

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this 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 this 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 Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 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 any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE				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)				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					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

Tactical Network Integration Test Framework

Lorraine Prior, Carl Fossa, David Ward, Jun Sun
 MIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02420
 Patrick Boehm (OPNET), Edward Kuczynski (DAG), John Cain (Akibia)
 Thomas Mak (PM WIN-T, US Army)

Approved by
 Army PAO
 25 5/16/11

Abstract— Mobile Ad-Hoc Networks (MANETs) will play a significant role in future tactical military networks. These tactical networks are required to support military operations and communications on-the-move in the mobile ad-hoc environment. They are also characterized by frequent changes in network topology, time varying bandwidth, interference and intermittent link blockage. The self-forming and self-healing nature of MANETs is therefore advantageous in a tactical military network.

The DoD's evolving tactical network architecture is expected to extend Global Information Grid (GIG) services to the forward edge. To accomplish this, it will be necessary to interconnect heterogeneous tactical networks. The interconnection of these networks requires a scalable routing architecture with routing protocols that can support the frequent topological changes of a tactical environment.

The potential challenges associated with connecting heterogeneous mobile, tactical networks motivated the development of the Tactical Network Integration Test Framework at MIT Lincoln Laboratory (LL). MIT LL has developed a simulation and emulation environment suitable to study network and routing architecture configuration at the "seams" between heterogeneous tactical networks. The Tactical Network Integration Test Framework is comprised of three separate test environments that we term simulation, high fidelity emulation, and scalable emulation. In all three of these test environments the "lower tier" of the network is a Wideband Networking Waveform (WNW) cloud that is represented in an OPNET model. The difference between the three test environments is the method used to represent the backbone or "upper tier" of the network. In the simulation environment, this upper tier is constructed with an OPNET model that is representative of a higher tier backbone network. In the high fidelity emulation environment, a small number of upper tier nodes are represented by actual networking equipment such as commercial routers and HAIPE devices in a laboratory environment. In the scalable emulation environment, a larger number of upper tier nodes are represented by Linux machines that are configured with Quagga software, actual algorithms and protocols. In this paper we will describe each of the environments and show that the combination of these three test environments allows us to look at both high fidelity interoperability on a small number of nodes as well as the impact of scalability given the general-purpose emulation environment. The validation of these three environments which was performed by comparing the results from similar tests that were performed on each of the three test environments to ensure that both

simulation and scalable emulation results match high fidelity results will be presented.

Keywords—MANET; tactical networks; OSPF; internetworking

I. INTRODUCTION

Tactical networks will provide valuable data from the forward edge while also consuming services from the DoD's Global Information Grid (GIG). The general assumption is that techniques and protocols currently used to interconnect fixed networks in the Internet will provide an acceptable internetworking solution in a mobile tactical environment. However there are many challenges associated with the internetworking of mobile, secure and heterogeneous tactical networks. The Tactical Network Integration Test Framework was established to study network and routing architecture configurations at the "seams" where these tactical networks will be connected. A primary goal of this framework was to provide a scalable environment in which to conduct repeatable experiments where the network impact imposed by mobile tactical networks challenges could be measured. These potential challenges include rapid changes in topology, links with variable qualities as well as intermittent link outages.

Figure 1 presents an overview of our approach to studying this network integration with simulation and emulation. The three framework test environments are complementary and each provides a capability that is advantageous to the overall study. Simulation provides a portable and scalable environment that allows the quick construction of network models. However given a tactical network the desired level of fidelity for the individual network components may not be available. The high fidelity environment, which is constructed with real commercial networking equipment, provides a performance baseline as well as insight into the configuration and operation of the actual protocols and elements in the network. But the high fidelity environment in this framework is limited to a set of four emulated nodes. Lastly, the scalable environment presents a general-purpose hardware emulation platform that is scalable and employs actual algorithms and protocols. However, the scalable environment cannot be constructed as quickly as the simulation. The approach taken in our framework is to execute the same network scenario with four

upper tier network nodes in all three environments. Network and routing architectures are tested in simulation then if acceptable, constructed in the emulation environments. The behavior and performance results from each environment are compared to ensure there was consistency in the experimentation and results before increasing the network size and executing the scenario again in both simulation and scalable emulation. Each of these three test environments will be described in further detail in Section II. The focus of this initial phase of study was to investigate the scalability of the OSPF routing protocol in the tactical environment and to observe the network performance at the seams between the upper and lower tier networks. The details of the scenarios executed and subsequent results will be presented in Section III. A summary and recommendations for future study are given in Section IV.




Figure 1. Network Integration Test Framework Overview

II. TEST ENVIRONMENTS

In all three environments the lower tier network is represented by an OPNET simulation model characteristic of a WNW network. The upper tier network topology is the same in all three environments. However, the method employed to implement the upper tier network is different across the three environments. In the simulation environment the upper tier is represented with an OPNET simulation. In the high fidelity environment, the upper tier consists of four nodes constructed with commercial-off-the-shelf networking equipment. In the scalable environment, the upper tier is comprised of algorithms and protocols running on general purpose processors.

A. Simulation Environment

The simulation test environment used OPNET Modeler to represent both upper and lower tier tactical networks. All simulation runs were conducted using OPNET Modeler 14.5 PL8.

The upper tier network model consisted of simulated satellite and ground-to-ground Line-of-Sight (LoS) data links. The satellite communications link was simulated with full mesh connectivity between all nodes. For both the upper and lower tier each simulated node was comprised of a secure or red local area router and a black wide area router. In each node, the two routers were separated by a simulated encryption device, which modeled the effects of encryption by adding/removing padding bits to packets traversing the

encryption device. The lower tier simulation is represented by an OPNET simulation model characteristic of a WNW network.

B. High Fidelity Environment

The upper tier network is represented in the high fidelity test environment by four nodes constructed with commercial networking components that are representative of the equipment found in a mobile tactical node. Here each of the four nodes was configured with a wide area network (WAN) router that was connected to an inline encryption device, which in turn was connected to local area network (LAN) router. The WAN router on each node was a Juniper 6350 router. This router's interfaces provided connectivity to a WAN emulated satellite and Line of Sight (LoS) networks. The inline encryptor, which encrypted the LAN traffic for transport across the WAN, was a General Dynamics KG-175A. The LAN router was a Cisco ASR1004 router. A second interface on this LAN router was connected to the lower tier simulation through the OPNET Modeler's System-in-the-Loop (SITL) feature.

The WAN satellite system was emulated with Anuec Ethernet Network Emulators. These units are inline devices that emulate real network conditions by injecting packet delay, packet delay variation, bit errors and packet drops into the traffic flows that traverse the unit. An Anuec emulator was attached to one interface on each WAN router and then to an Ethernet switch to create the satellite emulation. Each Anuec emulator was configured with a 125-millisecond packet delay to emulate the delay experienced in a ground to space satellite link. The round trip delay between any two mobile ground nodes was 500 milliseconds.

Three of the four upper tier nodes had an emulated (LoS) data link. This network was emulated using the PPPoE protocol with the credit based flow control mechanism as documented in RFC4938. The LoS emulation was comprised of Cisco 3845 and 2811 routers that acted as PPPoE servers, and Linux systems were configured with an open source implementation of the PPPoE client software based on Roaring Penguin's rpppoe package. The PPPoE environment was configured as a full mesh between the three LoS nodes. Each PPPoE session was provisioned with a throughput rate of 16 Mbps. Lastly, Cisco's Dynamic Multipoint VPN (DMVPN) feature was deployed on the upper tier LAN routers. The DMVPN provides a hub and spoke tunnel structure that allows the four LANs to appear as a contiguous network.

C. Scalable Emulation

In the scalable emulation environment, the upper tier network was emulated using open-source software running on a local grid of commodity servers, known here at MIT LL as the Network Link and Emulation Testbed (NLET). The available capacity of the NLET allowed us to run experiments with a network of 48 upper tier nodes in emulation. The lower tier network was simulated in OPNET Modeler; simulated WNW nodes interfaced with emulated upper tier nodes through OPNET Modeler's System-in-the-Loop feature.

- Operating System

The emulation was based on the Fedora 14 distribution of Linux. It utilized Linux containers, a new feature in the kernel, which in particular allows multiple networking stacks to exist under the same operating system instance. Processes that are launched in a container use an isolated networking stack associated with that container, which has its own routing table and neighbor table. This capability enabled us to emulate several complete upper tier nodes on a single server. (In contrast, hypervisor-based virtualization isolates entire operating system instances running on the same server, which each has an independent network stack. However, applications that require real-time performance such as our link modeling software would be adversely affected, due to clock skew and processor overhead from context switching.)

Each node in the upper tier network was emulated using two containers: one for the WAN colorless traffic, and one for the LAN red traffic. Each colorless container had one Ethernet interface that connected it to all of the other colorless containers over a virtual LAN (VLAN), which was used to model the upper tier wireless links. Each red container had one Ethernet interface that connected the upper tier red router to the red router of a corresponding lower tier WNW node in simulation, using a unique VLAN for each of these connections. These VLANs are provided by a configurable non-blocking Gigabit Ethernet switch, which interconnects the servers in the NLET.

Additional virtual network interfaces existed inside both red and colorless containers, including a virtual Ethernet pair device (VETH) that connected a red container to a corresponding black container. Figure 2 shows all of the network interfaces inside both containers of an emulated upper tier node, and how they were connected.




Figure 2. Network interfaces in an upper tier node in the Scalable Emulation

- Link Modeling

Extensible Mobile Ad-hoc Network Emulator (EMANE)[1] is an open-source application that emulates

wireless network connections between nodes using an underlying wired network. In its basic mode of operation (with the "Comm Effects Model"), it can impose arbitrary constraints on delay, delay variation, loss, duplication, and bandwidth between each pair of nodes as packets are sent through the network. Alternatively, it can use specific physical- and link-layer models, along with node movement data during an experiment, to provide higher-fidelity emulation of radios and environmental effects. EMANE version 0.6.2 was used in modeling both the satellite and LoS data links in the upper tier network.

Two data links were modeled in the upper tier network, a Line of Sight link and a satellite link. The LoS model operates as a Point-to-Point Protocol over Ethernet (PPPoE) client on each node, in order to emulate the credit-based flow control used in some radios. An open source PPPoE server implementation, rp-pppoe, provided the necessary interface to the model to allow it to send and receive network traffic from the operating system (rather than connecting the model to a commercial router appliance).

The satellite system was emulated with the EMANE Comm Effects Model. The model was configured to use a large delay characteristic of a satellite link, as well as bandwidth constraints to represent traffic shaping for different classes of nodes.

- Routing

The Quagga routing suite (version 0.99.17)[2] provided routing functionality on both the colorless and red networks. The mainstream Quagga release provides OSPFv2 routing, which was used throughout the upper tier red network in this study. The Mobile Routing branch of Quagga, maintained by Boeing, extends Quagga's standard OSPFv3 routing to support extensions for Address Families and MANET Designated Routers; this was used for the upper tier colorless WAN network.

- Encryption

The High Assurance Internet Protocol Encryptor (HAIPE) that bridges the colorless network and red network was modeled with a full mesh of generic IPsec Encapsulating Security Payload (ESP) tunnels, configured in the operating system on the red container with a freely available cipher (AES-CBC) and hash function (HMAC-SHA1). Despite the weaker encryption, this extends the packet header length comparably to a HAIPE device.

- Red Upper Tier Network Tunneling

The upper tier network uses Cisco's proprietary Dynamic Multipoint Virtual Private Network (DMVPN) feature to establish tunneling behind HAIPE devices. Each node in a DMVPN shares the same subnet identifier, abstracting the underlying logical topology of the network.

DMVPN was modeled using a full mesh of GRE tunnels configured in the operating system between each upper tier red container. Since OSPFv2 strictly follows the IP

subnetting model[3], each GRE tunnel was assigned a unique subnet identifier. This resulted in an artificial increase in the number of routes announced by the upper tier red routers, which grew exponentially with the size of the upper tier network. An alternative implementation for future studies is being considered that utilizes the IETF standard Next Hop Resolution Protocol (NHRP) on which DMVPN is based.

III. SCENARIOS AND RESULTS

The scalability of OSPF in a mobile tactical environment was our initial focus. Our experiment includes two mobile tactical networks configured in a single OSPF area. An upper tier network provides backbone like capacity to a lower tier network that operates at the edge. The main parameters that can be varied are network size, network topology and mobility patterns.

In our initial scenario the upper tier network contained only 4 nodes corresponding to the number of real nodes in the high fidelity emulation environment. The lower tier network, which was represented by the WNW OPNET simulation in each test environment, contained 40 nodes. These nodes were arranged in a grid-like network where each node, except the corner nodes have a neighboring node on the right, left, top and bottom. The distance between two adjacent nodes is approximately 1 km. During this initial phase of experimentation the upper and lower tier networks were interconnected on the red or secure LANs. The four upper tier nodes were connected to four distinct nodes in the lower tier network.

In this initial scenario the upper tier nodes were stationary and a pre-planned mobility pattern was executed in the lower tier simulation for each test run performed. For the pre-planned mobility pattern, we created four trajectories: DownLeftVar, DownRightVar, UpLeftVar, and UpRightVar. The detailed description of each trajectory is given in Table 1. For example, following UpRightVar trajectory, a lower tier node will stay stationary for 7 minutes. From the 7th minute to the 8th minute, the node will move at a speed of 46.6 miles per hour in the direction of (1, 1). From the 8th minute to the end of the simulation, the node will stay stationary.

Table 1. Simulated WNW node trajectories

Time:	0-7 minutes	7-8 minutes	8-15 minutes
DownLeftVar	stationary	direction (-1, -1); speed: 46.6 mi/hr	stationary
DownRightVar	stationary	direction (-1, 1); speed: 46.6 mi/hr	stationary
UpLeftVar	stationary	direction (1, -1); speed: 46.6 mi/hr	stationary
UpRightVar	stationary	direction (1, 1); speed: 46.6 mi/hr	stationary

Since the lower tier nodes are arranged in a square grid, we divided the nodes into four groups based on the quadrant the node is in. Nodes in the upper right quadrant will follow the trajectory UpRightVar. Similarly, nodes in the lower left quadrant will follow the trajectory DownLeftVar. Each

quadrant contained one of the four connections to the upper tier network. This pre-planned trajectory model allows us to check the impact of mobility on network performance in a predictable way.

The primary performance metrics collected were OSPF overhead traffic and ICMP Echo (or "ping") responsiveness between pairs of nodes in the lower tier network.

- OSPF traffic rate across gateway links: We measured the amount of OSPF traffic flow across the gateway link connecting an upper tier node and a lower tier node. In both emulation platforms, this measurement was taken using the Wireshark open-source protocol analyzer to capture OSPF packets during the test run; for simulation, the built-in statistic reporting was used. The OSPF traffic flow includes traffic both coming into and going out of the upper tier network.
- Percentage of ICMP Echo round-trip completions: The percentage of ICMP Echo packets that successfully complete a round trip between a lower tier network source and destination is measured, in order to approximate the percentage of time with end-to-end reachability across the network.

Consecutive test runs provided consistent results within each of the environments. Across the three environments there appeared to be a difference in the amount of OSPF initialization traffic. Inspection of packet captures revealed the discrepancy. While the synchronization of these intervals is consistent between the upper and lower tier routers in the simulation environment, they realistically depend on when OSPF routing starts on each network interface of the router, which we did not attempt to control in the emulation environments. If after starting the emulation, an OSPF Hello packet exchange is completed over the monitored gateway link, before any other gateway link, then all OSPF Link State Advertisements (LSAs) on both routers held would need to be sent across this monitored gateway link. However, if an OSPF Hello packet exchange is completed first on a different gateway link, then some or all LSAs may have already been received at the monitored gateway router through a different path before the OSPF Hello packet exchange completes on this monitored gateway link. The OSPF overhead level was comparable across the environments after the initialization process. This initialization difference is an example of the significance of operational validation[4][5] with real assets. Without this validation the implementation difference in router synchronization between simulation and emulation would not have been apparent.

A second scenario was performed with the same mobility pattern and with 80 nodes in the lower tier network. To compare OSPF overhead results between 40 and 80 lower tier network nodes, the overhead per second for a test run was summed cumulatively. The results for 40 and 80 scenarios for simulation are shown in Figure 3, high fidelity in Figure 4 and scalable in Figure 5.




Figure 3. Cumulative routing overhead in simulation after 100 seconds (40/80 Lower and 4 Upper Tier nodes)




Figure 4. Cumulative routing overhead in high fidelity after 100 seconds (40/80 Lower and 4 Upper Tier nodes)




Figure 5. Cumulative routing overhead in scalable after 100 seconds (40/80 Lower and 4 Upper Tier nodes)

The plots start 100 seconds into the scenario eliminating the initial OSPF overhead traffic from the cumulative summation. Also, note that there is an increase in the total OSPF traffic sent rate between 440-500 seconds. This increase is due to the partition of the lower tier network with pre-planned mobility. After the lower tier network partitioned, OSPF requires additional overhead traffic to reestablish routes through the upper tier network. These charts show the additional overhead incurred with the increased number of lower tier simulated nodes given the simple mobility pattern employed. A comparison of these charts validates the timing of events in the three environments. Initial convergence between the lower

and upper tier networks is apparent at 225 second in each chart and the OSPF traffic increase due to mobility is consistent across the three environments at 450 seconds.

In both the 40 and 80 node scenarios an ICMP Echo was initiated on a lower tier node in one quadrant and sent to a node in a second quadrant. The ICMP Echo round trip times as measured by the lower tier OPNET simulation in the high fidelity emulation are shown in Figure 6, and for the scalable emulation in Figure 7. When the simulated lower tier nodes begin to move 7 minutes or 420 seconds into the scenario, the lower tier network becomes fragmented and the upper tier provides a failover path. Note that for both 40 and 80 nodes in each environment, the round-trip time once the failover path is activated is substantially shorter as the ICMP packets transit the upper tier network. Also 420 seconds into the scenario as the mobility begins, the ICMP packets are not delivered for 36 seconds in the high fidelity 40 node run and 31 seconds for the 80 node run. Similarly, in the scalable environment packets are not delivered after mobility begins for 36 seconds in the 40 node run and 26 in the 80 node run.




Figure 6. ICMP Echo round-trip time in High Fidelity Emulation




Figure 7. ICMP Echo round-trip time in Scalable Emulation

A third scenario where the upper tier network was scaled to 48 nodes and the lower tier contained 80 nodes was executed in the scalable environment. In Figure 8, the peak bit rates attributed to OSPF overhead are comparable to the previous scenarios with 4 node upper tier networks, however the base of the base of the peaks are broader. These results will be used for comparison with a simulation of the same scale.




Figure 8. Instantaneous routing overhead in the Scalable Emulation after 100 seconds (80 Lower / 48 Upper Tier nodes)

IV. SUMMARY AND FUTURE WORK

The approach taken in this framework was to execute and observe a scenario in three different test environments, namely, simulation, high fidelity emulation and scalable emulation. The comparison of behavior and performance results from each of the environments allowed us to verify and validate the configuration and execution of the network and routing architectures under study with a small number of nodes before scaling to a larger number of nodes.

Each environment provided a capability that was advantageous. The simulation environment provided a tool to quickly configure and scale the network architecture under test as well as to introduce a simple mobility pattern in the lower tier network. The high fidelity emulation environment allowed us to observe and quantify the operation of real networking elements in our study. This information was used to verify models and algorithms within the simulation and scalable environments, respectively. Additionally, this insight into the operations of actual assets and protocol implementations provides information that can be transferred into the design and configuration of future simulation models as well as the open source algorithms implemented in the scalable emulation to provide more fidelity and interoperability in these environments.

In our experimentation, results that depicted a difference in the initial OSPF overhead traffic across the three environments led us to the realization of an operational difference between our simulation and emulation environments. This difference occurred due to the fact that OSPF timers are synchronized across routers in simulation but not synchronized across routers in the high fidelity and scalable emulation environments. While this is an important configuration parameter, this difference did not impact our focus of study. Traffic results also showed consistency across the three environments with respect to the timing of OSPF traffic bursts. The traffic peaks, which were due to mobility consistently, occurred at the 420-second interval in each environment.

The focus of the initial phase of study was to investigate the scalability of the OSPF routing protocol in the tactical environment and to observe the network performance at the

seams between the upper and lower tier networks. The initial phase of experimentation that we performed was configured with 4 upper tier nodes and 40 lower tier nodes. The number of upper tier nodes was chosen based upon the availability of real assets used in the high fidelity emulation environment. The purpose of this small scale scenario was to validate the network behavior and performance results across the simulation, high fidelity and scalable emulation environments. After verifying that the OSPF overhead and ICMP results in this small scale scenario were consistent across the three environments, the lower tier network size in the scenario was scaled in size to 80 nodes. Here again the OSPF overhead and ICMP time results were consistent across the three environments.

Access to the real networking assets provided the capability to understand the functionality of individual network elements as well as to construct an environment in which we were able to generate results on a small scale for comparison across simulation and emulation. Upon understanding the differences and consistencies in these results we scaled our scenarios and again validated the operational results across the applicable environments.

Future work will involve complex mobility patterns in each of the environments. The current scenarios have 4 gateways or connections between the upper and lower tier networks so investigation into increasing the number of gateways needs to be investigated. Also, implementing different routing protocols at the network seam between the upper and lower tier networks, such as BGP, would be desirable. There are areas that are important to the tactical environment such as the impact of deploying encryptors that dynamically discover peers versus encryptors that are statically configured. Another area involves the various communication waveforms that are proposed for the tactical environment that are implemented with radios that can communicate link metrics to the WAN router to which they are attached providing a means to calculate dynamic cost metrics in the implemented network routing protocol. These are a few of the potential areas of exploration and testing that could be conducted in the framework.

V. REFERENCES

- [1] Patel, Kaushik B., and Galgano, Steven M., "Emulation Experimentation Using the Extendable Mobile Ad-hoc Emulator," CenGen, Inc., Columbia, MD [Online], URL: <http://labs.cengen.com/emane/doc/training.html>
- [2] Quagga Software Routing Suite, GPL licensed IPv4/IPv6 routing software, [Online], URL: <http://www.quagga.net>.
- [3] Moy, J.T., "OSPF: Anatomy of an Internet Routing Protocol," Addison-Wesley 1998.
- [4] Kurkowski, S., Camp, T., and Colagrosso, M., "MANET Simulation Studies: The Incredibles," Colorado School of Mines, Golden, Colorado. Mobile Computing and Communications Review, Volume 9, Number 4.
- [5] Sargent, R., "Verification, validation, and accreditation of simulation models." In Proceedings of the 32nd Conference on Winter Simulation, pages 50-59, 2000.