

Using Modeling and Simulation to Examine the Benefits of a Network Tasking Order

Capt Matthew D. Compton, MA

Kenneth M. Hopkinson, PhD

Gilbert L. Peterson, PhD

James T. Moore, PhD

Air Force Institute of Technology

2950 Hobson Way

Wright-Patterson AFB, OH 45433

matthew.compton@afit.edu

kenneth.hopkinson@afit.edu

gilbert.peterson@afit.edu

james.moore@afit.edu

The Global Information Grid (GIG) is the military's computer and communications network which supports the myriad of military missions. Since military missions are highly planned, it is possible to collect knowledge of future conditions that could be of use to network designers whose goals include optimizing and protecting the GIG. This advanced knowledge includes which assets will be involved, what their capabilities are, where they will be, when they will be there, and particulars on the required data flows. A Network Tasking Order process is proposed as a means of collecting this information, analyzing the information to generate network taskings, and disseminating those taskings. Tactical integration of assets in mobile networks is introduced as another planning variable in the battlefield; not unlike logistical considerations such as fuel, ammunition, water, and so on used currently in operation planning. Modeling and simulation is used to support the proposed benefits.

Keywords: Global Information Grid, Network Tasking Order, Quality of Service, Command and Control

Disclaimer: The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

1. INTRODUCTION

The Global Information Grid (GIG) is defined in Department of Defense (DoD) Directive 8100.1 [1] as the “globally interconnected, end-to-end set of information capabilities, associated processes, and personnel for collecting, processing, storing, disseminating and managing information on demand to warfighters, policy makers, and support personnel.” In particular, the GIG includes all equipment involved in the transfer of information, that is, the military

communications network. Many subnetworks of the GIG, especially those in a wartime environment, are formed by mobile nodes.

Mobile nodes are by necessity wireless, and oftentimes have little infrastructure with which to connect in battlefield milieus. Therefore, nodes must have the capability to connect directly with each other and form what is called a Mobile Ad-hoc Network (MANET). Some of the advantages of MANETs are that they are rapidly deployable and can be self-organizing. Without infrastructure, each node must act as a router to ensure information can travel between nodes that are not directly linked. Along with the flexibility that MANETs provide, there are also some disadvantages.

Networks involving mobile nodes can have highly dynamic topologies. Without careful planning, the topologies that form can suffer from poor Quality of Service (QoS). The networks could have bottlenecks, or worse, be disconnected. Certain topologies may require data packets to make many hops to travel from a source node to a destination node, which leads to excessive delays. The loss or delay of information can have dire consequences. At the other extreme, a topology could be highly redundant and underutilized. When costs such as battery life, maintenance, or frequency allocation are considered, this is decidedly wasteful. It is in the DoD's best interest to mitigate these disadvantages as much as possible.

In a MANET, individual nodes often make their own decisions on how to connect to the network. Unfortunately, what may appear to be the best decision for a singular node may not be in the best interest of the network as a whole. For instance, one particular node may choose to send information through a path that has the fewest hops or the strongest signal, but as a result increases congestion at another node. This could result in delays or dropped packets for a higher priority data stream trying to route through the same congested node.

The GIG is composed of many parts. Individuals and mechanisms charged with network planning need to be aware of, to the greatest extent possible, what composes the GIG, where those pieces are, what their capabilities are, and how they are connected or communicating, both currently and into the future. Possession of this "GIG-awareness" would allow for tactical integration of assets as another planning variable in the battlefield; not unlike logistical considerations such as fuel, ammunition, water, and so on used currently in operation planning. There are a variety of ways in which such awareness can be leveraged to help glean improved performance out of the equipment being used and to aid in protecting that equipment from cyber attack.

The development and application of a Network Tasking Order (NTO) process is proposed as a means of enhancing "GIG-awareness" by taking advantage of the highly planned nature of military operations. Unlike most MANET research which often relies on random mobility models [2,3], in military scenarios the GIG could profit from foreknowledge of the general locations of assets, when they will be there, and the type of traffic they will generate. The NTO would be an analogue to and offspring of the Air Tasking Order (ATO), the daily tasking of air missions. The NTO could be used in the planning stage (as a pre-NTO) to help expose shortcomings or redundancies in the network. Also, the NTO could be useful in the execution stage to help quickly recover from unexpected events such as finding new routes for traffic when a node fails.

2. BACKGROUND

The idea of an NTO as put forth in this paper is not entirely new. Ranne and McKee [4] advocate that United States Strategic Command's (USSTRATCOM) Joint Task Force – Global Network Operations (JTF-GNO) and/or Air Force Network Operations (AFNetOps) should conduct

concept and prototype development with GIG NetOps Tasking Orders (GNTOs) as a means for command and control of the GIG. These GNTOs could be used “to communicate not only what to do and who does it with what assets; but also what to monitor and assess.” The authors envision three categories of GNTOs:

1. Standing Orders: for routine day-to-day operations,
2. Cyclical Orders: to communicate planning and resource allocation for specific periods of time, similar to an ATO, and
3. Dynamic Orders: to communicate near real-time direction for security and allocation issues.

In a recent email [5], McKee indicated that USSTRATCOM leadership liked the idea and are very interested in command and control of the cyberspace domain. JTF-GNO has implemented a version of the GNTO that is network defense focused, but not really integrated with anything. It is focused on computer/land networks with no real thoughts on air networks or a larger “cyber” perspective.

The concept of an NTO has also appeared in Stookey [6], which provides background and data to build a notional battlespace for testing and simulating the use of dynamic networks within the US military. Stookey elaborates on the necessity of developing an NTO to provide dynamic network routers a basis for making predictive decisions about where given nodes are spatially in a battlespace, what data links might be available, the bandwidth or throughput of such links, the bandwidth requirements of various data flows, and the priority of the data that might be destined to or coming from various nodes. Pecarina [7] envisions an NTO in which a Joint Forces Cyber Component Commander could assign weights of effort to different mission goals in cyberspace. In addition, he sees the NTO as a means of addressing a QoS issue. That is, the NTO will help to move to a point where information that is not needed or wastes time, bandwidth, and energy is blocked to allow the critical data to get through.

There are already military documents called NTOs. The 624th Operations Center (624 OC) was activated in August 2009 along with a new numbered air force, 24 AF. The 624 OC is 24 AF’s command and control operations center, responsible for the Air Force provisioned portion of the GIG for the purpose of Network Operations (NetOps) and Network Defense [8]. In the execution of that role, the 624 OC commander issues a variety of order types. Currently, these order types include Maintenance Tasking Orders (MTOs), NetOps Tasking Orders (NTOs), and Time Compliance Network Orders (TCNOs). Briefly, MTOs are for generalized cyber maintenance actions, NTOs are for urgent defensive actions, and TCNOs are for routine system vulnerability patching [8]. These various documents are more directed towards actions like moving to standard desktop configurations, disallowing thumb drive use, or blocking various file extensions at mail relays.

Finally, the 50th Network Operations Group at Shriever Air Force Base publishes a daily NTO to assist in allowing command and control of the Air Force Satellite Control Network which includes such as the Defense Support Program, the Navstar Global Positioning System, the Defense Satellite Communications System, NATO III, and Milstar. This schedule makes sure the satellite fliers have the needed ground antenna resources to perform routine tasks and to perform telemetry and other data transfers [9]. This version of the NTO appears most similar to the product of the NTO process proposed herein, but focuses on a much smaller portion of the GIG.

3. The NTO Process

The NTO is envisioned as an analogue to and offspring of the ATO, the daily tasking of air missions. In the future, the NTO may be more of a sibling to the ATO, providing a daily tasking of network missions to support the air missions. In the (maybe not so) distant future, the roles may reverse with the ATO created as a support for the NTO. In the meantime, the NTO creation process should follow a cycle similar to the ATO cycle.

ATOs are created by the Combat Plans Division of the Joint Air and Space Operations Center (JAOC). ATOs follow a three-day life cycle [10]. On any given day, the current day's ATO is being executed, the next day's ATO is being produced, and the subsequent day's ATO is being planned. The NTO process should closely follow, or be done in unison with the ATO process because the network is there to support the missions. The missions need to be known in order to plan the network. However, the process of creating the NTO can add insight that might result in changes to the ATO.

A large portion of the data needed for the NTO will come from the ATO itself. There are other planning documents such as the Space Tasking Order (STO), Operational Tasking Data Links (OPTASK LINK), and the Joint Communications Electronics Operating Instruction (JCEOI) that contain useful information as well. The STO is a document that is developed in parallel with the ATO. Its primary purpose is for tasking space assets with specific missions. Since satellites are used for communication, this information is of interest to network designers. The OPTASK LINK specifies data link procedures within a battle group, defines theater datalink architecture, and serves as a list of who can/may talk to whom. The OPTASK LINK provides unit locations, frequencies, duties, filter plans, and other link specific information [11]. Again, this is data that is of interest to network designers. The JCEOI is used for frequency allocation and deconfliction.

The content and structure of the ATO as well as the STO and OPTASK LINK are given in great detail in [12]. The ATO is a human/machine readable document written in United States Message Text Format (USMTF). This same format should be used for the NTO. Over 370 different types of messages use USMTF [12] which is familiar to most service members, and existing message browsers could easily be adapted to accommodate NTOs.

All of the missions in the ATO are grouped first by tasked country, then by tasked service, and after that by individual tasked units. It makes sense to keep this structure in an NTO so that the units that own each network component can easily find the pieces they are responsible for and configure them for the planned day. In addition, it is possible that communications will exist between the asset and its home unit that need to be planned for.

Once the ATO gets down to the individual unit level, all missions that a particular unit are responsible for appear sequentially. Within each mission, the number and type of aircraft along with call sign and primary configuration are given. The following information may also be listed: secondary configuration codes, Link 16 abbreviated call sign, Tactical Air Navigation system (TACAN) channel, primary Joint Tactical Information Distribution System Unit address, and identification friend or foe/selective identification feature (IFF/SIF) mode and code.

Along with knowing which assets will be involved in the network, it is important that network planners know what capabilities those assets have. There may be differences in networking capabilities between two assets of the same type. For example, the capabilities of an A-10 from one unit could be radically different from those of an A-10 from another unit. However, within a single unit the differences should be minimal. It would be beneficial to have

available a baseline capability for each asset type by unit. Some characteristics that would be valuable include:

- Types of interfaces (free space optical, radio frequency, etc.),
- Number of interfaces,
- Functional areas implemented for each data link,
- Frequencies or channels available,
- Transmission speed/range,
- Data forwarding capabilities,
- Encryption capability,
- Accepted protocols,
- Set-up time, and
- Queue characteristics.

This information should be placed into a centralized database at the JAOC so that when an asset is tasked, the networking capabilities of that asset can be automatically available to the network planners. To the best of our knowledge, no all-inclusive database containing this information currently exists. The Joint Spectrum Center (JSC) Equipment, Tactical, and Space (JETS) database comes very close. The JETS database contains detailed technical information about communications, radar, and electronic warfare equipment as well as operational parameters for each subsystem and component [13]. Though it would likely be classified, it ought to be fairly straightforward to compile such a database. Once constructed, the database could be kept current through updates from depots, program offices, or individual units.

The assets tasked in the planning documents can be cross-referenced with the capabilities database. The combined information can then be collated into a daily schedule for the network. This is a pre-NTO. No network specific taskings have been made yet. Analysis can now be performed upon this pre-NTO. By having a plan in place, it becomes more apparent where there could be single points of failure, gaps in connectivity, or bottlenecks. In such cases, it might be possible to change an orbit or add an extra asset. For example, one waypoint of an E-3's orbit could be adjusted slightly to allow periodic high bandwidth line-of-sight (LOS) communications to a ground unit that would otherwise be plagued with constant low bandwidth connections. Another example could be the addition of a communications relay mission for an unmanned aerial system (UAS) to linger over a certain location to act as a wireless router in support of a high priority mission. Feedback from the analysis can thus be used to make necessary changes to the ATO and other planning documents before they are published. Once sufficient analysis and feedback cycles have been performed, appropriate networking directives can be formed and published in a finished NTO. This NTO process is illustrated in Figure 1.

Missions in the ATO usually include a route with altitudes and speeds. Routes can either be a round trip to a target location with departure/return locations and times, one-way travel with departure/arrival locations and times, or orbit information with departure/return locations and times. In any case, given this information, there is some general idea of where an aircraft is going to be and when it will be there. When satellites fly over for limited but predictable time spans, the utilization of these resources can be planned for ahead of time. For example, a directional antenna can be prepositioned to the expected pointing angle so that it is ready for service when needed. Thus some simplified version of this information should be included in the pre-NTO.




Figure 1: NTO Data Flow

Each mission has a preferred mission type or designation. Mission type does not necessarily need to go into the pre-NTO, but it may give a clue as to the type of traffic. For example, a combat search and rescue (CSAR) mission will have different traffic characteristics than air reconnaissance or aerial refueling. The expected quantity and burstiness of traffic flows are important measures to include in the pre-NTO. These characteristics can be known through historical precedence.

By including the priority of certain traffic flows, routing agents can use this information to make decisions in situations of congestion. The agents can decide to allow the high priority information to pass through while dropping or delaying the lower priority information. Other alternatives are directing information over different routes or storing information to send at times of lower activity. Missions in the ATO are given priorities, and one can generally assume the transmissions of a mission would have corresponding priority. Thus, it will be imperative to carry these priorities into the pre-NTO to allow for ranking of traffic flows.

4. Example NTO Process

It may be useful at this point to illustrate what a translation of a single mission from an ATO into an NTO might look like. The example provided is academic and does not represent any real mission. In the ATO, each line of data, or set, is terminated by '//'; however, due to length it may wrap to fill multiple lines of text. Within a set, fields are separated by ','. Fields containing '-' are optional and no data has been entered. Figure 2 shows a few sets from an example ATO that pertain to a single mission.

The first three sets detail whom is being tasked in increasing specificity. The first set indicates that the tasked country (TSKCNTY) is the US. The second set designates the Air Force (F) as the service being tasked (SVCTASK). The third set specifies the unit being tasked (TASKUNIT) and its location. Here, it refers to the 23rd Fighter Squadron (23FS) at Spangdahlem Air Base in Germany. The location in this instance is given by the International Civil Aviation

Organization (ICAO) four character identifier ETAD. Location can also be specified by place name or by latitude/longitude.

```
TSKCNTRY/US//  
SVCTASK/F//  
TASKUNIT/23FS/ICAO:ETAD//  
AMSNDAT/N/D123HB/-/-/SEAD/-/-/DEPLOC:KGZ6/241200ZAPR/ARRLOC:  
    KDZ7/241300ZAPR//  
MSNACFT/1/ACTYP:F16CJ/SUPP01/2HARM/-/20001/30111//  
AMSNLOC/-/-/210/1//
```

Figure 2: Sample ATO mission

The fourth set contains aircraft mission data (AMSNDAT) with twelve fields of information. The first field is the residual mission indicator. An 'N' means that the mission is non-residual; the mission falls entirely within the ATO period. The next field is for the mission number identification, here it is D123HB. The third (Air Mobility Command mission number or event number), fourth (package identification), and fifth (mission commander) fields have been left blank. Field six holds the preferred type or designation for the mission. In this case, it is a Suppression Enemy Air Defense (SEAD) mission. Fields seven (secondary mission type) and eight (alert status) have been left blank. Field nine is used to specify the departure location of the mission if it differs from the location specified in the TASKUNIT set. An ICAO identifier, KGZ6, is given for the location. Field ten gives the day-time and month of departure, April 24 at 1200 Zulu. Field eleven provides the recovery location of the mission if other than the location specified in the TASKUNIT set. Again, an ICAO identifier, KDZ7, is given for the location. Finally, field twelve gives the day-time and month of recovery, April 24 at 1300 Zulu.

The fifth set holds individual aircraft mission data (MSNACFT) with seven fields of information. The first field gives the number of aircraft as 1. The second field provides the type and model of aircraft as an F-16CJ Fighting Falcon. The aircraft call sign, SUPP01, is placed in field three. In the primary configuration code field, 2HARM indicates the aircraft should be equipped with two AGM-88 high-speed anti-radiation missiles. Field five (secondary configuration code) has been left blank. Fields six and seven are both for IFF/SIF codes. In field six, 20001 indicates a mode 2 code (personal unit identity) with octal value 0001. In field seven, 30111 indicates a mode 3 code (normal air traffic control identity) with octal value 0111.

Finally, in the sixth set, additional mission location (AMSNLOC) information is given with five fields. This set provides mission location information for missions which have no specific target location, for example orbits or alerts. Fields one (day-time and month of start), two (day-time and month of stop), and three (mission location name) are left blank. Field four provides the vertical distance in hundreds of feet above mean sea level. A value of 210 indicates the mission should fly at 21,000 feet. The last field is the code for the priority assigned to a mission, which in this example is 1.

Notice that in a few short lines, there is some pretty detailed information about where this particular F-16CJ will be and when it will be there. Given the departure and recovery locations of the mission, intermediate locations can be interpolated. Similar information will be given for all other aircraft flying during the same time period. Presumptions can also be made regarding with whom this aircraft will be communicating.

Suppose that a capabilities database contains the following (hypothetical) information regarding F-16CJs from the 23rd Fighter Squadron:

- Equipped with one Improved Data Modem (IDM-302) capable of 16 Kbps digital communications over four independent channels accepting AFAPD, TACFIRE, IDL, and MTS protocols [14]
- The IDM-302 is interfaced with an ARC-164 Ultra High Frequency (UHF) Airborne Radio which features 10-85 watts amplitude modulation, 25 kHz channel spacing over a frequency range of 225.000 to 399.975 MHz [15]
- The aircraft includes an omni-directional UH-408 UHF Blade Antenna with -1 dB gain and 50 Ohm impedance [16].

Finally, based on historical records of SEAD missions, an average or expected data rate along with a peak data rate can be established. For the sake of the example, suppose the average traffic rate is 5 Kbps with bursts up to 15 Kbps. It is also possible to predict with whom the F-16CJ will be communicating. There will likely be communications between various aircraft in the same strike package. These other aircraft are identified in the ATO. Also, there should be communications with an E-3 Sentry Airborne Warning and Control System (AWACS) providing situational awareness of the battlefield and battle management. The AWACS closest to the F-16CJ during this mission can be identified from the ATO. There may also be communications between the F-16CJ and its home base and with the Joint Air Operations Center (JAOC).

All of this information is of interest to analysts who are trying to optimize the network and eliminate potential gaps in connectivity or bottlenecks. This information gets collated into a pre-NTO. The data in the pre-NTO can then be fed into tools such as the topology control algorithms or other programs that may exist or have yet to be created. The pre-NTO may have the appearance of Figure 3.

```

TSKCNTRY/US//
SVCTASK/F//
TASKUNIT/23FS/ICAO:ETAD//
AMSN DAT/N/D123HB/-/-/SEAD/-/-/DEPLOC:KGZ6/241200ZAPR/ARRLOC:
KDZ7/241300ZAPR//
MSNACFT/1/ACTYP:F16CJ/SUPP01/2HARM/-/20001/30111//
AMSNLOC/-/-/210/1//
MODEMDAT/1/TYPE:IDM302//
RADIODAT/1/TYPE:ARC164//
IFACEDAT/1/TYPE:UH408//
EXPCOMM/AVE:5/BURST:15/1//
COMMLINK/ICAO:ETAD/ICAO:JAOC/CALL:SKYWATCH43//

```

Figure 3: Sample pre-NTO mission

Notice that the first six sets in Figure 3 are repeated from the ATO. This is important information, most of which could be of use to analysts. Five new sets have been added using the same USMTF format. The first added set, MODEMDAT, provides information on the modem(s) carried by the asset. The first field in this set gives the number and the second field gives the type. Here it lists 1 IDM-302 modem. These fields can be repeated if there are multiple types of modems. The second and third added sets have a very similar format. The RADIODAT set gives the number and type of radio(s) on the asset. The IFACEDAT set lists the number and type of interfaces (or antennas) being used. It may be necessary to make these sets hierarchical if, for instance, different modems are paired with different radios and antennas. These three sets are derived from the information in the capabilities database.

The last two added sets are derived from historical precedence. The EXPCOMM set lists information on the expected communications traffic from this asset on this particular mission. The first field gives the anticipated average traffic rate in Kbps. The second field gives the anticipated maximum traffic rate in Kbps. The final field indicates the priority of the communications. In this case, the priority is identical to the mission priority listed in the AMSNLOC set. This set would also be a good place to put fields such as encryption requirements, type of traffic (voice, video, telemetry, etc.) being sent, protocols being used, and so on. The COMMLINK set lists some of the main communications partners for the asset. The field in this set is repeated for each expected traffic recipient. The labels in the field, ICAO and CALL, indicate how the recipient is identified. The first two fields in this example give the ICAO identifiers for the F-16CJ's home station and for the JAOC. The last field gives the aircraft call sign of an AWACS that will be coordinating the mission.

Finally, an example can now be shown of a tasking that could be generated and given to this particular mission. Suppose that it has been determined that during this mission, the F-16CJ should route all traffic destined for the JAOC through a KC-135 tanker with call sign FUEL03 in an orbit named BLUE 23 in the vicinity of the mission. This tasking may have the following appearance:

```

TSKCNTRY/US//
SVCTASK/F//
TASKUNIT/23FS/ICAO:ETAD//
TASKNODE/ACTYP:F16CJ/SUPP01//
1RTEDAT
/DEST           /START           /STOP           /NHOP           /LOC
/ICAO:JAOC      /241200ZAPR     /241300ZAPR    /CALL:FUEL03   /BLUE 23//

```

Figure 4: Sample NTO Tasking

Notice that the first three sets in Figure 4 are repeated from the ATO and the pre-NTO. These sets narrow down what is being given the tasking. The fourth set, TASKNODE, specifies which node in the network is being tasked. Here it is the F-16CJ with call sign SUPP01. The last three lines form what is known as a columnar set. Columnar sets are arranged in vertical columns under an appropriate column heading. The first line has the set name, which for columnar sets must begin with a number. Here the set is named 1RTEDAT for route data. The second line has the column headers which designate the type of information located in each column. The third line contains the information for this set. The information in this line is entered so that it falls under the proper column headers. The first column, DEST, indicates the final destination for the traffic that is being routed. The final destination here is referenced by the ICAO code for the JAOC. The next two columns, START and STOP, specify the day-time and month span over which this route is to be used. Here, the time span corresponds to the full mission duration found in the ATO, from 1200 to 1300 Zulu on 24 April. The fourth column, NHOP, indicates the next hop for traffic. The KC-135 is referenced using its call sign. The final column, LOC, indicates the location of the next hop. Here, BLUE 23 refers to the KC-135's orbit. Additional columns could be added as needed. For example, channel, encryption type, interface, etc. could be included in this set. As many information lines as needed can be used with columnar sets. This set resembles a routing table.

5. SIMULATION

A simple scenario is now provided to give an example of how *a priori* knowledge of network conditions identified during the planning stage of the NTO process could translate into increased QoS during the execution stage. The same scenario is examined from two points of view (POVs). First, the scenario is considered with a localized POV, where no NTO process is available to guide decisions made at the various nodes. Second, a more global POV is considered, where an NTO process has identified a problem and directives are issued to address it.

Scenario

The scenario consists of two individual sources generating information that needs to be sent to a common headquarters (HQ). HQ is far enough away that direct communication from the two sources is not available, but there is an intermediate node that can act as a router. Both sources have a 36 kbps connection to the router, and the router has a 36 kbps connection to HQ. Source 1 (S1), an MQ-1 Predator, will produce a high-priority video feed at a rate of 30.6 kbps over a thirty-minute time span that falls between the hours of 1400 and 1500. Source 2 (S2), a sensor net, produces a continuous stream of data at a rate of 7.2 kbps. S2's data is of lower priority, as it is collated at HQ and reviewed once daily. In addition to the router, there will also be an E-3 Sentry (AWACS) aircraft flying a 30-minute orbit during the hours from 1300 to 1800 in the airspace between S2 and HQ. This particular AWACS is known to be carrying equipment that allows wireless line-of-sight (LOS) networking at a capacity of 28.8 kbps. The orbit is such that it will not be in LOS of S1. It will be within LOS of both the S2 and HQ for 10 minutes each, with no overlap in these time spans. Thus, the aircraft cannot act as a router. However, the aircraft can be used as a data ferry, storing the information that is uploaded from S2 and downloading it later to HQ when it is in range. See Figure 5 for an overhead view of the scenario.




Figure 5: Overhead view of scenario

This scenario is first considered without an NTO process in place for guidance. S1's operators do not have any options for routing the video feed; they must utilize the router. S2's operators have a little more flexibility. Given the choice between sending its traffic to the router

or to the AWACS, the operators are likely to send their traffic through the router over the 36 kbps link. This seems like a reasonable choice given the larger bandwidth of this route and the delay that would be associated with data ferrying. However, S2's operators are likely not aware that there will be a 30-minute period where S1 is also sending data through the router. In addition, without an NTO, the router itself has no way of knowing the relative priorities of messages coming from S1 and S2. Even if the router is using a protocol such as Differentiated Services [17], S1 and S2 have not been directed what forwarding behaviors to put in their packet headers.

Next, the same scenario is envisioned with an NTO process in place. In the planning phase, analysts expose the potential for the combined data rates from S1 and S2 to overwhelm the router's queue. At this point, there are multiple courses of action available.

1. Increase the bandwidth on the links through the router.
2. Deploy a second router.
3. Turn off the data flow from S2 (neither store nor send) during the contention period.
4. Use the AWACS as a data ferry during the contention period.
5. Give the router a priority queue and mark the packets from S1 and S2 with relative priorities.

Options 1, 2, and 5 are rejected because they require a physical change to or addition of equipment. Option 3 is rejected because even though the data is of lower priority, the loss of data should be avoided. Since it is known that the information from S2 is not time-sensitive, the delay associated with option 4 is tolerable. This course of action approval is made at the operation level and given as a directive to the operators of S2 (and also of the AWACS) through the NTO. The increase in end-to-end delay for S2 is justified by the increased reliability of the overall network.

Design

The scenario is implemented in the open source ns-2 simulator (version 2.29) [18]. Traffic is generated with constant bit rate generators using UDP. The data rates at which S1 and S2 send packets can be achieved in many ways by adjusting packet size and the interval between packets. One extreme way would be to have S1 send a 3825-byte packet once per second and for S2 to send a 900-byte packet once per second. To make the simulation more realistic, we consider packet sizes of 8 bytes through 64 bytes in steps of 8 bytes for both sources. The time interval between packets is adjusted accordingly. Random jitter is added to the traffic generators to prevent having consistent simultaneous arrivals at the router of packets from the two sources. Because of the random nature of the traffic, results are gathered using 30 different random seeds. This necessitates 1,920 simulation runs each for the two versions of the scenario (without and with NTO). The various connections are implemented as simplex-links with 5ms delays and DropTail queues (with default buffer size). No noise or signal fading is simulated.

For the scenario with no NTO process in place (S2 using the router), S2 generates traffic for an hour with S1 generating traffic for the 30-minute span from 10-40 minutes. For the scenario with an NTO process in place (S2 using the AWACS), the timing is more complicated. S2 needs to send 30-minutes worth of data within a 10-minute window of opportunity. To allow some time for connections to be established between S2 and the AWACS, we arrange for the data to be sent in only 8 minutes. Consequently, 12.96 Mbits (or 1.62 Mbytes) of information in 8 minutes corresponds to an increased rate of 27 kbps. The worst case for end-to-end delay happens when the LOS contact between HQ and the AWACS ends just prior to when the LOS contact between S2 and the AWACS begins. This is the situation we model. S1 generates traffic for the 30-minute

span from 10-40 minutes. S2 sends traffic to the AWACS in two 8-minute intervals from 0-8 minutes and from 30-38 minutes. The AWACS relays the traffic to HQ in two 8-minute intervals from 22-30 minutes and from 52-60 minutes.

6. RESULTS

We begin with results from the scenario with no NTO. In Figure 6, we tabulate the mean percentage of packets that get dropped at the router from the two sources, broken down by packet size. Cells are shaded based on their value. The higher the percentage, the darker a cell is shaded. In Figure 6(a), we see that S1 experiences less loss when packet sizes are relatively similar. S1 suffers the least loss of 3.4518% when S1 packets are 40 bytes and S2 packets are 56 bytes. The largest loss of 4.7082% is when S1 packets are 32 bytes and S2 packets are 8 bytes. Conversely, S2 experiences more loss when packet sizes are relatively similar, as seen in Figure 6(b). S2 suffers the least loss of 2.4988% when S1 packets are 64 bytes and S2 packets are 16 bytes. The largest loss of 5.1471% is when S1 packets are 16 bytes and S2 packets are 24 bytes. The confidence intervals on these means are pretty tight. Figure 7 shows the 95% confidence intervals for the means when fixing the size of one source's packets. Figure 7(a) has S1's packets fixed at 56 bytes, and Figure 7(b) has S2's packets fixed at 56 bytes. These graphs are typical of the confidence intervals for all mixtures of packet sizes.

		Source 2 Packet Size (bytes)							
		8	16	24	32	40	48	56	64
Source 1 Packet Size (bytes)	8	3.51	3.56	3.78	3.95	4.05	4.04	4.42	4.40
	16	4.44	3.51	3.48	3.57	3.68	3.78	3.85	3.96
	24	4.67	4.18	3.51	3.46	3.49	3.56	3.65	3.71
	32	4.71	4.42	4.02	3.51	3.46	3.47	3.51	3.56
	40	4.68	4.59	4.29	3.92	3.50	3.46	3.45	3.49
	48	4.64	4.67	4.43	4.17	3.86	3.50	3.46	3.46
	56	4.63	4.69	4.56	4.33	4.07	3.80	3.50	3.45
	64	4.67	4.70	4.63	4.44	4.23	4.01	3.74	3.49


		Source 2 Packet Size (bytes)							
		8	16	24	32	40	48	56	64
Source 1 Packet Size (bytes)	8	5.07	4.94	4.48	4.11	3.91	3.91	3.08	3.17
	16	3.08	5.07	5.15	4.93	4.70	4.47	4.35	4.09
	24	2.60	3.62	5.06	5.15	5.06	4.92	4.73	4.61
	32	2.51	3.08	3.95	5.04	5.14	5.13	5.03	4.91
	40	2.59	2.73	3.37	4.16	5.06	5.13	5.15	5.07
	48	2.66	2.58	3.09	3.63	4.28	5.02	5.09	5.12
	56	2.64	2.52	2.83	3.27	3.82	4.39	5.00	5.09
	64	2.58	2.50	2.65	3.07	3.46	3.93	4.47	5.00

(a) Mean percentage of packets originating at S1 dropped at the router


(b) Mean percentage of packets originating at S2 dropped at the router

Figure 6: Mean percentage of packets dropped at the router for scenario with no NTO

It is important to note that for no combination of packet sizes does either source experience an acceptable amount of loss. The QoS is degraded significantly for both sources. The high-priority video feed losses from S1 are especially troubling. The suggested tolerance for data loss is below 1% for high-quality audio-video streaming and 2-3% for two-way interactive audiovisual services [19]. These tolerances are all surpassed for S1's data stream.



(a) 95% confidence intervals for the mean percentage of 56-byte packets originating at S1 dropped at the router




(b) 95% confidence intervals for the mean percentage of 56-byte packets originating at S2 dropped at the router

Figure 7: 95% confidence intervals for the mean percentage of packets dropped at the router for scenario with no NTO

		Source 2 Packet Size (bytes)							
		8	16	24	32	40	48	56	64
Source 1 Packet Size (bytes)	8	4.0102	4.0062	4.0054	4.0046	4.0041	3.9996	3.9911	4.0039
	16	4.0032	4.0089	4.0116	4.0028	4.0032	3.9984	4.0073	4.0003
	24	4.0057	4.0018	4.0018	3.9970	3.9946	3.9937	3.9921	3.9958
	32	4.0036	3.9960	3.9997	3.9992	3.9974	3.9982	3.9964	3.9895
	40	4.0080	3.9910	3.9978	3.9992	3.9992	3.9955	3.9939	3.9956
	48	4.0028	3.9993	4.0031	3.9984	3.9919	3.9843	3.9844	3.9892
	56	3.9962	3.9977	4.0053	3.9884	3.9894	3.9891	3.9826	3.9767
	64	3.9972	3.9973	3.9976	3.9991	3.9834	3.9819	3.9750	3.9723

(a) Mean percentage of total bytes dropped at the router



(b) Summary statistics on percentage of total bytes dropped at the router over all 1,920 runs

Figure 8: Mean percentage of total bytes lost at the router for scenario with no NTO

Since UDP does not retransmit lost packets, it may be of interest to see what the total loss from both sources is in terms of bytes. The table in Figure 8(a) shows the mean percentage of total bytes lost from both sources at the router, broken down by packet sizes. The cells are again shaded based on their value, but the range of values is rather narrow. Percentages range from

3.9723% to 4.0116% with a trend of higher percentages occurring when packet sizes are small and lower percentages occurring when packet sizes are larger. Figure 8(b) provides a summary of statistics on the percentage of total bytes dropped at the router over all 1,920 simulation runs. The overall average is a loss of 3.9967% of the total bytes sent, with a very narrow 95% confidence interval of 3.9947% to 3.9986%.

The results from the simulations run for the scenario with an NTO are very simple. No packets were lost from either source, regardless of the mix of packet sizes. Also, all data sent from S2 to the AWACS was subsequently sent from the AWACS to HQ. Bear in mind that no noise or signal fading was simulated; however, loss from those sources would mainly apply to the lower priority S2 data.

7. CONCLUSION

The main conclusion that can be drawn here is that decisions made with the assistance of an NTO process could positively impact the overall QoS of the GIG. The current situation is one of local scope. With no other information, given the choice between a higher bandwidth route with small delay and a lower bandwidth route with large delay, the first choice seems to be the better option. It is a greedy choice in the sense that it appears to offer the best service for S2. Having an NTO process in place allows for a more global view of the situation. It would allow for less obvious, but in the long run, better decisions to be made. Here, the second option can be seen as the better choice for it allows the high-priority data from S1 to flow uninterrupted, and all of the data from S2 still reaches the destination in exchange for additional delay. Not only is S2 sacrificing speed for the benefit of S1, but S2 also reaps the benefit of not losing any packets itself.

8. REFERENCES

- [1] U.S. Department of Defense, *Global Information Grid (GIG) Overarching Policy*, DoD Directive 8100.1, 2002
- [2] Camp, T., Boleng, J., & Davies, V. "A survey of mobility models for ad hoc network research" *Wireless Communication & Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, vol. 2, pp. 483-502, 2002
- [3] Bettstetter, C., Resta, G., & Santi, P. "The node distribution of the random waypoint mobility model for wireless ad hoc networks" *IEEE Transactions on Mobile Computing*, vol. 2, pp. 257-269, 2003
- [4] Ranne, W.K. & McKee, Jr., L.K. "Global Information Grid Netops Tasking Orders (GNTO) white paper", Smithfield, VA, 2006
- [5] McKee, Jr., L.K. Personal e-mail, July 9, 2009
- [6] Stookey, D.E. "A notional battlespace for simulating and testing dynamic wireless networks", Technical Report, Air Force Institute of Technology, WPAFB, OH, 2006
- [7] Pecarina, J.M. "Creating an agent based framework to maximize information utility", M.S. thesis, Air Force Institute of Technology, WPAFB, OH, 2008

- [8] Weida, J.A., *Air Force Guidance Memorandum on Command and Control (C2) of the AF-GIG*, AFGM13-01, Washington D.C., 1 Sep 2009
- [8] Imperial, M.J. Personal e-mail, December 18, 2009
- [9] Scott, W.B. & Covault, C. “High ground over Iraq” *Aviation Week*, The McGraw-Hill Companies, 2003
- [10] Chairman, Joint Chiefs of Staff, *Command and Control for Joint Air Operations*, Joint Publication 3-56.1, Washington D.C., 1994
- [11] Secretary of the Air Force, *Air and Space Operations Center*, AFOTTP 2-3.2, Washington D.C., 2004
- [12] The MITRE Corporation Technical Staff, USMTF Message Browser Help, 2004 Baseline Edition, 2004
- [13] Defense Information Systems Agency (Online), Information Systems (J6), www.disa.mil/jsc/info_systems_j6.html, January 14, 2010
- [14] Symetrics Industries (Online) IDM-302 Product Brochure, www.symetrics.com/content/products/idm/index.aspx, January 13, 2010
- [15] Raytheon Company (Online) AN/ARC-164 UHF Airborne Radio Product Data Sheet, www.ratheon.com/capabilities/rtnwcm/groups/ncs/documents/content/rtn_ncs_products_arc164_pdf.pdf, January 13, 2010
- [16] The Antenna Professionals (Online) UH-408 Antenna, www.rami.com/military/air/antenna-info.cfm?pid=49, January 13, 2010
- [17] Cisco Systems (Online), DiffServ – The Scalable End-to-End QoS Model, www.cisco.com/en/US/technologies/tk543/tk766/technologies_white_paper09186a00800a3e2f_ps6610_Products_White_Paper.html, January 13, 2010
- [18] University of Southern California Information Sciences Institute (Online), The Network Simulator – ns-2, www.isi.edu/nsnam/ns/, January 13, 2010
- [19] Miras, D., “A Survey of Network QoS Needs of Advanced Internet Applications (Working Document)”, Internet2 QoS Working Group, 2002

9. AUTHOR BIOGRAPHIES

Matthew D. Compton is a Captain in the U.S. Air Force stationed at Wright-Patterson Air Force Base, Ohio where he is working on a Ph.D. in computer science at the Air Force Institute of Technology (AFIT). His research interests lie in the areas of topology control, mobile ad hoc networks, and polymorphic networking. Captain Compton earned an MA in 1998 and a BA in 1996, both in Mathematics from the University of Oklahoma.

Kenneth M. Hopkinson is an Assistant Professor of Computer Science at AFIT. His research interests are in the areas of fault-tolerant and reliable distributed systems, networking, and simulation. Dr. Hopkinson received his Ph.D. at Cornell University in 2004, his MS from Cornell University in 2002, and his BS from Rensselaer Polytechnic Institute in 1997, all in Computer Science.

Gilbert L. Peterson is an Associate Professor of Computer Science at AFIT. His research interests include the areas of artificial intelligence and robotics. Dr. Peterson received his Ph.D. in Computer Science in 2001, his MS in Computer Science in 1998, and his BS in Architecture in 1995, all from the University of Texas Arlington.

James T. Moore is a Professor of Operations Research at the Air Force Institute of Technology. His research interests include optimization, integer programming, and meta-heuristics. Dr. Moore received his Ph.D. in Mechanical Engineering at the University of Texas Austin in 1998, his MS in Operations Research at AFIT in 1981, his MBA from the University of Wyoming in 1978, and a BA in Mathematics from the University of Oklahoma in 1974.