DYNAMIC END-TO-END QOS MANAGEMENT FOR ADVANCED RF TELEMETRY NETWORKS

Andrzej Cichocki¹, Mariusz A. Fecko¹, John Unger¹, Sunil Samtani¹, Larry Wong¹, Aleksander Kolarov¹, Mark Radke², Tom Young³

¹Applied Research, Telcordia Technologies
Piscataway, NJ

²Tybrin Corporation
Edwards AFB, CA

³USAF AFMC
Edwards AFB, CA

ABSTRACT[†]

We present iMANPOL – a dynamic end-to-end QoS management system for advanced RF telemetry networks with the red-black separation constraints. iMANPOL system encompasses network resource monitoring, allocation, and enforcement techniques to increase throughput and reduce end-to-end delay of telemetry traffic while protecting priority mission-critical flows. These goals are achieved through adaptive techniques for providing Differentiated Services, Admission Control Function, and Flow Preemption. The iMANPOL system has been implemented and tested in an emulated environment. The test results confirm that the admission control, particularly when coupled with preemption, can significantly increase the performance of priority flows in congested networks. An iMANPOL deployment in the integrated enhanced network telemetry would make more network resources available for high-priority tests and enable more dynamic test scheduling.

I. INTRODUCTION

Future network-based wireless telemetry solutions like Integrated Network Enhanced Telemetry (iNET) will enable networked access of telemetry data from multiple ranges simultaneously during a live flight test. Such a networked capability, however, poses a challenge because current telemetry networks require manual planning and configuration of the network and test articles. This process is time-consuming, error prone, and requires significant manpower with relatively advanced skills. It also increases the network downtime and does not allow dynamic re-configuration of the network when the test plans or requirements change. Automated tools are required to perform these tasks, and in particular, to manage resource sharing appropriately through defined Quality of Service (QoS) rules.

In this paper we present iMANPOL (iNET Management and Operations with Policy Controls) – a dynamic end-to-end QoS management system for advanced RF telemetry networks with the red-black separation constraints. iMANPOL system encompasses resource monitoring,

[†] This project is funded by the Test Resource Management Center (TRMC) Test and Evaluation/Science & Technology (T&E/S&T) Program through the U.S. Army Program Executive Office for Simulation, Training, and Instrumentation (PEO STRI) under Contract No. W900KK-09-C-0021. The Executing Agent and Program Manager work out of the AFFTC.

allocation, and enforcement techniques to promote the utility (in terms of QoS-related metrics such as flow throughput and delay) of the network while protecting priority mission-critical traffic. The iMANPOL approach uses two techniques: (1) Differentiated Services to protect mission-critical sessions by giving them preferential treatment inside the network, and (2) admission control and preemption to ensure that only traffic that is authorized and can be handled by available resources enters the network. The decision logic of iMANPOL is supported by the Network Resource and Performance Estimation (NRPE) module that learns the state of the black network (e.g., congestion/utilization) purely by applying end-to-end low-overhead loss and bandwidth probes.

Our design has several unique features. First, we support a wide range of actions such as flow admission, rejection, and preemption, but also the ability to temporarily defer delay tolerant flows. Second, the integration with the resource estimation module prevents premature reaction to transient network events such as link outages. Third, the decision logic is based on a set of rules that apply strict priority-based admission and preemption. This last feature offers additional protection to high priority traffic by freeing up resources when necessary.

The presented system's components (admission control logic, preemption, and network estimation) have been implemented and integrated on a Linux platform. We also performed a series of performance evaluation experiments using a testbed emulating a wireless telemetry network. The results of these tests confirm that the admission control and flow preemption are critical for the performance of priority flows in congested networks.

The remainder of this paper is organized as follows. Section II discusses telemetry networks and the problems posed by the RF environment. In Section III, the iMANPOL system is described, together with NRPE module (Section IV) and the admission control and preemption algorithms (Section V). Finally, Section VI presents the results of iMANPOL's performance tests.

II. TELEMETRY NETWORKS

Current telemetry systems gather measurements from the test article sensors and systems and broadcast them to the ground over an RF channel. This method of telemetering is referred to as serial streaming telemetry (SST). Once the data has been received on the ground, it is forwarded to a mission control room (MCR) usually over some manner of network transport. In the past this was predominantly ATM, but over the last few years the migration to IP as the ground network transport has been gaining traction. Generally a static QoS configuration of the ground network is sufficient to maintain the QoS requirements for telemetry streams and other services on a wired IP network.

A current Central Test and Evaluation Investment Program (CTEIP) project, Integrated Network Enhanced Telemetry (iNET), is working on redefining how telemetering is accomplished. iNET is moving telemetry from a broadcast, non-interactive paradigm to a network based, interactive one. As programs like iNET mature, they will provide enhancements to the current capabilities by providing a network uplink/downlink to the test article. It is envisioned that initially the network link to the test article will be mainly used for command and control, for status of the instrumentation package, and for the retransmission of lost telemetry data. As iNET matures and

evolves, it will eventually support telemetry data that originates as packets on the test article, routed over a multi-hop RF network, a ground network, and on to the MCR all over IP. As telemetry networks become more commonplace, they will experience the same QoS challenges that networking over RF brings with it.

In a wired IP network, typically link capacities do not change, and therefore a well engineered QoS configuration can remain static and still satisfy the intended QoS requirements. Unlike wired network links, wireless network link characteristics *can* change over time. Variations in RF channels capacity can be affected by many factors, including (1) changes in time slot allocations in network that utilize a Time Division Multiple Access (TDMA) for Media Access Control (MAC) which will impact available channel capacity as well as changing the timing of when packets are transmitted; (2) fluctuations in link capacities caused by changes in modulation used for a particular channel condition; or (3) fluctuations in channel capacity caused by changes in how the nodes access the channel.

Static QoS configuration that is not optimized for the current state of a wireless link may cause increased problems with the quality of the packet stream, such as increased jitter, increased latency and dropped packets. High priority traffic for applications like safety of flight data, streaming audio, video and interactive applications can be significantly impacted by these effects. When guaranteed delivery type applications are subjected to these effects, throughput and link efficiency are negatively impacted as well. The possible solution to this problem lies in dynamic modification of the QoS configurations based on varying link conditions to ensure that priority packet flows and link efficiency are not adversely affected by these changes. In order to do so, a solution needs to ascertain the instantaneous condition of the link and traffic flows, and dynamically make appropriate changes to the QoS parameters of the network.

In a telemetry network like iNET where the use of network encryptors protects sensitive telemetry data, complications due to security boundaries separating different portions of the network make implementing a dynamic QoS solution more difficult. Devices on either side of a security boundary have little or no insight into conditions on the other side (which is usually dictated by equipment and/or security policy). Thus, methods are needed to determine or estimate the current state of the network on an end-to-end basis when network encryptors are used. The iMANPOL system has been developed to answer the telemetry networks' needs for dynamic management of QoS parameters in response to varying link/network conditions, even where the transmission of some/all network state and traffic flow information is prohibited by security constraints.

III. IMANPOL

A high-level deployment architecture of iMANPOL is shown in Figure 1. The figure shows a notional diagram of an advanced telemetry network, consisting of red (colored) and black (colorless) networks. DiffServ is provided on the black side by the *QoS Control Agent (QCA)* co-located with every router. The agent supports dynamic Multi Level Precedence and Preemption (MLPP) policies via Active Queue Management (AQM) techniques. It monitors all the traffic originating from the local red enclaves as well as traffic traversing through the router. It uses the traffic statistics and the link bandwidth to determine the congestion. If congestion is

detected, the QCA configures the Weighted Fair Queuing (WFQ) and Random Early Drop (RED). These techniques actively manage the bandwidth and the drop profiles assigned to each traffic class/queue to protect high precedence traffic.

QCA works in concert with Admission Controllers, which perform *Edge-Based* (“Red Side”) *Admission Control* (EBAC) on a Test Article and Ground Station. The iMANPOL design does not have a centralized controller, i.e., each red enclave has its own instance of an admission controller that makes independent decisions. Nodes that need to send high priority traffic into the network make requests to their local Admission Control Function (ACF). The ACF determines whether the new request is allowed based on security policies and whether it can be supported by the network based on currently admitted flows and network resource usage. If the new request cannot be supported, either it is rejected (or possibly renegotiated), or other traffic of lower priority is preempted or downgraded based on current mission/test needs. This approach ensures that available network resources are allocated to the traffic that can best support the mission/test needs.

Figure 1: High level QoS deployment architecture

Consider the following relevant scenarios in advanced telemetry networks:

Request for bandwidth: During a test flight inter-maneuver period, a Test Engineer retrieves information of interest from the onboard SSR. Information is retrievable via a telemetry network, controlled by QoS parameters defined by or for the requestor. Another example is a need to transmit key test parameters to the cockpit display at a higher rate during stressing test sequences or to perform one-time data retrieval.

Reallocation of bandwidth: One of the TAs experiences an anomaly and needs to send additional data over a telemetry network to analyze the anomaly; hence, the bandwidth dedicated to other TAs must be temporarily reduced.

The iMANPOL would perform the following actions as part of the end-to-end QoS management:

- The control software on a Test Article makes QoS requests to the ACF for each QoS-sensitive session (e.g., to download telemetry data to the Ground Station), by specifying flow source and destination (addresses and ports), a set of requested traffic rates, and optional flow duration. The ACF is invoked to perform admission control (it admits or rejects the request depending on available network resources and request priority).
- If the request is admitted into the network, the session packets are marked with the appropriate priority and class (DSCP) and sent to the local router within the RF Network Element (RFNE). The RFNE forwards the packet to the local INE/HAIPE device, which would tunnel the packet in an encrypted tunnel through a multi-hop RF network to the destination INE device, e.g., at the RFNE in the Ground Station. Although the packet is encrypted, the DSCP markings are copied to the outer IP header and the DSCP value is visible on the colorless side. The new data flow then starts traversing the RF network.
- QCAs along the RF network path collect the traffic statistics for each DSCP to determine the congestion, as well as traffic loading across different classes. The traffic statistics is gathered for traffic originating in the local enclaves as well as those originating in remote enclave and traversing the router. QCAs will reconfigure the local router if needed to protect higher precedence traffic along congested links.
- If data rate needs to be increased for an existing flow, request for more bandwidth is sent to ACF, which makes admit/reject decision. The amount of bandwidth provided to individual TAs and their specific data flows are readjusted via ACF/QCA.
- ACF preemption module reduces the rate of (or terminates) lower priority flows if network resources do not allow accommodating high priority request. Additionally, QCA reduces queue weights for these lower priority flows. Once the need for additional bandwidth goes away, QCA increases queue weights to their previous state.

IV. NETWORK RESOURCE AND PERFORMANCE ESTIMATION (NRPE)

The COMSEC barrier prohibits direct communication between elements situated on the protected (i.e., red) network elements where tactical applications reside and the (black) radio network where the throughput-limiting loss-prone wireless links reside. From red-to-black, only the first 6 bits (the DSCP value) of the ToS field as plain text – Everything else in the original red IP datagram is encrypted while transported through the black radio network. From black-to-red, no information, not even marking by a router in the radio network of the ECN bit in the ToS field is communicated when the datagram is de-encrypted by COMSEC device at the destination platform. Thus, on the red side, the EBAC system must infer radio network resources and performance. This capability is provided by network resource and performance estimation (NRPE) module that learns the state of the black network (e.g., congestion/utilization) purely from end-to-end statistics by applying low-overhead loss and bandwidth probes.

Available bandwidth estimation: Existing techniques [1] for estimating available bandwidth are either biased and inaccurate or intrusively overload the end-to-end path with steadily increasing traffic to create congestion [2]. Our approach is based on extending the packet

dispersion technique [1] by using not simply the spacing between pairs of packets, but the sum of the spacing's between pairs of packets in a longer probe sequence. As described in Ref.[3], this approach reduces the estimation noise, and while the results of the probe do not in general converge to the available bandwidth, they do converge to a useful lower bound on available bandwidth. While packet dispersion techniques typically consider the spacing between pairs of packets, our technique uses a “train” of packets injected into the network periodically from the source node to the destination node. This also enables a bursty loss estimate to be formed. Since the capacity estimate is dependent on relative time, no time synchronization is necessary between the source and destination.

Figure 2: Available bandwidth estimation

Performance results: We have performed extensive testing of the NRPE capability using both emulated and real-life wireless networks. In an emulated network, we estimated the available bandwidth varying between 2380 kbps and 180 kbps. Different probe settings gave sufficiently accurate estimates for both the lower (150 kbps) and the higher (2200 kbps) bound. The probe overheads were 1.1% and 0.84%, respectively. In addition to the emulated scenarios, the bandwidth probe has also been tested outdoor (Ft. Monmouth, NJ, 2009) using an experimental radio network based on Handheld, Manpack & Small Form Fit (HMS) radios in the CSMA mode. The experimental HMS radios had two modes of operation – high data rate at 120 kbps and low data rate at 70 kbps. We observed that the bandwidth probe is able to accurately assess the transmission rate of

the radio (Figure 2). Secondly, the bandwidth probe gives a reasonably accurate lower bound on available bandwidth over a range of load up to 50% of the system capacity.

V. ADMISSION CONTROL FUNCTION (ACF) AND ACF ALGORITHMS

As shown in Figure 1, flow admission control represents the primary mechanism by which EBAC can promote QoS for network traffic transported via encrypted tactical wireless networks. The basic idea of an Admission Control Function (ACF) is to selectively admit, deny or preempt network application flows so as to promote the overall utility of the network of the traffic while protecting the end-to-end performance of priority flows. Among the challenges of EBAC in the telemetry environment is to compute good flow admission decisions given the limited insight into network resources limited by wireless link capacity.

Our design has several unique features. First, we support a wide range of actions such as flow admission, rejection, and preemption, but also the ability to temporarily defer delay tolerant flows. Second, the integration with the resource estimation module allows prevents premature reaction to transient network events such as link outages. Third, the decision logic is based on a set of rules that apply strict priority-based admission and preemption. This last feature offers additional protection to high priority traffic by freeing up resources when necessary. A fundamental principle enforced by the ACF is flow admission/preemption based on flow-priority. That is, flows deemed to be high-priority are likely admitted while flows deemed to be low-priority are admitted only if its network path is deemed to be lightly utilized. Each flow, therefore, is associated with one of several possible flow-priority levels/settings. The set of

flow-priority settings are strictly ordered in terms of their priority/precedence. The precedence values in the current implementation are (from the highest): Flash Override (1), Flash (2), Immediate (3), Priority (4), Routine (5) and Best Effort (6). The ACF is reconfigurable to consider other priority schemes.

Application proxies: To address application-specific admission control needs and to provide a unified admission request descriptor, the EBAC system also incorporates application proxy functionalities. Application proxies are entities in the system that can make admission requests to the EBAC system. An application proxy can be as simple as a shell script that makes the request to the EBAC system and only starts the application if the request was accepted. When the application terminates, the shell script would send a termination notification to the EBAC system to inform it that the flow has terminated. The EBAC system would then clean up the resources associated with the flow.

Triggering events: The ACF functionality is associated with a number of specific, discrete triggering events due to the interactions with network applications and the NRPE functionality. Table 1 below identifies some of the key events/scenarios that can trigger admission decision processing. The events are anticipated to occur on relatively coarse-grained time scales (i.e., on the order of several seconds, or more). This time scale is reasonable given the binary nature of admission control decisions and the fact that a larger frequency of triggering events might be disruptive to the network applications.

Table 1 Admission Control (ACDL) Decision Logic Triggers

Trigger	Description	Triggered Action
New flow request	A request from a local network application to send or receive a network flow	Accept/Deny/Preempt decision logic for the flow request
Flow termination	Notification (either explicit or implicit) from a local flow source or receiver that the communication session has ended	Check for whether recently denied flow requests or preempted flows should be allowed or promoted
Resource increase	NRPE component detects an improvement in network resource availability	Check for whether recently denied flow requests or preempted flows should be allowed or promoted
Resource decrease	NRPE component detects a degradation in network resource availability	Check for whether an accepted flow should be preempted

Admission control algorithms: The first algorithm implemented in iMANPOL is deterministic. The decision it makes is based on priority of the flow, requested bandwidth and expected duration. It always admits flows having one of two highest priorities (Flash Overdrive and Flash), and flows that have Immediate priority and are short. Lower priority flows are accepted only if the available bandwidth reported by the NRPE is sufficient to accommodate them. However, this algorithm exhibits a few undesirable properties: some flows may never be admitted, and, when coupled with preemption method, it may lead to oscillations (i.e., admitting and preempting flows in a cycle). To alleviate these problems, we have developed a probabilistic admission algorithm. It also accepts high priority flows and short flows without further consideration. However, for the lower priority flows it computes an admission probability based on the flow priority, network capacity slack (that is, the difference between available bandwidth reported by the NRPE and bandwidth requested by the flow) and the

reported packet loss rate. This probability is used to make a decision about admission of the flow. Thus, the same flow in the same network conditions may or may not be admitted, which lessens the likelihood of oscillations and makes admission of lower priority flow more fair (for example, less dependent on the order of requests).

Preemption: The other fundamental scenario the ACF addresses is a change in network resources, in particular, a worsening of the network state. The deterioration of network performance and resources may be due to, for example, a link failure event that results in severe loading on links required for traffic reroute, an unplanned influx of network application traffic that overloads otherwise correctly provisioned link capacities, or reduced capacity caused by channel degradation. In such cases, to protect mission-critical traffic, the ACDL selectively preempts flows to help alleviate the network congestion. The preemption algorithm is executed in a number of stages separated by a random delay. In each stage the ACF is limiting the bandwidth and possibly terminating flows with higher priority than in the previous stage. That is, if the NRPE reports high congestion, first the Best Effort flows will be preempted, then the Routine etc. However, high priority flows are always protected, and will not be preempted regardless of the network situation. While the concept of incrementally preempting flows starting from the lowest-priority flow and delaying subsequent preemption of higher priority flows helps prevent oscillations, they remain a possibility. The random delay in checking the network utilization and reacting is a further step to prevent the oscillations, as is the introduction of the probabilistic admission algorithm. Note also that strict priority-based preemption rules ensure that offered load and throughput for high-priority traffic is steady and relatively well protected.

VI. EXPERIMENTS

To test the behavior of iMANPOL system, and algorithms, we have performed a series of experiments. The testbed consisted of six networked machines configured as shown in Figure 3.

Figure 3 Topology of the testbed

The links represent the RF black side of the network. Conceptually, Nodes 1, 2, 3 and 4 are Test Articles and contain a red-side vehicle network, where the applications and the ACF functionality operate. Nodes 5 and 6 contain only a black side router (inaccessible from the red side of Nodes 1, 2, 3 and 4). The link between these two nodes is made the bottleneck by artificially limiting its rate to 5Mbps. All test flows were running between the “red” nodes on the edge

(Node 1 and 4 pair). The NRPE was installed and run on the same two nodes to provide end-to-end network situation estimation. In our testbed, we do not have an INE device to provide red-black separation; however, we emulate this constraint by not revealing any link statistics to the ACF and NRPE modules. The test cases were a mix of flows with various priorities, bandwidth requirements and durations, as shown in Figure 4. The tests were designed to emulate diverse network conditions that could arise randomly, or as a result of a specific workload. All flows were UDP, implemented by running ‘iperf’ utility. Since the bottleneck link is limited to 5 Mbps, there is not enough bandwidth to accommodate all of them, and the ACF will actually have to perform well to make the network usable.

The experiments were run in five different ACF configurations:

1. no admission control;
2. deterministic admission control;
3. deterministic admission control coupled with preemption;
4. probabilistic admission control;
5. probabilistic admission control coupled with preemption.

Test case	No. concurrent flows	Characteristic of test flows	Aggregated requested bandwidth
T1	10	8 low priority flows, 2 high priority, all the same duration (200 sec) and bandwidth requirements (2Mbps)	20,000
T2	10	8 high priority flows, 2 low priority, with random duration (100/200 sec) and bandwidth (100kbps/2Mbps)	10,500
T3	10	mix of flows with randomly generated priorities, durations and bandwidth	10,500
T4	10	started sequentially with increasing priority (from Best Effort to Flash Override), random duration and bandwidth	10,500
T5	10	started sequentially with decreasing priority (from Flash Override to Best Effort), random duration and bandwidth	12,400

Figure 4 Test flows – Requested bandwidth shown in kbps

In order to compare the performance of the admission control and preemption algorithms, we defined a performance metric based on the ratio of the average of measured bandwidth to the requested bandwidth of the flows. Since the ACF (and in general, iMANPOL system) has been designed to enhance the performance of the high priority flows, the average is weighted by the priority of the flow. Specifically, for each flow f with priority p ($p_{min}=6$, $p_{max}=1$) the requested bandwidth B_f and measured bandwidth b_f , the Priority Performance PP_f metric is defined as

$$PP_f = \begin{cases} 0 & \text{if } \frac{b_f}{B_f} < \text{threshold} \\ \left(\frac{p_{min}}{p}\right)^2 & \text{otherwise} \end{cases}$$

where threshold is set to 80%, as studies [4] indicate this to be the lowest acceptable value. The highest priority flow receiving full requested *bandwidth* will thus have the PP value of $p_{min}^2 = 36$, while the same flow receiving only half of the desired bandwidth will have the PP value of 0. The results of our tests are presented in Table 2.

Table 2: Test results – Priority Performance metric

Test case	No ACF	Deterministic ACF	Deterministic ACF plus preemption	Probabilistic ACF	Probabilistic ACF plus preemption
T1	0	0	36	0	45
T2	90	100	81	99	90
T3	0	1	58	56.25	58
T4	1	37	49	1	49
T5	0	45	49	45	45
Average	18.2	36.6	54.6	40.25	57.4

As expected, the ACF significantly improves the Priority Performance metric. Even the basic deterministic ACF yields **101%** PP improvement. A probabilistic ACF coupled with preemption achieves the best results, yielding **215%** and **57% PP** improvement over the no-ACF case and the deterministic ACF, respectively. Note that for test cases T1, T3, and T5 no flow achieved acceptable performance without admission control, and that in case T1 (two high priority flows mixed with 8 heavy, but low priority ones) only preemption allows us to ensure that the high priority flows actually receive the minimum acceptable bandwidth. Note also a significant difference between performance of the system without admission control, with admission control, and with admission control coupled with preemption. However, the particular algorithm used (deterministic or probabilistic) is of lesser importance with respect to the defined metric.

VII. CONCLUSION

The iMANPOL system provides a dynamic end-to-end QoS management system for advanced RF telemetry networks with the red-black separation constraints. It employs adaptive techniques for providing Differentiated Services, Admission Control Function, and Flow Preemption. The iMANPOL system has been implemented and tested in an emulated environment. The results confirm that the admission control, particularly when coupled with preemption, can significantly increase the performance of priority flows in congested networks. In the near term, we will perform more performance tests that will (1) incorporate dynamic changes of the bottleneck link capacity, (2) assess the impact of wireless losses, (3) defer delay tolerant applications that will avoid rejecting such applications up front while possibly admitting them later when network resources improve, and (4) quantify the increased stability of the system as provided by the probabilistic admission decision logic.

The iMANPOL capabilities offer important benefits for integrated enhanced network telemetry. More network resources could be made available for high-priority tests, which would increase the volume of critical test data collected and processed in real time. The dynamic adaptation of network QoS configuration when the test plans change and to multiple on-going tests as needed would enable a more dynamic test scheduling. This has the potential to increase the number of tests completed, faster scheduling and completion of tests; and more tests to be run per unit time.[‡]

VIII. REFERENCES

- [1] C. Dovrolis, P. Ramanathan and D. Moore, "What do packet dispersion techniques measure?" Proc. IEEE Infocom 2001, April 2001, pp. 905-914.
- [2] M. Jain and C. Dovrolis. End-to-end Available Bandwidth: Measurement Methodology, Dynamics and Relation with TCP Throughput. In IEEE/ACM TON, August 2003
- [3] D. Shur, and A. Zelezniak, "Bandwidth and Performance Monitoring in Virtual Networks using Active Probes", Invited talk, IEEE NJ Coast Section Seminar, December 11, 2002.
- [4] K.C. Mansfield, J.L. Antonakos. Computer Networking from LANs to WANs: Hardware, Software, and Security. Boston: Course Technology, Cengage Learning. P501.

[‡] Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Test Resource Management Center (TRMC) and Evaluation/Science & Technology (T&E/S&T) Program and/or the U.S. Army Program Executive Office for Simulation, Training, & Instrumentation (PEO STRI).