This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Suri, Niranjana; Nilsson, Jan; Hansson, Anders; Sterner, Ulf; Marcus, Kelvin; Misirlioglu, Levent; Hauge, Mariann; Peuhkuri, Markus; Buchin, Boyd; In'T Velt, Ronald; Breedy, Maggie
The angloval tactical military scenario and experimentation environment

Published in:
2018 International Conference on Military Communications and Information Systems, ICMCIS 2018

DOI:
[10.1109/ICMCIS.2018.8398729](https://doi.org/10.1109/ICMCIS.2018.8398729)

Published: 27/06/2018

Document Version
Peer reviewed version

Please cite the original version:
Suri, N., Nilsson, J., Hansson, A., Sterner, U., Marcus, K., Misirlioglu, L., Hauge, M., Peuhkuri, M., Buchin, B., In'T Velt, R., & Breedy, M. (2018). The angloval tactical military scenario and experimentation environment. In *2018 International Conference on Military Communications and Information Systems, ICMCIS 2018* (pp. 1-8). IEEE. <https://doi.org/10.1109/ICMCIS.2018.8398729>

The Angloval Tactical Military Scenario and Experimentation Environment

Niranjan Suri
US Army Research Laboratory &
Florida Institute for Human and
Machine Cognition (IHMC)
Adelphi, MD USA
niranjan.suri.civ@mail.mil

Jan Nilsson
Swedish Defence Research
Agency (FOI)
Linköping, Sweden
jan.nilsson@foi.se

Anders Hansson
Swedish Defence Research
Agency (FOI)
Linköping, Sweden
anders.hansson@foi.se

Ulf Sterner
Swedish Defence Research
Agency (FOI)
Linköping, Sweden
ulf.sterner@foi.se

Kelvin Marcus
US Army Research Laboratory
Adelphi, MD USA
kelvin.m.marcus.civ@mail.mil

Levent Mısırlıoğlu
MilSOFT Software Technologies
Inc.,
Istanbul, Turkey
lmisirliloglu@milsoft.com.tr

Mariann Hauge
Norwegian Defence Research
Establishment (FFI)
Kjeller, Norway
mariann.hauge@ffi.no

Markus Peuhkuri
Aalto University
Espoo, Finland
markus.peuhkuri@aalto.fi

Boyd Buchin
Rohde and Schwarz
Munich, Germany
boyd.buchin@rohde-schwarz.com

Ronald in 't Velt
TNO,
The Hague, The Netherlands
ronald.intvelt@tno.nl

Maggie Breedy
Florida Institute for Human and
Machine Cognition (IHMC)
Pensacola, FL USA
mbreedy@ihmc.us

Abstract— The Anglova scenario is designed to support experimentation with tactical networking environments and provides node mobility and network connectivity for a realistic battalion-sized military operation consisting of three vignettes, including the deployment of armored vehicles, surveillance of the maritime domain and an urban operation with a naval component. Altogether, the scenario includes 283 nodes and lasts over four hours with detailed mobility and pathloss data for each node. Also included with the scenario are radio models for the different radios that would be part of a heterogeneous network. The scenario has been developed by the NATO IST-124 Research Task Group and released into the public domain in order to facilitate experimentation with networking protocols and algorithms by the community at large. While primarily designed for the Extendable Mobile Ad-hoc Network Emulator (EMANE), the scenario can also be adapted to other experimentation environments.

Keywords—military scenario; tactical networking; heterogeneous networking; network emulation; network experimentation

I. INTRODUCTION

Experimentation is often critical to evaluating and comparing new algorithms and protocols in the networking domain. Simulation- and emulation-based evaluation are two of the most popular approaches to experimentation, with both of them requiring scenarios to drive their behavior. While numerous “toy” scenarios exist, these do not provide sufficient fidelity to conduct experiments, especially in a militarily-relevant context. A related problem is that militarily-relevant

scenarios are often not publicly available, thereby hindering the ability for different researchers to use a common baseline in order to be able to compare results with each other. To address this problem, the NATO IST-124 Research Task Group (RTG) on “Heterogeneous Tactical Networks – Improving Connectivity and Network Efficiency” developed the Anglova scenario to facilitate efficient research collaboration among the partners in the group and released it into the public domain. This paper provides an overview of this scenario and characterizes the mobility and communications patterns afforded by the scenario. The paper also describes two different deployment models of the scenario, along with the software tools that manage the experimentation lifecycle. It is the hope that the scenario and accompanying tools will encourage experimentation and analysis while at the same time providing a common baseline to facilitate comparison of results. The objective of this paper is not to describe a specific experiment, but to describe the scenario in sufficient detail for others to use it as part of their experiments. The IST-124 group hopes that the scenario will be a valuable contribution to the research community at large, by providing a realistic baseline for experimentation.

II. MOTIVATION AND SCENARIO OVERVIEW

The NATO IST-124 RTG’s primary objectives are to investigate approaches and mechanisms to improve connectivity and network efficiency in heterogeneous tactical networks. To this end, exploration and experimental evaluation of alternatives is important. Five different approaches to experimentation were considered: simulation, emulation,

laboratory evaluation with actual networking hardware (e.g., tactical radios), limited field experimentation, and finally live military exercises. Each of these approaches has its relative advantages and disadvantages, particularly at different states of development and maturity of the ideas, algorithms, and implementations (more details in [1]). For the NATO working group, emulation-based experimentation was considered to be the best compromise as it allows actual software components that would be used in the field to be evaluated. While it does involve additional hardware costs compared to simulation, it does not involve the exorbitant costs of the other alternatives that involve actual radio equipment and potentially military personnel. It has also been shown [2] that it is difficult to do credible comparison of research results if the results are produced with different simulators. This is also an argument for choosing emulation-based experimentation focusing primarily on one emulation environment. However, any emulation-based experiment (or even simulation-based experiments) needs a scenario to drive the behavior of the nodes that are part of the experiment. In particular, the scenario must define the mobility of the nodes over time, the characteristics of the communication links between the nodes over time, and the characteristics of the radios that are providing the network links between the nodes. Furthermore, the scenario must be sufficiently large and sufficiently diverse to enable meaningful experimentation. Finally, the scenario must also define the data or message traffic that is to be exchanged between the nodes.

Another motivation for the development of the scenario was to be able to release it to the public as open source data. While there are multiple scenarios in existence, they are typically protected by the nations that developed them and are difficult to obtain access to and even more difficult to share across different nations and the broader community including academics and other researchers. Furthermore, the scenarios that were (somewhat) available were either toy scenarios or synthesized scenarios on the order of 10s of nodes that did not include the desired heterogeneity at the network level.

To address the above limitations, we decided to develop the Anglova scenario, which depicts an operation conducted by an army battalion task group and a naval task group. The scenario commences upon reception of reconnaissance data that raises an alert about an attack by insurgent forces against coalition forces in the operational zone. The scenario's tactical domain is located in the fictitious area of Fieldmont in Anglova, where the Coalition HQ (CHQ) of the Military Contingent (MC) is based. The scenario is divided into three vignettes, with the first vignette focusing on the intelligence preparation of the battlefield, the second vignette on the deployment of the battalion and movement to the objective as well as surveillance by the naval task group, and the third vignette on an urban operation including a medevac by a helicopter to a naval ship. Overall, the scenario involves 283 nodes and runs for approximately four hours (it is possible to run subsets of the scenario as desired). A recent example of the successful use of the scenario to evaluate the performance of OLSR is in [3]. While the scenario has been successfully used by the IST-124 group for experimentation, the objective of this paper is to describe the scenario so that other researchers in the broader

community may also be able to use it for their own experimentation.

III. SCENARIO CHARACTERIZATION

Characterizing the Anglova scenario is important in order for experimenters to understand and anticipate expected behavior at the RF level. The following subsections focus on the second vignette (which has the most complicated dynamics due to the motion and terrain), the naval component of the scenario, and the third vignette, which is the urban operation. The last subsection describes the radio models that were developed for the scenario. Vignette 1 is not fully implemented in the emulation testbed yet, thus we describe Vignette 2 and 3 in this paper.

A. Vignette 2: Troop Deployment

The second vignette covers the deployment of the coalition forces, a battalion consisting of six companies, into the operational zone. The task for the battalion is to stage an attack against a hostile force that is advancing into the operational zone. The area selected for the troop deployment vignette primarily consists of hilly terrain covered by forests. The troop mobility pattern was sourced directly from a NATO exercise, and is characterized by movements over a rectangular area of 13 km by 33 km mainly utilizing large and small roads. The speed of the vehicles varies, with speeds up to 60 km/h on the main roads. The battalion starts by moving in a single column on one of the main roads from the CHQ (see Figure 1; the CHQ is indicated with a cyan-colored star). After about 10 km, the battalion splits up over two main roads and after about 25 km splits up further onto many paths grouped in companies. Towards the end, the battalion finally splits up to the level of platoons. Altogether, the Vignette 2 mobility pattern is 7800 seconds (130 minutes) long [1].

The battalion consists of six companies: four mechanized companies with 24 vehicles each, one command and artillery company with 22 vehicles, and one support and supply company with 39 vehicles. Together, there are 157 vehicles, with each of them being a network node. In addition, an *Unmanned Aerial Vehicle* (UAV) node was also added to this vignette (with a trajectory in the shape of two connected green circles in the figure) – it can act as a relay and provide persistent surveillance capabilities.

Figure 1: Illustration of the movements (direction is from left to right (north to the south)) in the battalion.

To calculate the path loss between the nodes in Vignette 2, a Uniform Theory of Diffraction (UTD) propagation model by

Holm[4] was used. The model uses a digitized terrain database for the calculation to include large-scale fading effects. In Figure 2 to Figure 4, the path loss is used to calculate the theoretical maximum connectivity in the 157-node network for the Vignette 2 deployment phase. This is the theoretical best case connectivity that can be achieved and serves as a benchmark that can be used to compare the performance of medium access protocols (MAC) and routing protocols deployed in the scenario. Note, the UAV and the CHQ were not included in this analysis of the network. Three different waveforms to connect the nodes were investigated. The connectivity is illustrated by showing how the fraction of nodes at h hops distance from each other varies over time (however, note that routing protocols may use something other than hop count as their metric). The average of h is taken over all nodes in the network. The hop distance is theoretically calculated, with the assumption that there is a communication link between two nodes if the path loss value is less than a system gain G_{sys} that varies for the waveforms. The three investigated waveforms are: 25 kHz, 17.5 kbit/s narrowband (NB) waveform with $G_{sys}=156$ dB; 250 kHz, 175 kbit/s medium band (MB) waveform with $G_{sys}=146$ dB; and 1.25 MHz, 875 kbit/s wide band (WB) waveform with $G_{sys}=139$ dB.

Figure 2: The NB 25 kHz waveform at the 50 MHz frequency band.

Figure 3: The MB 250 kHz waveform at the 300 MHz frequency band.

Figure 4: The WB 1.25 MHz waveform in the 300 MHz frequency band.

As can be seen in Figure 2 the 25 kHz waveform at 50 MHz keeps the network connected for the whole deployment phase and no more than two hops are required. For the other two waveforms at 300MHz (Figure 3 and Figure 4), the network is not always fully connected as there are a number of unreachable nodes at various times. After around 2000 seconds, the network start to be stretched out and the number of hops increases. Towards the end of the vignette, the network is fragmented with a few nodes behind the main part of the battalion that cannot be reached using the WB waveform. However, within the main part of the battalion, almost all nodes can be reached with a maximum of three hops. This is the reason why four or more hops seldom occurs. The graph in Figure 5 shows the fraction of nodes at different hop distances averaged over the vignette. In the figure, it is shown that at least a few paths need 4 or more hops. Furthermore, it can be seen that the number of hops for the overall network increases from the NB 25 kHz to the MB 250 kHz to the WB 1.25 MHz waveforms, which is to be expected.

The link dynamics in terms of lost links per node is also analyzed over Vignette 2 for the three different waveforms. As before, the NB waveform uses the 50 MHz frequency band, the MB and the WB waveform use the 300 MHz frequency band. In order to understand the link dynamics, the following simulation is performed; we use OLSR's HELLO protocol to monitor gained and lost links. A HELLO message is transmitted from each node with a given retransmission interval, the default value in OLSR is two seconds.

To decide whether a link exists, we use the basic link quality estimation included in the OLSRv1 RFC [13]. This method estimates the reliability of a link based solely on OLSR HELLO packets. The method assigns weight 1 to all received hello packets and weight 0 to all lost hello packets over a link. To obtain a measure of the link quality, denoted Q , the weight sequence is exponentially filtered according to [13], resulting in values in the range between zero and one. With standard OLSR parameter settings, a link is classified as reliable if Q is larger than an upper threshold set to 0.8. When

a link is classified as reliable, it will remain reliable until Q becomes lower than a lower threshold set to 0.3.

Figure 5: The fraction of nodes at different hop distances averaged over vignette 2.

Figure 6: Average number of lost links per node and per second over the two hour long Vignette 2.

We denote the minimum required number of consecutive correctly received HELLO packets to establish a link by M . For the OLSR standard setting, the algorithm will consider a new link reliable if three consecutive HELLO packets are received, i.e. $M=3$. If two consecutive packets are lost on a reliable link, the link will be considered unreliable. To obtain $M=6$ the upper threshold of Q is adjusted. The M value determines how cautious the routing protocol would be to consider that a link is established. With a large value on M it takes longer to establish a link. Fewer links would exist in the network and the connectivity would be lower than with a small value of M . Furthermore, the number of lost links per node and per second is lower with a larger value than with a smaller value for M .

Figure 6 plots the number of lost links per node per second, averaged for all the 157 nodes. Note that the range is better with the NB waveform than with the WB waveform. As a consequence, on average, the node-degree of the NB network is 120 links, but only about 45 links in the WB network. See [5] for more insight in the correlation between node degree and stable connectivity for different relative mobility. At different

points of time in Vignette 2, the difference between the NB and the WB waveform can be large. When $M=3$, at about 6000 seconds into Vignette 2, only 0.15 links per node are lost for the NB waveform but as much as 0.6 links per node are lost for the WB waveform. It can also be seen that fewer links are lost when $M=6$, than when $M=3$. With 0.6 lost links per second and 157 nodes, about 95 links in the network are lost per second. This is a considerable number of topology changes in the network. Many lost links per second reduce the packet delivery ratio in the network. As the HELLO interval is 2 seconds, it may take up to 4 seconds to detect that a link is lost. Therefore, even if a link in reality is lost during this time period, a node may still try to use it to send a packet. We can conclude that there is a high rate of topology changes in the network and that the dynamics varies depending on the waveform used and the elapsed time of Vignette 2.

A 1000 second long segment in the latter phase of the Vignette 2 has been used to investigate the scalability and performance of OLSR in [3]. Also by using this segment, the effects of small-scale fading on the stability of the links is analyzed in [6].

The objective of the two presented analyses is to provide an understanding of the expected behavior of the network in the Anglova scenario. We hope that others will exploit this scenario for a variety of experiments with routing protocols as well as other tactical networking protocols. Characterizing the behavior of the links at two different frequencies (50 MHz and 300 MHz) and three different bandwidths (NB at 25 kHz, MB at 250 kHz, and WB at 1.25 MHz) provides the necessary background information that can be used by others to evaluate the performance of their protocols/components over the Anglova scenario.

B. Naval Task Group of the Scenario

The Naval Task Group of the Anglova scenario is part of both Vignette 2 (troop deployment) and Vignette 3 (urban operation). One Task group is formed along the coast of Anglova. The Task Group is under the operational control of Fleet Commander / Maritime Interdiction Force (MIF) Commander located at the Coalition Head Quarters. The task group consists of one command ship holding the flag officer and 20 other surface vessels. There is also one multipurpose helicopter, which provides medevac duties within the task group. Each of the task units perform operations within LOS of at least one other ship in order to utilize Line of Sight (LOS) connectivity with the group, so that they can take advantage of the V/UHF frequency band for communications. However, in some situations, HF communications is utilized when LOS is not possible.

Pathloss generation for the naval platforms in Vignette 2 was accomplished by using the open source SPLAT! (Signal Propagation, Loss, And Terrain analysis tool) program [7] using the Longley-Rice model. Topographical information was imported from the 3-arc second SRTM data, which is publicly available. These calculated pathloss values for the ships is also available as part of the Anglova distribution.

C. Vignette 3: Urban Operation

The third vignette takes place within an urban area located at the end of Vignette 2 and covers an urban counter-insurgency operation involving three platoons (72 nodes), 10 unattended ground sensors, one aerial sensor (Aerostat), two UAVs (tactical and data harvest), three satellites, the 21 navy ships that are continuing the maritime mission, and the multi-purpose helicopter tasked for medevac. Vignette 3 is split into three parts. Part one involves vehicular and ground troop movement to a known insurgent location (Figure 7). Part two includes the neutralization of the insurgents and an IED (also Figure 7). Part three involves a medevac of the wounded from the urban environment to a naval ship and finally concludes with the platoons returning to the CHQ (Figure 8).

Figure 7: Vignette 3: Platoon movement to insurgency location.

Figure 8: Vignette 3: Insurgency/IED neutralization and medevac of wounded.

Unlike Vignette 2, the mobility for Vignette 3 was generated by using Google Earth and the US Naval Research Laboratory's (NRL) Network Modeling Framework (NMF) to produce detailed scripted node mobility models. NMF was also

used to generate the EMANE emulation inputs for node connectivity. Pathloss for all of the nodes in Vignette 3 was generated using SPLAT! and the Longley-Rice propagation model.

D. Radio Models

Radio model configurations were generated for both narrowband and wideband radios. The actual EMANE files are available for download as part of the Anglova scenario. This section describes the characteristics of the radio models.

1) 25 kHz narrowband radio model

We have based the radio models on the following assumptions: Vehicle mounted tactical radios typically have an output power of 50 watts (i.e. 47 dBm) with a typical noise figure of 12 dB or better. Antenna gain, cable loss, and connector losses typically sum up to 0 dB. When two radios are on the same vehicle and operating in the same frequency band (i.e., co-site operation), they have a desensitization of 6 dB or better. However, this aspect is not currently included in the emulation models.

Fading has to be taken into account. Fading consists of two components, slow and fast fading. A typical fading margin is 8 dB. As the emulation is based on a field measurement campaign, the precomputed propagation model includes slow fading. Fast fading is not currently modeled and not taken into account in the emulation. The thermal noise figure is -144 dBm/kHz, which is included in EMANE.

To represent a typical narrowband radio, characteristics similar to the N2 mode of the physical layer of the NATO narrowband waveform¹ were chosen as it provides a good compromise between data rate and range. This mode has a bandwidth of 25 KHz. Thus, the thermal noise power is -130 dBm ($10 * \log_{10}(25) = 14$ dB). Adding the noise figure provides the receiver sensitivity, which is -118 dBm. As typical frequencies in the military VHF band (30 to 88 MHz) the frequencies 50 MHz, 51 MHz and 52 MHz were chosen.

Simulation results of an approximated model of the NATO NBWF mode N2 requires the following signal to noise ratio (SNR) threshold values for the given block error rates (BER) (see [8] for more information):

BER	100%	90%	60%	30%	10%	0%
SNR	< 8.7 dB	9.0 dB	9.3 dB	9.7 dB	10.1 dB	≥ 10.9 dB

The average block size was estimated as follows. An average transmission requires two slots. The NATO NBWF PHY for mode N2 has a PHY data rate of 31.5 kbps. The number of user bits per mode scales linearly to the NATO NBWF PHY data rate. For mode N2, the factor is 31.5 kbps / 20.0 kbps. The average block size thus is 564 bits = (284 bits + 432 bits) / 2 * (31.5 kbps / 20.0 kbps).

Mode N2 has a MAC data rate of 17.5 kbps and uses a dynamic MAC, which was emulated within EMANE using the

¹ STANAG 5630 Edition 1 "Narrowband Waveform for VHF/UHF Radios", Ratification Draft, NATO
UNCLASSIFIED

TDMA MAC model, a generic TDMA scheme that supports schedule distribution and updates in real-time using events.

2) 250 KHz Medium Band Radio Model

The medium band radio model is deduced from the narrowband radio model by increasing the bandwidth and the data rate by a factor of 10 to 250 kHz and 175 kbps. Thus, the thermal noise power is -120 dBm ($10 * \log_{10}(250) = 24$ dB). Adding the noise figure provides the receiver sensitivity, which is -108 dBm. As typical frequencies in the military UHF band (225 to 400 MHz), the frequencies 300 MHz, 301 MHz, 302 MHz and 303 MHz were chosen for the different medium band networks within the scenario.

3) 1.25 MHz Wideband Radio Model

As an alternative to the 250 KHz medium band radio model, a 1.25 MHz wideband radio model was also developed by increasing the bandwidth and the data rate by a factor of 50 to 1.25 MHz and 875 kbps. Thus, the thermal noise power is -113 dBm ($10 * \log_{10}(1250) = 31$ dB). Adding the noise figure provides the receiver sensitivity, which is -101 dBm. The frequencies 300 MHz, 302 MHz, 304 MHz and 306 MHz were chosen for the different wideband networks within the scenario.

All of the EMANE configuration files for these radios are available as part of the overall Anglova scenario distribution.

IV. EXPERIMENTATION ENVIRONMENT

This sections provides some detail on EMANE and the deployment for the experimentation environment.

A. EMANE Overview

The Extendable Mobile Ad-hoc Networking Emulator (EMANE) [9] was developed by CenGen (now AdjacentLink) under sponsorship by the US Naval Research Laboratory (NRL), the Office of the Secretary of Defense (OSD), and the US Army Research Laboratory (ARL). It is freely available on GitHub and provides a flexible framework for emulating multiple types of radios. EMANE primarily addresses the physical (PHY) and media access (MAC) layers of the network stack and includes three different radio models – a generic RF Pipe model, an IEEE 802.11abg model, and a TDMA model. Each of these models offers various customizations of the radio parameters, antennas, and error models. While models of some specific tactical radios exist, they are not available in the public domain and hence will not be discussed in this paper.

EMANE supports a flexible deployment model – with centralized control, fully distributed control, or a hybrid of the two. Each node in the emulated network is represented within EMANE by an instance of a Network Emulation Module (NEM). The NEMs instantiate the PHY and MAC layers of the radio being emulated. One NEM represents one network type. For example, if a router has several radio interfaces, each interface has its own NEM. Each EMANE instance can contain one or more NEMs. In the centralized model, all the NEMs are instantiated within a single instance of EMANE. In the fully distributed model, there is one EMANE instance for each NEM. In the hybrid model, there are multiple instances of EMANE, each with a subset of the overall NEMs. In the

hybrid and distributed deployment models, the NEMs communicate with each other using an Over The Air (OTA) channel, which uses UDP multicast over a control network (which is independent of the emulated network interfaces). The NEMs also react to control events that are sent to them over the EMANE Event channel, typically also via UDP multicast. For example, when running the Anglova scenario, the position and pathloss updates are sent to EMANE via this control channel.

The following two subsections describe the two different deployments that have been setup for experimentation in IST-124. The first one describes a fully distributed deployment that utilizes an overarching experimentation framework called DAVC. The second deployment is a hybrid deployment based on the VMware ESXi virtualization platform. Finally, the last section describes open issues that were encountered with EMANE during the course of experimentation.

B. Dynamically Allocated Virtual Clustering (DAVC)

The NATO IST-124 uses the US Army Research Laboratory's Dynamically Allocated Virtual Clustering Management System (DAVC) to deploy the Anglova scenario using the distributed EMANE emulation model. In this emulation model, the EMANE software is installed within VMs that execute the applications that are the subject of the experimentation and whose performance is being evaluated.

DAVC is a web based virtualization service and cloud-operating environment that creates complex virtual experimentation clusters that can be used for simulation-based, emulation-based, and hybrid field/emulation experimentation [10]. DAVC deploys networked clusters composed of VMs tailored to user specifications. The DAVC management system abstracts away test-bed infrastructure configuration through automated provisioning processes that configure the virtual networking for each VM [10]. Clusters created by DAVC are heterogeneous, so each VM can have different OSs, application sets, and hardware attributes such as RAM, CPU cores, hard disk, and network interfaces. DAVC users can register custom VMs as templates that can be used within their experimentation clusters.

IST-124 created a custom Ubuntu 16.04 VM to represent nodes within the Anglova scenario. This template VM includes the applications necessary for running the Anglova scenario including EMANE, the Multi Generator (MGEN) [11], and the OLSRv1 and OLSRv2 routing protocols [12]. The VM also includes the various EMANE radio models, mobility, and path-loss configuration files specific to the Anglova scenario. Custom scripting to bootstrap the Anglova scenario and emulation environment is also available in the VM. Once registered with DAVC, the VM is then used as a template within a DAVC experimentation cluster to run the Anglova scenario.

Using DAVC, the Anglova scenario is distributed with a 1-to-1 mapping with each Anglova node running within a single DAVC virtual cluster node. The entire 283-node scenario can run within a 284-node DAVC cluster as shown in Figure 9. When deployed in this manner, node 284 acts as the experimentation orchestration node and is responsible for executing the bootstrap scripting that launches the various

applications and EMANE on the remaining 283 nodes. Also shown in Figure 9 are two networks DAVC auto-configures for the 284-node cluster. The first network is an out-of-band control/debug network that allows node 284 to communicate with and execute experimentation instructions on the other 283 nodes. The second network is the experimentation network that EMANE uses to overlay the emulated Anglova radio network channels upon.

Figure 9: The 283 node DAVC cluster and networks. The emulated EMANE network is overlaid on the DAVC exp network while experimentation commands flow across the DAVC control network.

The emulation environment provides the flexibility to deploy subsections of the entire 283 node scenario within a DAVC cluster. If a researcher is only interested in experimenting with the 159 mobile nodes contained in Vignette 2 discussed in Section III.A, the bootstrap scripting can map only those nodes contained within that portion of the Anglova scenario to 159 DAVC cluster nodes where the corresponding EMANE configuration files will be executed. This type of deployment model that uses only a subsection of the Anglova scenario was used in the experiment performed in [3]. The authors were interested in evaluating the scalability of the OLSRv2 routing protocol on company and multi-company sized topologies ranging from 24 nodes to 96 nodes, therefore they deployed several 96 node DAVC clusters and configured the experimentation scripting to only run a portion of the scenario that contained four 24 node companies.

The DAVC model also supports experimentation concurrency, which allows multiple experiments to be conducted repeatedly in parallel. The experiments performed in [3] are an example of this model. Multiple DAVC clusters were deployed where each hosted a different experimental test case. Some experiments were run in parallel where each experiment executed a different version of the OLSR routing protocol and different parameters were provided to the experimentation scripting that varied the MGEN background traffic generation model, network traffic shaping, and the number of Anglova nodes involved in the experiment. This

feature of the DAVC model has the effect of reducing overall experimentation runtime.

C. VMware ESXi

In addition to the DAVC deployment model, the Anglova scenario was also deployed using a second, alternate configuration using the VMware ESXi virtualization platform. This deployment uses the hybrid model for EMANE configuration, where N virtual machines have one EMANE server VM that runs all of the emulation components for those VMs. An example configuration is shown in Figure 10. Note that the test VMs are designated test nodes (TN- n) and there could be multiple per physical server, with a recommendation of one VM per CPU core. Each test VM is also configured with at least two network interfaces, one for the application data and one for control traffic. Each physical server also contains an EMANE Server VM, which runs all of the EMANE components. One advantage of this deployment model is that the test VMs do not have to run any of the EMANE components.

Figure 10: Anglova Deployment over ESXi

The network shown in green is configured as a virtual switch within ESXi, with each test VM connecting one or more interfaces to this virtual switch. Each connection is assigned a different VLAN id. The EMANE Server VM connects to this switch using a VLAN trunk, which uses the 802.1Q protocol. Within the EMANE Server VM, multiple virtual interfaces are created, one per VLAN, which are then managed by EMANE. For this configuration, any traffic generated in the data interfaces of the TNs will be transferred to the corresponding virtual Ethernet interfaces inside the EMANE Server, where each interface is mapped to one NEM. At runtime, EMANE reads from each interface, applies the necessary communications effects, and writes the data out to

the interfaces that should receive the data. The network shown in blue is for EMANE data (OTA and events) exchanged between the different physical servers. Finally, the network shown in red is the control network used by a user to log into the VMs, start/stop the experiments, and collect data.

D. Open Issues with EMANE

During the course of experimentation, four open issues were encountered that still need to be address by the emulation framework. Note that these issues are not with the Anglova scenario, but with the underlying EMANE framework. The first issue was with the RFPipe implementation of EMANE. In particular, the radio models for Anglova were first developed to use RFPipe. However, the RFPipe does not implement any MAC algorithm and does not realize any interference effects. While the radio model allows the definition of a capacity limit as well as realize latency and reliability effects, multiple nodes are allowed to transmit at their individual rate limits and the receivers could receive, in aggregate, data rates higher than the capacity of the radio. Furthermore, RFPipe does not realize typical wireless communications effects such as hidden neighbors. RFPipe is good for point-to-point links, but cannot be used as-is for a radio network sharing a common channel.

Another issue with the radio models was discovered when using the TDMA model to emulate the narrowband links within the scenario. Unlike the RFPipe, the TDMA model includes a MAC and limits the data rates correctly. However, the TDMA implementation requires that all the NEMs using the TDMA model have synchronized clocks with an accuracy that is a function of the TDMA time slice. For the Anglova scenario, the accuracy of NTP was insufficient for the TDMA model to work in a distributed deployment. The hybrid deployment works if the NEMs for all of the nodes that are part of a TDMA network are deployed within the same EMANE instance.

We observed the following issue with the 802.11abg model within EMANE: When multiple senders are transmitting at the maximum data rate to a single receiver, the incoming data rate at the receiver exceeds the maximum channel rate by a small fraction (as a function of the number of transmitters). This implies that EMANE's 802.11 model does not correctly implement the 802.11abg standard.

Finally, the last issue was with the EMANE event generation model, which allows a control application to generate events that control the emulation environment. These events are sent to the NEMs via UDP multicast. As mentioned earlier, the Anglova scenario has a playback component that generates the position and pathloss events to send to EMANE. Given the size of the Anglova scenario and the update rate of 1 Hz, the EMANE receiver seemed to be too slow to receive and process the events without losing some of the packets due to the UDP receive buffer being too small. This problem was solved by splitting the scenario into smaller subsets that were played back independently but still synchronized so that EMANE could support all the nodes simultaneously.

V. CONCLUSIONS

The Anglova scenario and all of the related files for setting up and running experiments using EMANE are available for download at <http://www.ihmc.us/nomads/scenarios> and also at <http://www.arl.army.mil/nsrl>. For IST-124, having this common scenario has worked very well for efficient collaboration across multiple countries and organizations, especially when hosted in a cloud environment with tools such as DAVC to setup experiments and provide access. We hope that the details provided in this paper are useful to other researchers that are interested in using the Anglova scenario for their own experimentation.

ACKNOWLEDGMENT

Other IST-124 group members who have contributed ideas, comments, and suggestions to the Anglova scenario include Piotr Lubkowski (POL), Anne Marie Hegland (NOR), and Taner Kurtulus (TUR). We also thank King Lee (USA) who was a significant enabler and contributed much to the scenario.

REFERENCES

- [1] N. Suri et al., A Realistic military scenario and emulation environment for experimenting with tactical communications and heterogeneous networks, in Proceedings of the 2016 International Conference on Military Communications and Information Systems (ICMCIS), Brussels, 2016, doi: 10.1109/ICMCIS.2016.7496568.
- [2] S. Kurkowski, T. Camp, and M. Colagrosso. MANET Simulation Studies: The Incredibles, in ACM SIGMOBILE Mobile Computing and Communications Review, vol. 9, no. 4, pp 50-61, 2005.
- [3] K. Marcus et al., Evaluation of the Scalability of OLSRv2 in an Emulated Realistic Military Scenario, in Proceedings of the 2017 International Conference on Military Communications and Information Systems (ICMCIS), Oulu, Finland, 2017.
- [4] P. Holm, UTD-Diffraction Coefficients for Higher Order Wedge Diffracted Fields. IEEE Trans. Antennas Propagat., vol. AP-44, pp. 879-888, June 1996.
- [5] E. M. Royer, P. M. Melliar-Smith, and L. E. Moser, "An analysis of the optimum node density for ad hoc mobile networks," in proceedings IEEE ICC, Helsinki, Finland, pp. 857-861 vol.3, 2001.
- [6] S. Sterner, U. Uppman., On the Robustness of OLSR in a Mobile Tactical Scenario in Rural terrain, in Proceedings of the 2017 International Conference on Military Communications and Information Systems (ICMCIS), Oulu, Finland, 2017.
- [7] J. Magliacane. SPLAT!. Internet: <http://www.qsl.net/kd2bd/splat.html>, [April 20, 2018].
- [8] B. Libæk and B. Solberg, "A simulator model of the NATO Narrowband Waveform physical layer," FFI Notat 2011/00533 , 2011.
- [9] AdjacentLink, LLC. Extendable Mobile Ad-hoc Network Emulator (EMANE). Internet: <https://github.com/adjacentlink/emane/wiki>, [April 20, 2018].
- [10] K. Marcus and J. Cannata. Dynamically Allocated Virtual Clustering Management System, in Proceedings of SPIE 8742, Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR IV, May, 2013.
- [11] US Naval Research Laboratory. Multi-Generator (MGEN). Internet: <http://www.nrl.navy.mil/itd/ncs/products/mgen>, [April 20, 2018].
- [12] Optimized Link State Routing (OLSR). Internet: <http://www.olsr.org>, [April 20, 2018].
- [13] T. Clausen and P. Jacquet, Optimized link state routing protocol (OLSR), IETF Network Working Group, RFC 3626, October 2003.