

15th ICCRTS

“The Evolution of C2: Where Have We Been? Where Are We Going?”

Paper ID 114_S

Title: Mission-Related Execution and Planning Through Quality of Service Methods

Authors:

Vinod D. Naga¹, Maj, USAF – DOCTORAL STUDENT – vinod.naga@us.af.mil
John M. Colombi¹, PhD – john.colombi@afit.edu
Michael R. Grimaila¹, PhD – michael.grimaila@afit.edu
Kenneth M. Hopkinson², PhD – kenneth.hopkinson@afit.edu

Point of Contact:

Vinod D. Naga¹, Maj, USAF – DOCTORAL STUDENT – vinod.naga@us.af.mil
937-255-3636 x7126

1- Air Force Institute of Technology
Dept. of Systems and Engineering Management (AFIT/ENV)
2950 Hobson Way
Wright-Patterson AFB, OH 454333

2- Air Force Institute of Technology
Dept. of Electrical and Computer Engineering (AFIT/ENG)
2950 Hobson Way
Wright-Patterson AFB, OH 454333

Copyright © 2010 by Vinod D. Naga, John Colombi, Michael R. Grimaila and Kenneth Hopkinson.
Published and used by ICCRTS with permission. 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.

Mission-Related Execution and Planning Through Quality of Service Methods

Abstract

Moving information across the distributed Global Information Grid is critical to Command and Control. Often, the flow of this information is challenged by other flows and dynamics common to military operations with tactical communications constraints. Quality of Service (QoS) is one of the available mechanisms to manage these resources and improve the state of these critical information flows. In order to manage the resources in concert, a framework is needed to bring mission emphasis and priority to the QoS deployment. Besides improving the performance of missions, emphasizing missions in QoS arrangements will enhance mission assurance since failure of a single mission will limit success of all missions which have dependence. This paper provides a brief review of QoS concepts. We emphasize the importance of decomposing missions and then allocating resources by mission area. We investigate methods for using mission and context to establish effective Service Level Agreements (SLA) leading to tailored QoS. The paper ends by demonstrating a mission-oriented QoS experiment.

Introduction

Quality of Service (QoS) is a construct by which performance can be specified, apportioned, measured and maintained across large scale networks. This capability is ultimately useful if QoS operations across the networked system of systems work towards the same goals as the organizations which implement the QoS mechanisms. With this synergy, remote command and control (C2), information retrieval as well as streaming of sensor data over the networked system of systems may be managed for mission success and under dynamic conditions. In this paper, we review the use of networks for military objectives and briefly outline some QoS principles. We then examine the functional decomposition of organizational objectives and functional allocation of these objectives to physical elements. Aligning network elements to mission objectives provides a basis to implement QoS mechanisms for the benefit of overall mission accomplishment. We also catalogue research which has built mission-aware and context-aware QoS methods for Network-Centric systems.

Network-Centric Warfare (NCW) and Network-Centric Operations (NCO) are concepts introduced by Admiral Cebrowski and others (Alberts, Garstka & Stein 1999). These concepts continued to carry influence within the Department of Defense (DoD) research and operational communities. The growth of information networks has been encouraged by shrinking budgets and the introduction of information-intensive systems including a renewed appreciation for information types from diverse sources. Research in the field of computer networks, wireless networks and mobile sensors have established hundreds of algorithms and variations on algorithms to schedule and manage network resources. These algorithms optimize or use heuristics based on network performance, or other quality attributes such as bit-error-rate, throughput, packet arrival rate, end-to-end delay, and traffic priority, among others. While these algorithms focus on performance, they are often motivated by requirements dictating mission performance. Intermediate requirements such as stability, reliability and adaptability take into consideration stakeholders' missions and provide goals for network performance. Fortunately,

QoS mechanisms offer a great deal to help with specifying, establishing and achieving these network requirements.

Another challenge in military tactical networks is the variation in configurations of elements depending on the set of missions for that area of operation. These elements, many of which are mobile, may consist of various combinations of sensors, shooters, decision-making platforms and communications nodes. Besides addressing the longterm planning challenges of interoperating when this constellation is not fully known a priori, there is the additional challenge of establishing and maintaining effective communications among these elements when the core of the network is also mobile (Graham 2008). Other changes to the network may manifest themselves similarly to changes with physical movement. Node failures, denial of service, jamming, interference, link dropouts and other unexpected conditions require network adaptation and reprovisioning of resources much like the mobility case. QoS enhancement techniques must factor in the complex problem of managing connectivity to optimize across the set of potentially competing missions.

Anunciado proposed an approach to establishing dynamic Command and Control (C2) nodes in an enterprise environment (Anunciado 2006). The execution of macro and micro policies by network components is analogous to the significance of requirement interdependencies. Macro policies are stipulated by stakeholders and set the high-level operation of the enterprise. Micro policies then follow by directing specific nodes to interact in order to realize the goals of the macro policy. This concept of aggregate actors who have the interests of the enterprise in mind, but take action in directing local elements is often repeated throughout the literature (Estrin et al. 1999; Intanagonwiwat et al. 2002; Intanagonwiwat et al. 2003; Ye, Lai and Farley 2003). QoS and mission characterization can be the method to specify and execute macro and micro policies.

Mission Organization and Decomposition

Functional Allocation

Network Centric Operations (NCO) has motivated many studies which looked at assessing mission effectiveness against the communications necessary to execute the mission. The Boeing Company performed a study looking at Command and Control, Intelligence, Surveillance and Reconnaissance (C2ISR) requirements for communications channels by performing a functional allocation among entities in that system of systems (Carson et al. 2005). They augmented time-sensitive targeting (TST) models with parametric representations of communication systems and measured the effectiveness of the TST missions. Through empirical methods they established upper and lower boundaries for communication performance impact on mission performance. The researchers concluded that certain ISR elements were ineffective for the mission unless certain data rates were available.

Other studies examine the functional allocation process and look for ways to improve it. Related functions which are readily automated include strategy to task decomposition, prioritization and task scheduling, and establishing metrics and performance assessment (Brown 2001). If properly configured automated decision tools perform these functions for network configuration beforehand it effects better performance on the network. A systems engineer may perform functional allocation and functional decomposition as key activities for designing a system. Synthesis and decomposition are specific functions which have an algorithmic formulation that

assist in developing capabilities by aggregating system elements and recursively partitioning a system into distinct entities (Ravichandar, Arthur & Broadwater 2006).

One method of function decomposition is the Function Analysis Systems Technique (FAST) (Wixson 1999). Here, functions and components needed to accomplish a mission are identified. The technique incorporates a model known as the FAST Model to visually depict a process flow of functions and reasons or purposes against a waterfall timeline depicting precedence. FAST is compatible with the method known as Value Engineering where value of additions and functions is quantified.

Decomposing Military Missions

The size and complexity of military endeavors requires that in order for military missions to be successfully accomplished, they must be decomposed. Militaries also assemble subordinate organizations which are outfitted to perform anticipated tasks. The United States military prescribes force and resource decomposition to accomplish joint tasks and to achieve joint missions in a variety of manuals for operations (U.S. Department of Defense 1995; 2002; 2003; 2008; 12th Air Force 1996) and also in prescribing systems engineering processes (U.S. Department of Defense 2004; Sage 1992). These guides largely promote the Strategy-to-Objective-to-Mission-to-Task sequence when planning military operations.

Other militaries follow similar procedures, evidenced in the military mission decomposition study conducted for a NATO Defence Requirements Review (Armstrong 2005). Another military example of identifying superior and subordinate missions involves the Defense Information Systems Agency (DISA). DISA provides communications capabilities and infrastructure for much of the DoD through a complex planning and organizing activity (U.S. Department of Defense 2004; 2005). This activity includes satellite communications (SATCOM) which is costly but can provide modifiable communications quickly even where prior terrestrial infrastructure does not significantly exist. SATCOM often carries the full range of strategic to tactical communications and includes both military and leased use of commercial satellites. DISA manages the SATCOM access program by identifying, documenting and prioritizing requirements in major layers, detailed with multiple attributes and individual priorities (U.S. Department of Defense 2007). A similar process is in development for the very agile Battlefield Airborne Communications Node (BACN) (Richards 2009).

Since flexible, agile and capable communications systems exist and are being further developed for military applications, it is of even greater importance that decompositions of military missions be applied directly to the communications resources and devices of the GIG.

Quality of Service

Defined

Quality of Service (QoS) has many definitions which depend on an application's QoS context. In most cases, QoS refers to a level of performance which is managed in a three-part process within the structure of Service Level Agreements (SLA) containing Service Level Specifications (SLS) on an information technology (IT) network (Doshi et al. 2006). The first part begins when the level of performance is expressed as a need and communicated from application to network in the form of a requirement. Next, the network then uses the requirement to schedule and reserve

resources. Finally, the network responds back to the application confirming the requirement and communicating the reservation of needed resources to fulfill the requirement (Ghosh et al. 2005). After these steps, the application may operate within the bounds of its stated requirements with assurances that the network will support its operations. There are many nuances to this process involving how the requirements are communicated, what levels of service may be requested and promised, how the reservation is maintained and managed as both new requests arrive and previous requests age, and how various network entities participate in the process (Yurcik and Banerjee 1995). A well-designed process includes graceful degradation and upgrading of service so applications relinquish and take resources in a coordinated manner benefitting the mission.

Quality of Service provides a method by which applications can manage and track performance of their data streams, adapt and use the network as it changes and institute fairness and resource sharing policies. These sharing policies allow applications to request the resources they need as well as allow the network to manage its own resources.

Quality of service in the network context is negotiated using various metrics. Bandwidth is the common term most frequently attributed to application requirements and network capabilities. Bandwidth refers to the size of a data channel in terms of how much information can transit the channel per unit time. Actual bandwidth at a given time is sometimes difficult to measure. Therefore, other metrics can infer performance in measurable terms more appropriately than bandwidth. These metrics may include end-to-end delay or latency, end-to-end packet delay variation, bit error rate, packet loss rate, packet loss ratio, queuing delay and queue size and sometimes remaining available bandwidth.

Service Level Agreements and Specifications

Service Level Agreements (SLA) provide the chief means by which QoS is established by providing a single place where requirements and user expectations as well as provider promised service levels are recorded (Doshi et al. 2006). The user community, those who are engaged in a common mission, enacts requirements on the network in the Mission Service Level Agreement (MSLA). The MSLA leads to individual SLAs with and among the different providers. The MSLA includes mission-related, technical and geographical information. The SLA includes technical specifications (SLS) of network performance. The SLA is the result in satisfying the optimization problem of satisfying many requirements which often exceed available resources. An SLA which maximizes a mission effectiveness functions is the ideal driver of QoS mechanisms.

Service Quality

Quality of Service may also exist in other contexts. Any system or SoS which has interdependencies may exercise QoS-like activities. The definition provided above pertains to IT systems. However, service quality is the originating concept of QoS and is the level of performance which one entity expects from another, including non-IT SoSs.

Service quality may also be reflected in the context of a system's purpose or an organization's mission. Putting level of service values and requirements in a similar context to overall organizational performance provides a common framework by which service quality may be provided for the benefit of overall performance. As goals are stipulated by the overall organization, the cost to achieve the goal will also be expressed and is subdivided amongst the

participating sub-organizations. The cost provides a reference by which the performance of the sub-organization can be measured within the framework of the overall organization's goals. Likewise, the overall organization's goals may be connected to the service quality provided by the sub-organization. Reaching those goals provides justification to incur costs or enhance certain services. These service and cost relationships between organizations are similarly represented between systems in a SoS framework. Just as they do for organizations and teams, the purposes or missions of the system of systems may also lead to service quality determinations.

Features Delivered and Requirements

Several features are provided by Quality of Service mechanisms. These include predictable service delivery (Kurose 1993), adaptability, fairness (Sisalem 1997), and network resource management.

Routing. The first step in establishing a reservation for the purpose of QoS is identifying the most efficient route or path between source and destination. Routers along this path must respond to requests and reserve resources. In addition, when the network topology changes due to outages or if nodes move out of communications range or if traffic leads to congestion, rerouting may be needed and reservations established along new routes. The capability to identify efficient routes, maintain those routes and establish more efficient routes is a critical prerequisite to achieving QoS (Yurcik and Banerjee 1995).

Signaling. Under the integrated services (*intserv*) model, signaling refers to the request and control information which is transferred between applications and routers in order to establish QoS (Firoiu et al. 2002; Kurose & Ross 2007) commensurate with the overall SLA. The request for bandwidth or a level of service is communicated in the form of a requirement through a signaling protocol. Routers along the efficient path adjust the request based on their capacity and move it to the destination. The destination acknowledges back to the source confirming the QoS arrangement is in place (Zhang et al. 1993).

Intserv signaling can be categorized as in-band or out-of-band. In-band signals are embedded in packets carrying the data for which they are signaling. Out-of-band signals transit in their own packets and therefore may not use the same route or move as efficiently as the data. Either type of signaling can be used to establish, maintain or tear-down a QoS connection.

IP Header Marking. Under the *diffserv* model, the IP header of the datagram containing data packets contains information to deliver QoS. Differentiated Service Code Point (DSCP) or other structures identify the class of service to which the packet belongs. Routers read the DSCP and make operations decisions on each packet versus others and their DSCP. Establishing classes and assigning applications and their data flows to those classes are policy decisions managed as part of an admission policy for QoS (Firoiu et al. 2002; Kurose & Ross 2007).

Admission Policy and Control. In order to improve the level of service for select applications and data streams, network entities must be able to restrict traffic. This ability to restrict traffic is known as admission control. The rules by which the restrictions are implemented involve selectivity and are summarized in an admission policy. Figure 1 demonstrates how controls and policy follow from an established SLA and SLS. Within an overall admission policy, a QoS

scheme implements its own admission policy on network devices. This admission policy describes the QoS scheme and ensures that higher priority services can receive improved levels of service (Firoiu et al. 2002; Liebeherr, Patek & Burchard 2003). Admission policy defines rules which implement the SLA.

Admission control implements admission policy and is handled mainly by routers and other devices which manage traffic on the network (Mirhakkak, Schult & Thomson 2000). Routers prioritize packets based on the class (*diffserv*) or data stream (*intserv*) to which they belong then use priority to schedule them. Highest priority packets are passed immediately while lower priority packets are queued and eventually dropped if queuing exceeds available resources. If SLS items such as delay or throughput are not satisfied, then the QoS scheme must be revisited.



Figure 1 - SLAs are composed of SLSs. Admission policy and controls, based on SLSs, are executed in network devices.

QoS with Mission Context

Service Level Agreements for Mission

End-to-end QoS also must contend with the challenge of shared networks. Rarely does data travel solely on networks owned and managed by a single organization. In many cases, a data stream, and the request which generated it, will traverse a variety of networks between source and destination. These networks exist to provide transport for this data but ultimately perform better when requests are satisfied within the context of all known requests as well as all available resources. Likewise, the applications requesting the resources are better served when the responses to their requests are realistic and executable. SLAs provide a mechanism for managing this process on the Global Information Grid (GIG) (Doshi et al. 2006). User communities of interest (COIs) are defined by their common purpose and orient themselves with a mission. Network service providers (NSP) establish network service domains (NSD) which contribute resources for the user COI's mission. The SLA is the tool by which the user COI communicates their request and against which the NSP assigns the resources of the NSD. The user COI includes a mission planner who establishes a mission service level agreement (MSLA). The MSLA is then negotiated with a middleman, an intermediate party established by a central authority. From the approved MSLA, the central authority negotiates SLAs with and among various NSDs. Doshi provides descriptions of roles for these entities and aspects which should be included in the SLA (Doshi et al. 2006). The author also describes metrics at the MSLA and SLA levels which may serve as content for actual MSLAs and SLAs and provide measures gauging how well the network is performing.

Meeting the information flow requirement of all users is difficult because it requires knowing the information flow requirement for a particular user, all its competing users and their relative

importance. Often the needed source information can be found in the SLA. Investigating how users are organized into user COIs and then examining the user COI's missions allows for prioritizing competing requirements. This can also provide insight to compare user COIs with differing missions and differing service characteristics. These requirements may be captured in the form of information exchange requirements (IERs) like the MSLA and influence the formation of network SLAs. A systems engineering methodology invoking appropriate architectural information such as DoDAF (U.S Department of Defense 2009) may be used to develop the mission IERs, SLA and especially the final QoS provisioning as shown in Figure 2.



Figure 2 - Bridging Mission Characterization to QoS Algorithms

Translating mission IERs to physical nodes and the QoS mechanisms is a challenge in distributed systems like those found on the GIG. A layered view of the system (Wong-Jiru 2007) may be well-suited to follow the layered view of the process depicted in Figure 1. The distributed real-time embedded systems perspective provides a segmentation of systems and sub-systems. This segmentation permits development of a layered perspective of missions and systems, applications and resources. The layered perspective provides a way to attack the challenge of effecting QoS provisioning starting with a declared set of missions. Methods such as those used for business process modelling (BPML), Universal Modeling Language (UML) and Systems Modeling (SysML) may be used to accomplish the translation of mission IERs to physical nodes. These activities involve difficult and non-trivial steps.

Distributed Real-Time Embedded Systems (DRE)

There is an important distinction between single-node systems and distributed systems. In a single-node system, the environment is largely closed and all the conditions with which an embedded system must contend are contained within the system and the single environment in which the system operates. Distributed systems involve portions of a system which are embedded among a number of nodes. These nodes often exist in various environments resulting in varying conditions and some unpredictable situations for the distributed systems embedded in them. Distributed real-time embedded systems (DRE) (Loyall et al. 2005; Loyall and Schantz 2008) exhibit system of systems characteristics and are designed specifically to deal with complexities inherent to such environments.

Many previous efforts underscore the importance of designing and managing end-to-end QoS in order to achieve application or user performance requirements. Achieving end-to-end QoS in a DRE is a challenge but one approach is to divide the QoS and resource management into layered views (Loyall and Schantz 2008). Dividing QoS management into mission/system layer, application string layer and resource layer views allows the management effort to also be divided amongst these functions. The mission/system layer view is a high level view which includes mission goals, available system applications and all available resources. The mission/system layer view is also aware of the importance of applications and how to allocate resources to achieve the mission goals. The application string layer understands the resource needs to

achieve an application string's end-to-end needs. This layer recognizes whether applications have the ability to share resources with other applications as well as identifies those which consume all available resources much like the network SLA. The resource layer view manages access to resources such as CPU, memory and local bandwidth. At the resource layer, the goal is to allow resources to be requested and then meet those requests in a fast, effective process without focusing on applications or mission priorities. The resource layer is the heart of any good QoS mechanism. A framework which uses these layers in concert facilitates the critical communications needed between layers to manage movement of prioritized information (Mitchell et al. 2008). A generalized representation of the SLA-to-Admission Control flow of Figure 1 is provided in Figure 3. Here, resources are shared according to policy based on sub-layers of application interdependencies and mission service level agreements.



Figure 3 - Resources are shared according to policies based on application interdependencies. The application agreements follow from requirements identified in mission service level agreements.

Three approaches to delivering this framework include negotiation, hierarchical and static (Loyall and Schantz 2008). Negotiation involves applications requesting, using and returning resources to the larger system. Systems operating under the negotiation approach must have disincentives at the resource layer to prevent requesting and consuming too many resources. Hierarchical approaches require a management entity to determine importance and hierarchy which are reflected in the system policies. These policies govern application behavior and access to resources. The hierarchical approach is useful in many network deployments which follow a subnet-to-subnet configuration. In (Ghosh et al. 2005) a QoS-based Resource Allocation Model is implemented in a scalable fashion to govern flows over a large-scale network which has a hierarchical topology. In the static approach, all allocation decision-making is done beforehand and explicitly programmed into the system. The static approach appeals to dynamic systems such as DRE systems only to the extent to which the dynamics have been anticipated and programmed into the system.

QoS Management System

Just as a system's applications can be embedded in distributed nodes to form DREs, users are often engaged in multiple communities of interest (COIs). While a mission defines the boundary of a user COI, an individual user may participate in multiple missions and consequently be a member of multiple COIs. These COIs are dynamic objects with members joining and leaving. The influence of these members can change the focus of the COI. Information management systems are often integrated with COIs to manage, control, allocate and utilize assets and entities which are also part of distributed systems. A QoS Management System (QMS) which connects information management systems and COIs to provide enhanced QoS driven by mission first is

described in (Loyall et al. 2007). The QMS described includes a connectivity monitor and the following core modules: System Resource Manager, Local Resource Manager and QoS mechanisms. The connectivity monitor provides visibility into resource usage to the higher level QoS managers. The information for the connectivity monitor flows up through the layered system so intermediate QoS performance measures are available for tasks such as predicting impacts due to QoS changes. Alternative QoS monitors are provided in (Jiang et al. 2000) where placement of monitor and choosing monitor type and synchronizing returned information to data flows are explored to great depth. The System Resource Manager manages QoS policies and factors in overall mission goals, available system resources and policies with respect to utilizing those resources. The Local Resource Manager receives the resource allocations from the System Resource Manager and implements the allocations by managing and operating QoS mechanisms. The QoS mechanisms control interfaces, create messages and reformat data to achieve the resource actions.

A separate QoS Management solution offers a switching hybrid system model which provides insight into the operation of the real system (Abdelwahed et al. 2003). With this model, QoS decisions can be made and implemented in a static fashion with updates coming as the model dictates relevant changes to the system. In this context, hybrid represents the case where both time and event driven dynamics affect the system.

Stanley et al. (2005) introduced the Mission Service Automation Architecture (MSAA) which correlates information flows to operational capabilities providing the ability to prioritize network traffic. MSAA incorporates the concept of IT service codes as defined in the Information Technology Infrastructure Library (ITIL) as a means to align IT services with customer mission requirements. The MSAA requires a configuration management database and collection of independent software agents to provide an understanding of the purpose of information flows. Data tagging is used to provide the needed visibility by inserting IT service management codes into the headers of network packets.

The Air Force Research Laboratory Information Directorate (AFRL/RI) conducts research in effective QoS management as well. They developed an Advanced Technology Demonstration (ATD) where the Airborne Warning and Control System (AWACS) MultiSensor Integration (MSI) tracker was influenced by track mission importance (Lawrence et al. 1997). Tracks were selectively updated by the MSI using time, precision and accuracy metrics to rate the tracks using user-defined policies and time-value/time-utility function profiles (Cho 2006). The QoS manager uses the resulting metrics to make decisions on which tracks should absorb resources and on which tracks certain tracking applications should continue to operate. This was very much a hierarchical management approach acknowledging a system divided into mission, application string and resource layers. The study found improved performance on key tracks when track volume increased to resource limits and acknowledged this as serving the goals of adaptive resource management, QoS-based graceful degradation and dependable distributed real-time applications.

A promising negotiation framework in development at AFRL/RI is QoS middleware which includes the resource identification and management capabilities for use by other systems. This QoS middleware concept further extends the DRE notion.

QoS Middleware

An often preferred method to implement the layered views and management of resources is through deployment of middleware (Dasarathy et al. 2005; Loyall et al. 2005). AFRL/RI provides the leadership on Joint Battlespace Infosphere (JBI) under which various capabilities are provided in middleware (Loyall et al. 2005; Mitchell et al. 2008; Loyall et al. 2009). AFRL/RI conducts premier research in the areas of command and control, intelligence and information systems processing. In 1997, they helped execute the DARPA Quorum program which looked for novel solutions and improvements to QoS challenges.

Middleware such as CORBA or a service-oriented architecture (SOA) solution can be deployed within each node of a system of systems and carry out many functions critical to managing infrastructure and resources. The middleware contains algorithms which work at each of the three level views and manage access to resources for the node and the applications which run on it. Bryan et al. (2005) presented a method to integrate the scheduling of resources, management and monitoring of resource using solutions wrapped in CORBA.

Providing infrastructure such as middleware provides some QoS solutions for distributed military systems on the GIG. Challenges still remain for the dynamics and mobile situations which are commonly part of military deployments. A recent effort to provide QoS services in middleware for tactical military users is Quality of Service Enabled Dissemination.

QoS Enabled Dissemination (QED)

Quality of Service Enabled Dissemination (QED) builds on previously developed QoS tools and entities as well as the CORBA-based middleware to meet quality requirements of users and their missions reliably, in real-time and with resilience to dynamics of tactical environments (Mitchell et al. 2008; Loyall et al. 2009). QoS Enabled Dissemination was funded in-part by AFRL/RI to deliver services which fit within JBI and Net-Centric Enterprise Services (NCES) for the GIG. Often, information management and information dissemination through publish-subscribe-query are decoupled processes. QoS Enabled Dissemination is designed to provide priorities and direction for the information management as well as context for queries and provided information so they can be prioritized and delivered. The QED system is designed to provide

- timely delivery of information needed by tactical users in mobile scenarios
- tailored and prioritized information based on mission needs and importance
- rapid response to priority shifts and unfolding situations
- operations robust to failures and intermittent communications.

The QED prototype provided a QoS administration interface, aggregate and local QoS managers, differentiated queues for submission and dissemination services and client-side monitoring of QoS activities. Demonstration results provided motivation for furthering context and mission aware QoS work (Loyall et al. 2009). Further work needed in this genre of study include

- optimization heuristics to achieve aggregate QoS for publish-subscribe systems
- service level agreements to satisfy client preferences when clients are decoupled
- capture of actionable mission-level abstractions in policy

As a part of JBI and NCES, QED and the work that follows will provide QoS services for users of the GIG.

Mission-Oriented Quality of Service Performance

Design of the Experiment

Network Centric systems can be configured to suit the missions which use the network. By identifying the missions using the network, their relative priority and the missions' information exchange requirements, the mission manager can ensure the network delivers needed performance to the mission. A primary mechanism for accepting the configuration and delivering this performance is QoS. The following experiment demonstrates the performance loss if network QoS is improperly configured and also demonstrates the advantage of enacting a QoS configuration designed for a mission.

The network pictured in Figure 4 serves as the basis for this experiment which was conducted using OPNET Discrete Event Simulation. Server local area networks (LANs) on the left provide information at the request of users in the Center LANs on the right. A server LAN exists for the following mission areas: operations (ops), intelligence (intel), logistics (logs) and personnel. A Center LAN also exists for these four mission areas. For the purpose of this experiment, requests from a mission area center LAN go to the server LAN for the same mission area. Likewise, all responses from a mission area server go to the Center LAN for that mission area. Request traffic is infrequent and for the purpose of this experiment, considered negligible. The traffic of interest flows from a Server LAN to a Center LAN and this experiment will examine how to maintain good performance for the flows necessary for mission areas.

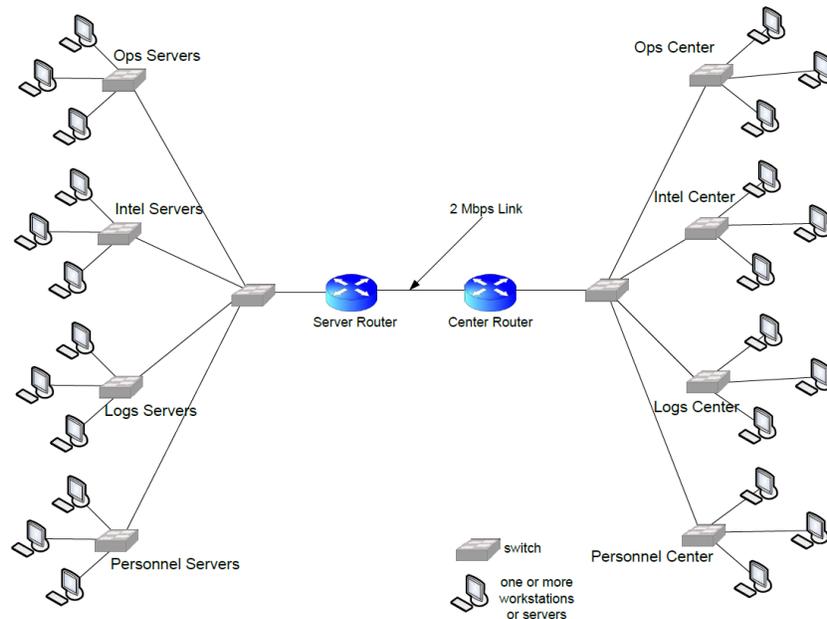


Figure 4 - Configuration of network for experiment. Information flows from Server LANs to Center LANs must transit the 2Mbps link. The routers implement PQ QoS and identify traffic using DSCP classes.

All links in the network except one have more than adequate capacity for the information flows they will carry. The critical link is depicted in the middle of the diagram in Figure 4. This link

has 2 megabits/second (Mbps) capacity and connects the Server Router and Center Router. No other paths are available between a Server LAN and Center LAN. Most networks, including SoS networks and systems operating as DRE on cooperative networks like the internet, are adequately provisioned in the core. If a capacity constraint exists, it is typically at the edge of the network where the user and/or the server is located. Networked SoS's such as those in military operations sometimes present an exception to this situation. A military network may have a limited-bandwidth connection at the edge but it may also have connection constraints deeper in the network due to limited resources, mobility requirements, interference, network attack and other such concerns mentioned in this paper's introduction. The limits of this experiment's 2Mbps link will simulate all civilian and military network situations in a general way. The constrained link represents inadequate capacity somewhere in the information flow's critical path. Use of this link can be managed effectively using QoS mechanisms which will be exercised by the applications at the edge and the routers connected by the link. More complex scenarios are possible which provide multiple paths, multiple QoS-enabled routers and constraints at both ends of the network. These scenarios will be explored in future work.

In this experiment, the military campaign proceeds through a series of phases. In each phase, a mission area is declared responsible for the primary objective of the phase. The phases for the campaign are each depicted in Table 1 along with the mission areas listed in order of priority.

Table 1 – Phases for military operational campaign with mission area priorities for each phase.

Campaign Phase:	Equal	Logistics	Operations	Intelligence
Priority 1		Logs	Ops	Intel
Priority 2	All Mission Areas	Intel	Intel	Ops
Priority 3		Ops	Logs	Logs
Priority 4		Personnel	Personnel	Personnel

In order to establish a performance baseline, the first campaign phase is named 'equal'. In this phase, all mission areas have equal precedence and priority. In the next phase, build-up of forces is the key priority so logistics takes precedence. The next phase represents commencement of combat operations and ops therefore takes precedence. In the last phase, an intelligence requirement is identified as critical and must be supported with all available resources. Having commander declared objectives and critical mission areas for the phases, the subordinate commanders detail information exchange requirements (IERs) in Table 2.

Table 2 - IERs for each campaign phase and mission area in priority order.

Campaign Phase:	Equal	Logistics	Operations	Intelligence
IER 1	Ops traffic: 300 kbps	Pre-combat logistical staging	Ops command and control	Streaming ISR feeds
IER 2	Intel traffic: 300 kbps	Real-time intel updates	Real-time intel updates	Ops command and control
IER 3	Logs traffic: 300 kbps	Pre-staging ops plans	Ops support logistics	Ops support logistics
IER 4	Pers traffic: 300 kbps	Health and wellness info	Health and wellness info	Health and wellness info

Using the IERs as an MSLA, the mission managers update the service level agreement (SLA) with SLS's listed in Table 3. Network engineers then design a QoS provisioning scheme which is described below.

Table 3 - Service Level Agreement (SLA) with specifications (most critical for campaign phase listed first) stipulated by mission managers for each campaign phase.

Campaign Phase:	Equal	Logistics	Operations	Intelligence
SLS1	Ops traffic: 300 kbps	Logs End-to-End delay < 0.1 sec	Ops E-to-E delay < 0.1 sec	Intel E-to-E delay < 0.1 sec
SLS2	Intel traffic: 300 kbps	Logs traffic received > 95%	Ops pkt delay variance < 0.2	Intel pkt delay variance < 0.1
SLS3	Logs traffic: 300 kbps	Ops End-to-End delay < 0.3 sec	Ops traffic received > 99%	Intel traffic received > 99%
SLS4	Pers traffic: 300 kbps	Intel End-to-End delay < 0.3 sec	Intel E-to-E delay < 0.1 sec	Ops E-to-E delay < 0.2 sec
SLS5		Personnel traffic received > 50%	Intel traffic received > 80%	Ops pkt delay variance < 0.3
SLS6			Logs traffic received > 50%	Ops traffic received > 99%
SLS7			Pers. traffic received > 20%	Logs traffic received > 25%
SLS8				Pers. traffic received > 20%

Mission traffic for the declared mission area in each phase is to be given highest priority among all mission traffic. It is also expected that mission traffic for the declared mission area will increase significantly upon entering that phase. The expected mission area traffic rates for each phase are given in Table 4.

Table 4 - Mission traffic flow rates and priorities for each mission area. Mission flow rates vary as campaign phases change. Configured priorities for mission areas on network devices lag campaign phase changes and are depicted here as well.

critical mission config. mission	equal		logs		ops		intel		configured priority
	equal	logs	logs	ops	ops	ops	ops	intel	
operations flowrate (kbps)	300	300	500	1400	1400	1400	1400	1400	operations 2
intelligence flowrate (kbps)	300	300	50	50	50	500	500	500	intelligence 1
logistics flowrate (kbps)	300	300	1600	800	800	800	800	800	logistics 3
personnel flowrate (kbps)	300	300	50	50	50	50	50	50	personnel 4

The configured mission area priorities for the QoS scheme and campaign phases are depicted in Table 4. This experiment assumes there is some lag in adjusting the network provisioning for all transitions except the first transition to logistics. The campaign phase - network configuration status (equal-equal, equal-logs, logs-logs, ops-logs, ops-ops, intel-ops, intel-intel) is also provided in the top row of Table 4. For some time at the beginning of each phase, the network runs in the previous configuration and then adjusts to a configuration better-suited for the primary mission area of that phase. The network achieves a steady-state and steady-state data is gathered during each step depicted by a column in Table 4 and representing a unique combination of campaign phase and network configuration.

This experiment serves to examine the upper bounds of performance on the constrained link and therefore only UDP data streams are utilized. Acknowledgements and backoff timers may reduce delivery rates so TCP connections and responses are not used in this simulation. Real operations utilize a combination of TCP and UDP packet streams so future simulations should introduce mission-oriented applications using TCP. Packet size averages 500 bytes and follows an exponential distribution. Packet arrivals have an average based on flow rate and are also exponentially distributed. (OPNET 2003)

This experiment implements QoS using four main assured forwarding (AF) DSCP classes plus the expedited forwarding (EF) class and priority queuing (PQ). In a PQ arrangement, an AF queue is serviced only when the higher priority AF queue is empty. The EF class is always

forwarded immediately. Routers have a 100 packet capacity queue for EF class and another 100 packet capacity queue for general, best-effort traffic. Packet capacity is 20, 40, 60 and 80 for the AF classes in rising priority. In the “equal” campaign phase, all mission areas are assigned to the same class. For all other phases, the declared mission area is given the primary class designation. The configured priority for each mission area is given in Table 4.

Each mission area is also granted a command link running at Integrated Services Digital Network (ISDN) speed of 128kbps. Each command link runs a constant UDP stream occupying this bandwidth while active. Each command link is assigned to the EF class and therefore receives top priority regardless of campaign phase. Combined, the command links occupy more than a quarter of the bandwidth on the constrained link. Simulations were run with and without the command link active.

In addition to the mission flows, each mission area also has a 200kbps link for low priority background traffic. These background traffic links also run a constant UDP stream but only receive best-effort service by all network devices including the QoS-enabled routers.

Experiment Results

The critical task in this experiment is for the Server Router to manage access to the constrained link using a QoS arrangement configured for campaign phase. Performance is measured against the SLA provided in Table 3. The challenge is evident in observing the figures in Table 5 and Table 6.

Table 5 - Command links inactive. Volume of flow for DSCP-labeled and non-labeled flows by campaign phase - network configuration steps. The final two rows convey the amount by which the link capacity is exceeded.

critical mission	equal	equal	logs	ops	ops	intel	intel
config. mission	equal	logs	logs	logs	ops	ops	intel
total mission flow (kbps)	1200	1200	2200	2300	2300	2750	2750
total prioritized flows (kbps)	1712	1712	2712	2812	2812	3262	3262
all traffic (kbps)	2512	2512	3512	3612	3612	4062	4062
% link exceeded by pri.	-14%	-14%	36%	41%	41%	63%	63%
% link exceeded by all	26%	26%	76%	81%	81%	103%	103%

Table 6 - Command links active. Volume of flow for DSCP-labeled and non-labeled flows by campaign phase - network configuration steps. The final two rows convey the amount by which the link capacity is exceeded.

critical mission	equal	equal	logs	ops	ops	intel	intel
config. mission	equal	logs	logs	logs	ops	ops	intel
total mission flow (kbps)	1200	1200	2200	2300	2300	2750	2750
total prioritized flows (kbps)	1200	1200	2200	2300	2300	2750	2750
all traffic (kbps)	2000	2000	3000	3100	3100	3550	3550
% link exceeded by pri.	-40%	-40%	10%	15%	15%	38%	38%
% link exceeded by all	0%	0%	50%	55%	55%	78%	78%

Table 5 lists flows and use of the link without the ISDN command links and Table 6 lists these figures with the ISDN command links. The fourth row of each table provides the amount by which the DSCP-labeled packets exceed the link and the fifth row provides this figure for all network traffic. As the mission-related traffic increases with ongoing campaign phases, use of the link becomes increasingly contentious.

The result of this overcrowding is evident when examining the use of the queues on the router. The EF class traffic (command links) uses Q5 and best-effort traffic uses Q_default. Whether command links are active or not, the lowest three classes of traffic have much of their packets discarded when mission traffic rates rise. This is to be expected with a PQ QoS configuration.

With such high packet drop rates, it is extremely important that the network configuration support the campaign phase and declared mission area. Table 7 and Table 8 convey the result of the router queue performance in terms of mission area. Without the command links active, the primary mission area is able to deliver all of its packets. Secondary mission areas begin to suffer however as mission rates rise with ongoing campaign phases. When command links are active, the performance impact is even more evident during all phases. It is also clear that the primary mission area is underserved when the network QoS is not configured for the declared mission area. Once the network QoS configuration is adjusted, performance to the primary and secondary mission areas is returned.

Table 7 - Command links inactive. Mission area packet delivery success during each campaign phase - network configuration step. Performance suffers as mission rates rise with primary mission area performance near SLS threshold with improper QoS configuration.

phase	equal-equal	equal-logs	logs-logs	ops-logs	ops-ops	intel-ops	intel-intel
ops	100%	100%	68%	80%	100%	100%	97%
intel	100%	100%	100%	100%	100%	94%	100%
logs	100%	100%	100%	100%	67%	4%	4%
personnel	100%	100%	10%	3%	3%	0%	0%

Table 8 - Command links active. Mission area packet delivery success during each campaign phase - network configuration step. Very limited bandwidth availability makes QoS configuration absolutely critical for primary mission area performance.

phase	equal-equal	equal-logs	logs-logs	ops-logs	ops-ops	intel-ops	intel-intel
ops	100%	100%	0%	42%	100%	100%	60%
intel	100%	100%	4%	100%	95%	15%	100%
logs	100%	100%	92%	100%	5%	0%	0%
personnel	100%	99%	0%	0%	0%	0%	0%

Figure 5 and Figure 6 demonstrates the ultimate impact to mission area performance in terms of end-to-end delay. Whether or not the command links are active, end-to-end delay and packet delay variation do not meet the needed specification when the QoS arrangement is not configured for the correct campaign phase. Lower priority mission areas experience indefinite delay as mission rates rise and primary mission areas cannot deliver the performance specified for their phase in the SLA.

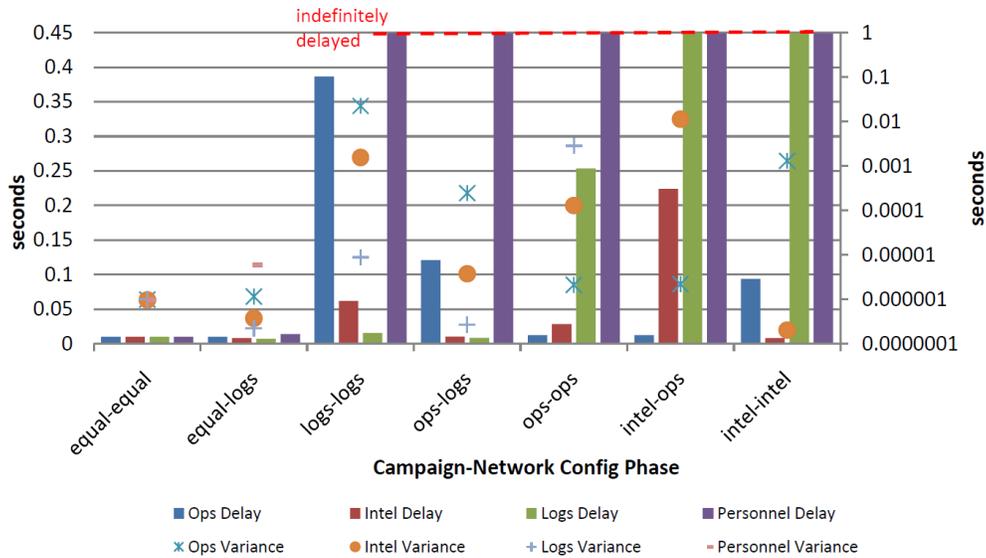


Figure 5 - Command links inactive. Average end-to-end delay and variation for each mission flow in each campaign plan - network config phase. Performance falls below specification in misconfigured phases.

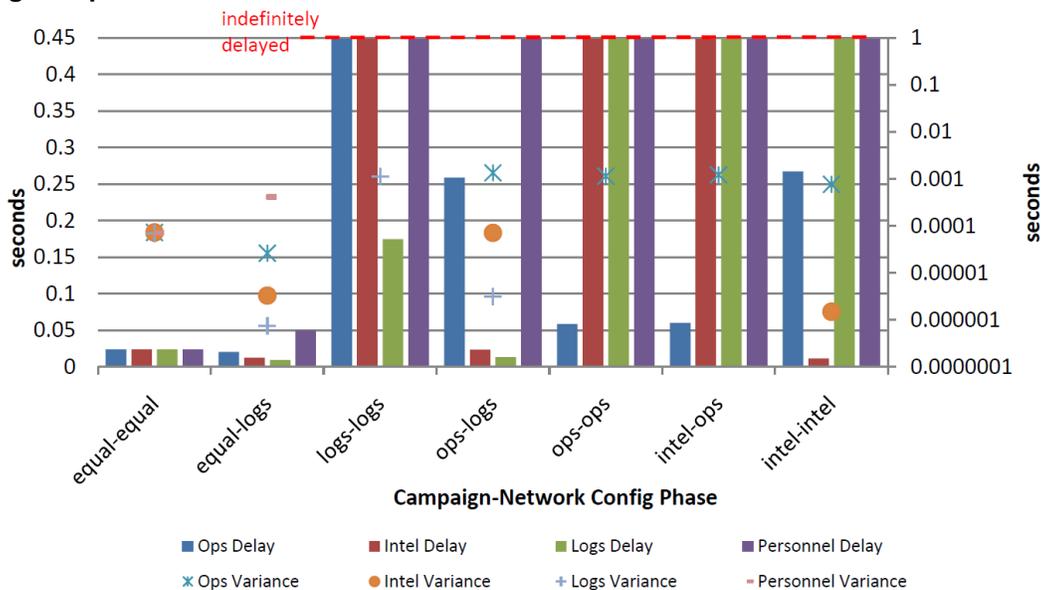


Figure 6 - Command links Active. Cost of network misconfiguration increased with command links - increased average end-to-end delay and variation for each mission flow in each campaign plan - network config phase.

While it is an improvement to find superior performance for the primary mission area by aligning QoS, it is still a concern that other mission areas cannot meet their minimal performance goals. By failing these non-primary areas, the overall campaign phase is likely to fail. Other QoS protocols and an increased focus on minimum performance levels for all mission areas will prevent these failures.

Figure 7 and Figure 8 convey a measure of value based on the number of packets delivered by each mission area. The value statistic for campaign phase – network configuration c is given by

$$V_c = \sum_{m=1}^{M \text{ mission areas}} C \omega^{\alpha_m} N_m \quad (1)$$

where C is a normalization constant, ω^{α_m} is a weighted priority level based on priority α ($1 \leq \alpha \leq 4$ corresponding to the lowest to highest priority classes) and N_m is the number of packets delivered to the destination for mission area m . Command link packets are not included in the value statistic calculation since delivery of those packets is near 100% when the links are active. A similar statistic was previously used to calculate performance of a scheduling routine based on satisfied requests (Theys et al. 2000; Naik, Siegel & Chong 2001; Dharwadkar, Siegel & Chong 2001).

Packets belonging to the declared mission area for a phase are designated priority $\alpha = 4$ and therefore have more value when delivered. When the network QoS is not properly configured, the value delivered by the network is degraded as demonstrated in the “equal-logs” and “ops-logs” phases. When network QoS is adjusted for the campaign phase, the value is increased as demonstrated in the “logs-logs” and “ops-ops” phases.

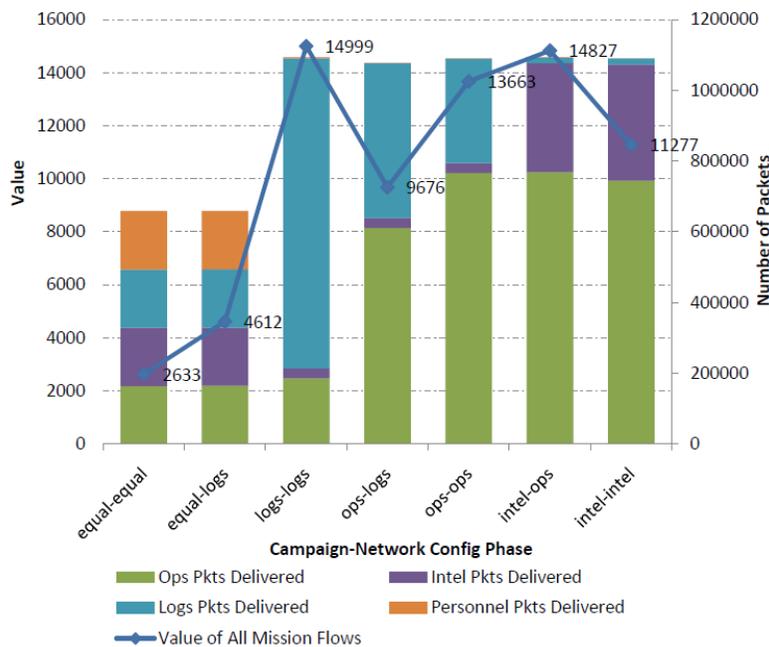


Figure 7 - Command links inactive. Value increases when campaign plan and network configuration are aligned. Bars show volume of a mission area's packets delivered a phase.

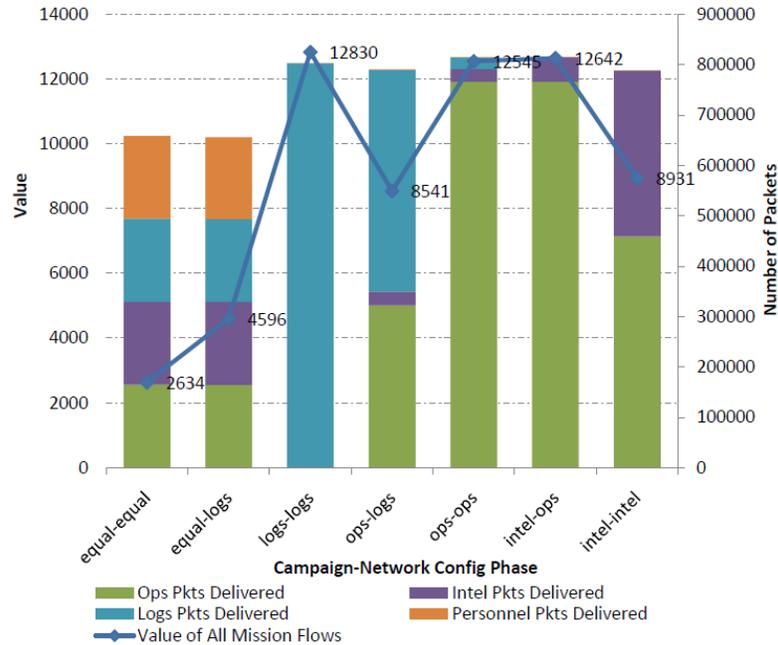


Figure 8 - Command links active. Overall delivered value of non-command link data is less than without command links but aligning campaign phase and network configuration increases value.

It appears the transition from “ops-ops” to “intel-ops” and then to “intel-intel” violates our conclusions about value. It is useful to examine the bars in these figures which show the volume of packets delivered for each mission area. When the “intel-ops” phase begins, the volume of intel mission area packets produced and delivered increases tenfold. Still having priority $\alpha_{intel} = 3$, $V_{intel-ops}$ increases measurably with the increased delivery of intel mission area packets over logs mission area packets. This increase occurs in spite of the mismatched QoS configuration. When the configuration is corrected in the last phase, “intel-intel”, $V_{intel-intel}$ drops below even $V_{ops-ops}$. $V_{ops-ops}$ should exceed $V_{intel-intel}$ because the volume of ops packets in “ops-ops” is greater than those of intel packets in “intel-intel” (the same holds true for $V_{logs-logs}$ vs $V_{ops-ops}$). The reason for the decrease from “intel-ops” to “intel-intel” also lies in the volume of ops packets. When top priority in the QoS configuration shifts to intel, the ops packets have lesser priority and contribute less value to $V_{intel-intel}$. In spite of the increased value of intel packets, the number of delivered ops packets in “intel-intel” are the same or less than in “intel-ops” and this decrease has heavy impact on $V_{intel-intel}$. To prevent this issue, an improved value statistic may vary ω from one mission-area to another based on a mission manager’s determination of importance.

Experiment Summary

Using a simple network model, we demonstrated the impact of implementing a QoS configuration without properly aligning it to the current mission area within the campaign. DSCP and PQ are useful tools for QoS but will heavily restrict lower priority traffic in favor of higher priority. Command links represented highest priority traffic and forced mission area performance down to the specification threshold when QoS configuration was aligned with campaign phase. This demonstrates the criticality of aligning the QoS configuration with missions if useful bandwidth is to be kept in reserve. Finally, important performance metrics such as overall packet delivery, end-to-end delay, packet delay variation and value were significantly affected by QoS actions. QoS configurations which are not aligned with missions

in campaign phases limited the performance of primary mission areas while aligned configurations served to protect performance of primary mission areas even as network activity increased. It is difficult to assure overall mission success, but it is well-understood that mission failure often results from failure of any single aspect of the mission. By prioritizing network performance using mission IERs, not only is the primary mission well-served but all other missions can avoid failure modes as well.

Further Experimental Work

Additional simulations using networks with multiple constrained links and multiple QoS-enabled routers will increase the understanding of mission-oriented QoS. Future work will include additional information flows including TCP and will involve more typical application behavior to increase the fidelity of the model. Additional steps will incorporate DRE-like situations by embedding applications for one mission area within a community focused on a different mission area. Finally, adding detail to the value statistic such as a weight base which varies from mission area to mission area will permit its use in more varied situations.

Summary

Military and civilian contingency operations such as disaster first responders and command and control systems must always wrestle with the challenge of deploying an adequate network which remains tailored even as situations and objectives change. In this paper, we have presented various methods for decomposing missions in order to apply systems engineering principles for network design. We have also explored the QoS concept and some recent research in developing a construct for making QoS mission-aware. We detailed an experiment examining network performance under varying mission objectives and QoS configurations which demonstrated the criticality of aligning QoS to the mission. This research contributes to engineering a modern network which can serve various users belonging to different communities of interest and who are conducting operations on a dynamic system of systems.

References

- U.S. Department of Defense. (2003). Joint Planning Pub. JP 3-30: Command and control for joint air operations.
- . (1995). CJCSI 3500.01B Joint Training Policy for the Armed Forces of the United States
- . (2002). CJCSM 3500.04C Universal Joint Task List (UJTL) v4.2
- . (2004). Defense Acquisition Guidebook. Office of the Deputy Under Secretary of Defense for Acquisition and Technology ODUSD(A&T), Based on DoD Instruction 5000.02.
- . (2005). DoD Directive 3020.40 "Defense Critical Infrastructure Program (DCIP)".
- . (2006). Joint Planning Pub. JP 6-0: Joint communications system.
- . (2007). Chairman of the JCS Instruction 6250.01C: Satellite Communications. 30Apr07.
- . (2008). Joint Planning Pub. JP 3-0: Joint operations.
- . (2009). DoD Architecture Framework v2.0.
- 12th Air Force. (1996). AFFOR Air Operations Center Standard Operating Procedures.
- Abdelwahed, S., Neema, S., Loyall, J., & Shapiro, R. (2003). A hybrid control design for QoS management. Proceedings of the 24th IEEE International Real-Time Systems Symposium, 366.
- Alberts, D. S., Garstka, J., & Stein, F. P. (1999). *Network centric warfare : Developing and leveraging information superiority*. Washington, DC: National Defense University Press. Retrieved from http://www.dodccrp.org/files/Alberts_NCW.pdf
- Anunciado, D. H. (2006). Framework for the dynamic establishing of command and control (C2) nodes in an enterprise environment. Paper presented at the *2006 Command and Control Research and Technology Symposium (CCRTS): The State of the Art and the State of the Practice*, San Diego, CA.
- Armstrong, S. (2005). *Mission task analysis for the NATO defence requirements review*. No. RTO-TR-SAS-027 AC/323(SAS-027)TP/39
- Brown, M. J. (2001). Re-characterizing goals, objectives and metrics to support application of advanced analytical technologies in military command and control. *6th ICCRTS: Collaboration in the Information Age*, Annapolis, MD. , 6(6)
- Bryan, K., DiPippo, L. C., Fay-Wolfe, V., Murphy, M., Zhang, J., Niehaus, D., et al. (2005). Integrated CORBA scheduling and resource management for distributed real-time embedded

systems. Paper presented at the Real Time and Embedded Technology and Applications Symposium, 2005. RTAS 2005. 11th IEEE, 375-384.

Carson, R. S., Kennedy, P. D., Butler, M. C., Crowe, T., Erikson, P. C., Miller, A. P., et al. (2005). Effects of functional allocations and associated communications on C2ISR mission effectiveness. *10th ICCRTS: The Future of Command and Control*, McLean, VA. , 10 264-268.

Cho, H., Ravindran, B., Wu, H. and Jensen, E.D. On multiprocessor utility accrual real-time scheduling with statistical timing assurances. in *Proc. of the IFIP International Conference on Embedded and Ubiquitous Computing*, August 2006, pp. 274–286.

Dasarathy, B., Gadgil, S., Vaidyanathan, R., Parmeswaran, K., Coan, B., Conarty, M., et al. (2005). Network QoS assurance in a multi-layer adaptive resource management scheme for mission-critical applications using the CORBA middleware framework. Paper presented at the *Real Time and Embedded Technology and Applications Symposium, 2005. RTAS 2005. 11th IEEE*, 246-255.

Dharwadkar, P., Siegel, H. J., & Chong, E. K. P. (2001). A heuristic for dynamic bandwidth allocation with preemption and degradation for prioritized requests. *Distributed Computing Systems, International Conference on*, 0, 0547. Retrieved from <http://doi.ieeecomputersociety.org/10.1109/ICDSC.2001.918985>

Doshi, B., Kim, P., Liebowitz, B., Park, K. I., & Wang, S. (2006). *Service level agreements and QoS delivery in mission oriented networks* (White Paper No. 06-0084). Boston, MA: MITRE. Retrieved from http://www.mitre.org/work/tech_papers/tech_papers_06/06_0084/

Estrin, D., Govindan, R., Heidemann, J., & Kumar, S. (1999). Next century challenges: Scalable coordination in sensor networks. Paper presented at the *MobiCom '99: Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking*, Seattle, Washington, United States. 263-270. Retrieved from <http://doi.acm.org/10.1145/313451.313556>

Firoiu, V., Boudec, J. L., Towsley, D., & Zhang, Z. (2002). Theories and models for internet quality of service. *Proceedings of IEEE, Special Issue on Internet Technology, August*.

Ghosh, S., Rajkumar, R., Hansen, J., & Lehoczky, J. (2005). Scalable QoS-based resource allocation in hierarchical networked environment. Paper presented at the *Real Time and Embedded Technology and Applications Symposium, 2005. RTAS 2005. 11th IEEE*, 256-267.

Graham, S. R. (2008). Challenges to the construction of an evolvable mobile military network. Air University Air Force Institute of Technology Blue Dart.

Intanagonwiwat, C., Estrin, D., Govindan, R., & Heidemann, J. (2002). Impact of network density on data aggregation in wireless sensor networks. Paper presented at the *ICDCS '02: Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS'02)*, 457.

- Intanagonwiwat, C., Govindan, R., Estrin, D., Heidemann, J., & Silva, F. (2003). Directed diffusion for wireless sensor networking. *IEEE/ACM Trans.Netw.*, 11(1), 2-16. Retrieved from <http://dx.doi.org/10.1109/TNET.2002.808417>
- Jiang, Y., Chen-Khong, Y. J., Treurniet, J. R., & Simons, N. R. (2000). Providing quality of service monitoring: Challenges and approaches. *International Journal of Network Management*, 10(6), 323. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.21.7974&rep=rep1&type=pdf>
- Kurose, J. (1993). Open issues and challenges in providing quality of service guarantees in high-speed networks. *SIGCOMM Comput. Commun.Rev.*, 23(1), 6-15. Retrieved from <http://doi.acm.org/10.1145/173942.173943>
- Kurose, J. F., & Ross, K. W. (2007). In Hirsch M., Hartman S. (Eds.), *Computer networking: A top down approach* (4th ed.) Addison Wesley.
- Lawrence, T., Hurley, P., Wheeler, T., Kanevsky, A., Maurer, J., Wallace, P., et al. (1997). Quality of service for AWACS tracking. Rome, NY: AFRL/IFGA.
- Liebeherr, J., Patek, S. D., & Burchard, A. (2003). Statistical per-flow service bounds in a network with aggregate provisioning. In *Proceedings of IEEE Infocom 2003*, 1680-1690.
- Loyall, J., Schantz, R., Corman, D., Paunicka, J., & Fernandez, S. (2005). A distributed real-time embedded application for surveillance, detection, and tracking of time critical targets. Paper presented at the *Real Time and Embedded Technology and Applications Symposium, 2005. RTAS 2005. 11th IEEE*, 88-97.
- Loyall, J., Sharma, P., Gillen, M., & Schantz, R. (2007). A QoS management system for dynamically interoperating net-centric systems. *Proceedings of the {SPIE} conference on defense transformation and net-centric systems*, Orlando, FL april 9-12.
- Loyall, J., & Schantz, R. (2008). Dynamic QoS management in distributed real-time embedded systems. In I. Lee, J. Leung & S. Son (Eds.), *Handbook of Real-Time and Embedded Systems*, CRC Press.
- Loyall, J., Carvalho, M., Martignoni, A. I., Schmidt, D., Sinclair, A., Gillen, M., et al. (2009). QoS enabled dissemination of managed information objects in a publish-subscribe-query information broker. In *Proceedings of the {SPIE} Conference on Defense Transformation and Net-Centric Systems, Orlando, FL April 13-17*,
- Mirhakkak, M., Schult, N., & Thomson, D. (2000). Dynamic QoS for mobile ad-hoc networks. MITRE Technical Report.
- Mitchell, G., Loyall, J., Webb, J., Gillen, M., Gronosky, A., Atighetchi, M., et al. (2008). A software architecture for federating information spaces for coalition operations. Paper presented at the *Proceedings MILCOM 2008*, San Diego, CA.

- Naik, A., Siegel, H. J., & Chong, E. K. P. (2001). Dynamic resource allocation for classes of prioritized session and data requests in preemptive heterogeneous networks. In *Proceedings of the 2001 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2001)*, Las Vegas, Nevada. , II(June 2001) 787-796.
- OPNET (2003). Modeling variability in background traffic. Product Documentation 2003.
- Ravichandar, R., Arthur, J. D., & Broadwater, R. P. (2006). Reconciling synthesis and decomposition: A composite approach to capability identification. *14th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS)*, abs/cs/0611072
- Richards, D. L. (2009, 28 Jul 2009). BACN bleeds. Retrieved from <http://www.gcic.af.mil/news/story.asp?id=123160760>
- Sage, A. P. (1992). Systems engineering. (pp. 86-92) Wiley.
- Sisalem, D. (1997). End-to-end quality of service control using adaptive applications. Paper presented at the *IFIP Fifth International Workshop on Quality of Service (IWQOS '97)*.
- Stanley, J.E., Mills, R.F., Raines, R.A. and Baldwin, R.O. (2005), Correlating Network Services with Operational Mission Impact, Military Communications Conference (MILCOM 2005), Atlantic City, New Jersey, Oct 2005, pp. 1-7.
- Theys, M. D., Beck, N. B., Siegel, H. J., & Jurczyk, M. (2000). Evaluation of expanded heuristics in a heterogeneous distributed data staging network. In *Proceedings of the 9th Heterogeneous Computing Workshop, 2000. (HCW 2000)*, 75-89.
- Wixson, J. R. (1999). Function analysis and decomposition using function analysis systems technique. *International Council on Systems Engineering Annual Conference (INCOSE '99)*, Brighton, England
- Wong-Jiru, A., Colombi, J, Suzuki, L. and Mills, R. (2007) Graph theoretical analysis of network centric operations using multi-layer models. In *Proceedings of the Conference on Systems Engineering Research, 2007 (CSER 2007)*.
- Ye, Nong, Ying-Cheng Lai, and Toni Farley. (2003). Dependable information infrastructures as complex adaptive systems. *Journal of Systems Engineering* 6, (4): 225-37.
- Yurcik, W., & Banerjee, S. (1995). "Local quality of service provisioning to meet end-to-end requirements." In *ATM Networks, Proc. of the SPIE Conference on Emerging High Speed LANs and WANs*.
- Zhang, L., Deering, S., Estrin, D., & Shenker, S. (1993). "RSVP: A new resource reservation protocol." *IEEE Networks Magazine*, 8-18.

Biographies

Vinod D. Naga is a Systems Engineering PhD candidate at the Air Force Institute of Technology. His research focuses on improving quality of service for DoD networks. He previously served as a technical liaison officer to a non-DoD agency, a field operations officer, a C2ISR technology transition manager and a program manager for E-8C JSTARS moving target indicator radar technology at the Air Force Research Laboratory.

Dr. John Colombi is an Assistant Professor of Systems Engineering at the Air Force Institute of Technology. He leads sponsored research in architectural analysis, quality of service and Human Systems Integration. Before joining the faculty, Dr. Colombi led command and control systems integration efforts, systems engineering for the E-3 AWACS aircraft, served at the National Security Agency developing biometrics and information security and ran communications networking research at Air Force Research Laboratory.

Dr. Michael R. Grimaila is an Associate Professor of Systems and Engineering Management at the Air Force Institute of Technology. His research interests include mission assurance, cyber damage assessment, embedded system security, and risk mitigation. Dr. Grimaila serves as an Editorial Board member of the Information System Security Association (ISSA) Journal and consults for a number of DoD organizations.

Dr. Kenneth Hopkinson is an Assistant Professor of Computer Science at the Air Force Institute of Technology. His research interests focus on simulation, networking, and distributed systems. More specifically, he looks at fault-tolerant and reliable approaches to mobile wireless networks as well as security and reliability in critical infrastructures with a concentration on communication-based protection and control systems in the electric power grid.