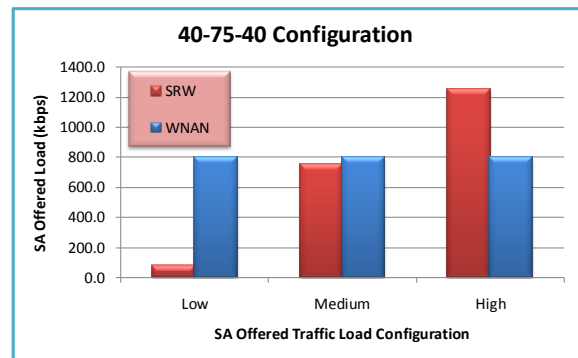


Figure 7: SA Offered Traffic Load for the 42-75-42 Configuration.

The primary figure of merit measured with the tri-federate capability developed in this effort is the achievable message completion rate (MCR). The MCR is measured as a function of the offered unicast SRW federate1--> SRW federate 2 traffic load for the “low”, “medium”, and “high” SA traffic levels. The MCR performance results generated for the 9-9-9 and 42-75-42 configurations are shown in Figure 8 and Figure 9, respectively. As shown, the achievable MCR decreases for both configurations as the SA background traffic load increases and as the unicast traffic load increases.

There are a number of useful applications for the multi-federate capability developed in this effort. For example, this capability could be used to help define the minimum SA traffic period possible in a multi-waveform network within the context of achieving a given MCR for alternative user application traffic profiles. Another application could be to compare the capabilities of alternative transit network waveforms (e.g., WNAN, WNW, etc.) in the presence of varying traffic profiles and user configurations. The capability can be readily

extended to support applications requiring an increased number of federates as well as waveforms where existing models exist to represent their behavior.

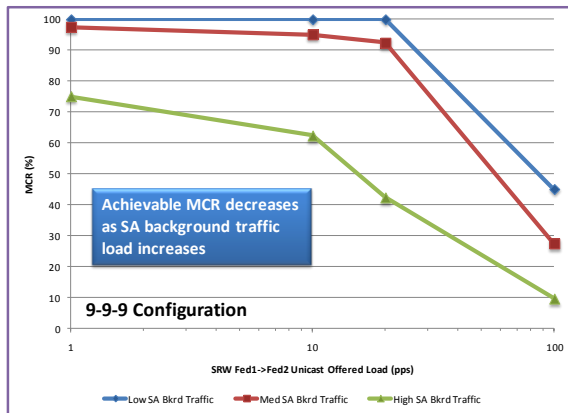


Figure 8: MCR as Function of Unicast Offered Traffic Load for 9-9-9 Configuration.

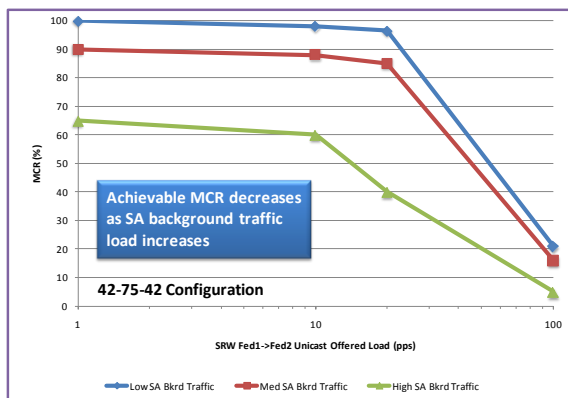


Figure 9: MCR as Function of Unicast Offered Traffic Load for 42-75-42 Configuration.

5. Conclusions

The MNMI seeks to exploit the power of HPC and scalable software to: (1) develop the fundamental knowledge required to enable *a priori* prediction of the behaviors of diverse and dynamic networks; (2) understand the design trade-offs and impact of various technologies under a wide variety of dynamic adverse conditions; and (3) quantify the impact of network technologies both technically and operationally to make acquisition decisions.

While the long-term software solutions to achieve these goals will likely require a reformulation of the waveform models that exist today, we have developed a near-term capability that allows users to

exploit HPC capabilities using existing waveform models today. The MNMI will continue to evolve this near-term capability as it pursues its longer-term objectives.

Acknowledgements

The views expressed herein are those of the authors and do not represent official positions of the United States Department of Defense, the Army Research Laboratory or the MITRE Corporation. The authors would like to thank Dr. Raju Namburu for his vision and leadership for the MNMI. The work reported here also benefits from the valued contributions of: Geoff Kelsch, Jim Smallcomb, Peyton Oliveri, Lisa Hoffman, Joe Greco and Ken Renard.

References

- [1] D.S. Alberts, J.J. Garstka, and F.P. Stein. *Network Centric Warfare – Developing and Leveraging Information Superiority*, 2nd Edition, CCRP. August 1999.
- [2] R.R Namburu, “Network Interdisciplinary Computing Environment for Mobile Network Modeling,” DoD HPC Users group meeting, Seattle, WA, June, 2009.
- [3] NetDMF Overview. [http://netdmf.org/index.php/NetDMF Overview](http://netdmf.org/index.php/NetDMF%20Overview), accessed 1/4/2011.
- [4] G.M. Amdahl. “Validity of Single-Processor Approach to Achieving Large-Scale Computing Capabilities”, in: *Proceedings of AFIPS*, 483-485, April 1967.
- [5] J.L. Gufstafson. “Reevaluating Amdahl’s Law”, *Communications of the ACM*, 532-533, May 1988.
- [6] M.L. Van De Vanter et al. “Productive Petascale Computing: Requirements, Hardware, and Software”, Technical Report, SMLI TR-2009-183, Sun Microsystems, July 2009.
- [7] D.W. Bauer, C.D. Carothers, and A. Holder. “Scalable Time Warp on Blue Gene Supercomputers”, in: *Proceedings of the 2009 Workshop on Principles of Advanced and Distributed Simulation*, 35-44, June 2009.
- [8] K.S. Perumalla. “ $\mu\pi$: A Scalable and Transparent System for Simulation MPI Programs”, in: *Proceedings of the 3rd International ICST Conference on Simulation Tools and Techniques*, 2010.
- [9] R.M. Fujimoto. *Parallel and Distributed Simulation Systems*, John Wiley and Sons, 2000.

- [10] D.W. Bauer and C.D. Carothers. "Scalable RF Propagation Modeling on the IBM Blue Gene/L and Cray XT5 Supercomputers", in: *Proceedings of the 2009 Winter Simulation Conference*, Dec 2009.
- [11] C.D. Carothers and K.S. Perumalla. "On Deciding Between Conservative and Optimistic Approaches on Massively Parallel Platforms", in: *Proceedings of the 2010 Winter Simulation Conference*, 678-687, Dec. 2010.
- [12] The ns-3 network simulator, <http://www.nsnam.org/>, accessed 1/5/2011.
- [13] P.G. Dunn. "Multiple OPNET Simulation Environment using HLA", in: *Proceedings OPNETWORK*, 2004.
- [14] M.J. Dooley, J.D. Dallaire, and J.H. Reaper. "A Simple Approach to Distributed Simulation in OPNET", in: *Proceedings OPNETWORK*.
- [15] P. Knight, et al. "Analysis of Independent Throughput and Latency Benchmarks for Multiple RTI Implementations", in: *Proceedings of the 2002 Fall Simulation Interoperability Workshop*, 02F-SIW-068, Nov 2002.
- [16] P. Knight, et al. "WBT RTI Independent Benchmark Tests: Design, Implementation, and Updated Results", in: *Proceedings of the 2002 Spring Simulation Interoperability Workshop*, 02S-SIW-081, March 2002.
- [17] R. Fujimoto and P. Hoare. "HLA RTI Environments in High Speed LAN Environments", in: *Proceedings of the 1998 Fall Simulation Interoperability Workshop*, 98F-SIW-085, Nov 1998.
- [18] S. Ferenci and R. Fujimoto. "RTI Performance on Shared Memory and Message Passing Architectures", in: *Proceedings of the 1999 Spring Simulation Interoperability Workshop*, 99S-SIW-074, March 1999.
- [19] J.S. Steinman. "WarpIV Kernel: Real Time HPC-RTI Prototype", in: *Proceedings of the 2005 Spring Simulation Interoperability Workshop*, 05S-SIW-071, March 2005.
- [20] M.D. Myjak, D. Clark and T. Lake. "RTI Interoperability Study Group Final Report", in: *Proceedings of the 1999 Fall Simulation Interoperability Workshop*, 99F-SIW-001, November 1999.
- [21] J. Huang, J. Hao, X. Zhao, and G. Jianxing. "Basic Study on RTI Interoperability", in: *Proceedings of the 2010 Second International Conference on Computer Modeling and Simulation*, 2010.
- [22] U.S. Army Research Laboratory. DoD Supercomputing Resource Center. Harold Overview. <http://www.arl.hpc.mil/Systems/harold.html> ; accessed: 12/20/2010.
- [23] Open MPI. Open Source High Performance Computing. <http://www.open-mpi.org/>; accessed: 1/10/2011.
- [24] Soldier Radio Waveform (SRW) 1.0C Waveform Design Specification (WDS) ITT DOCUMENT # 8224935 CDRL No. A001 Rev. H, Jan 2009
- [25] WNaN Network Architecture Design Document, BBN DOCUMENT # WNaN2-SYAD-02-01, 13 August 2009
- [26] <http://www.alionscience.com/Technologies/Wireless-Spectrum-Management/TIREM/>
- [27] U.S. Army Research Laboratory. DoD Supercomputing Resource Center. Guide to the PBS Queueing System. <http://www.arl.hpc.mil/SystemSoftware/pbs.html>; accessed: 12/20/2010.

Author Biographies

BRIAN RIVERA is the chief of ARL's Tactical Network Assurance Branch. He is also the Simulation Lead for the MNMI. He received a Doctor of Philosophy in Electrical and Computer Engineering from the Georgia Institute of Technology in 2007. His research interests include secure mobile wireless networking, intrusion detection for mobile ad hoc networks and network science where he is examining the interactions between cognitive and communication networks.

JOHN A. TUFAROLO is an NDIA Certified Modeling and Simulation Professional and is a consulting engineer to MITRE from Pragmeering Corporation. He has more than 24 years of experience in simulation system and simulation technology design, development, integration, and testing supporting development, testing, and application of military training and analysis simulations. Mr. Tufarolo served as Chair, Vice-Chair, and Information Director for the Association for Computing Machinery (ACM) Special Interest Group on Simulation (SIGSIM), and is a member of the ACM and SIGSIM. Mr. Tufarolo has a B.S in Electrical Engineering from Drexel University, and a M.S. in Systems Engineering from George Mason University.

GARY COMPARETTO is a Senior Principal Engineer with the MITRE Corporation. He received the B.S. in E.E. from Bucknell University, M.S. in Nuclear Engineering from Penn State University and M.S. in E.E. from Georgia Tech. He has extensive

technical expertise in performance modeling and simulation of tactical mobile ad-hoc terrestrial and SATCOM networks. Mr. Comparetto has published over 40 technical papers and has authored chapters in two graduate level text books. His primary interests include communications network performance, SATCOM link analysis, and waveform design.

VINAY LAKSHMINARAYAN is a Senior Engineer with the MITRE Corporation, with more than 15 years of experience supporting M&S of communications networks for several US Army programs and commercial industry clients. He received a Bachelor's degree in Electronics & Communications Engineering from the University of Mysore, India and a Master's degree in Computer Science from the University of Missouri. His interests include M&S of commercial and DoD communications networks.

MOHAMMAD MIRHAKKAK is a Principal Engineer with the MITRE Corporation where he has worked since 1990. He received the B.S. and M.S. degrees in electrical engineering from Tehran University, Tehran, Iran, and the Doctor of Science (D.Sc.) degree in computer science from The George Washington University, Washington, DC. In recent years, he has been actively involved in Modeling and Simulation (M&S) of network protocols for use in DoD network environments.

ERNEST H. PAGE is a member of the technical staff of The MITRE Corporation. He received the Ph.D. in Computer Science from Virginia Tech in 1994. The author of over 40 papers on distributed

simulation and modeling methodology, he currently serves on the editorial boards of Society for Computer Simulation International (SCS) *Simulation*, *SCS Journal of Defense Modeling and Simulation*, and the *Journal of Simulation*. He has held the position of Secretary (1995-1997), Vice Chair (1997-1999) and Chair (1999-2001) of the Association for Computing Machinery (ACM) Special Interest Group on Simulation (SIGSIM).

NANCY SCHULT is the Associate Department Head of the Communications Network Engineering and Analysis Department at MITRE, in McLean, Virginia. She has a Ph.D. in Computer Science from the University of Virginia. Her current research interests include network protocols for wired and wireless networks, quality of service issues in mobile ad hoc networks, and performance analysis of large-scale wired/wireless networks.

DAVID YOO is a Lead Engineer with the Army Programs Directorate of the MITRE Corporation. He received the B.S. degree in computer engineering from The College of New Jersey (formerly Trenton State College), Ewing, NJ, in 2002. He received the M.S. and Ph.D. degrees, both from Drexel University, Philadelphia, PA, in 2005 and 2006, respectively.

Table 1: Offered Traffic Loads for the 9-9-9 and 42-75-42 Configurations.

Run #	Model	Configuration	# nodes	SA packet size (bytes)	SA period (sec)	Unicast packet size (bytes)	Unicast period (sec)
1	WNAN	Transit Federation	75	56	3.1		1
2	SRW	Source & Destination Federations	42	30	5	1024	0.1
3							0.05
4							0.01
5	WNAN	Transit Federation	75	56	3.1		1
6	SRW	Source & Destination Federations	42	30	0.5	1024	0.1
7							0.05
8							0.01
9	WNAN	Transit Federation	75	56	3.1		1
10	SRW	Source & Destination Federations	42	30	0.3	1024	0.1
11							0.05
12							0.01
13	WNAN	Transit Federation	9	56	1		1
14	SRW	Source & Destination Federations	9	30	1	1024	0.1
15							0.05
16							0.01
17	WNAN	Transit Federation	9	56	0.06		1
18	SRW	Source & Destination Federations	9	30	0.4	1024	0.1
19							0.05
20							0.01
21	WNAN	Transit Federation	9	56	0.03		1
22	SRW	Source & Destination Federations	9	30	0.07	1024	0.1
23							0.05
24							0.01