REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE		3. DATES COVERED (From - To)
MAY 2008	Conference Paper Postprint		January 2008 – April 2008

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
QOS-LI: QOS LOSS INFERENCE IN DISADVANTAGED NETWORKS - PART II	In House

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
62702F

6. AUTHOR(S)
Vidyaraman Sankaranarayanan, Shambhu Upadhyaya, and Kevin Kwiat

5d. PROJECT NUMBER
4519

5e. TASK NUMBER
22

5f. WORK UNIT NUMBER
49

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

AFRL/RIGG   University of Buffalo
525 Brooks Road   Computer Science and Engineering Dept.
Rome, NY 13441-4505   Buffalo, NY 14260

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/RIGG
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TP-2009-60

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for public release; distribution unlimited PA# WPAFB-2008-0350 Date Cleared: 12-February-2008

13. SUPPLEMENTARY NOTES
This paper was presented at and published in the proc. of the IASTED Conference on Wireless and Optical Communications,Quebec City,Quebec Canada, 26-28 May 2008. This work is copyrighted. One or more of the authors is a U.S. Government employee working within the scope of their Government job; therefore, the U.S. Government is joint owner of the work and has the right to copy, distribute, and use the work. All other rights are reserved by the copyright owner.

14. ABSTRACT
The Quality-of-Service (QoS) in disadvantaged networks is a function of many parameters, including the nature of the physical link over which the disadvantaged network operates. While there exist mechanisms (like specialized versions of TCP) to account for their nature and operate optimally over disadvantaged networks, their operation under adversarial conditions have not been investigated. In the authors' prior work, they presented a game theoretic framework to infer the nature of a QoS loss in disadvantaged networks. In this work, they present the translation of the theoretical framework to a satellite based disadvantaged network. This paper shows the feasibility of the game theoretic formulations in satellite networks through simulations in Opnet.

15. SUBJECT TERMS
Disadvantaged Networks, Game Theory, QoS, Resource Allocation, Satellite Networks, Multi-Armed Bandit Problem

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Kevin A. Kwiat
U	U	U	UU	8	19b. TELEPHONE NUMBER (Include area code) N/A

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

QoS-LI: QoS Loss Inference in Disadvantaged Networks – Part II

Vidyaraman Sankaranarayanan and Shambhu Upadhyaya
Computer Science and Engineering
University at Buffalo
Buffalo, USA
Email: {vs28, shambhu}@cse.buffalo.edu

Kevin Kwiat
Air Force Research Laboratory
525 Brooks Road
Rome, NY 13441
Email: kwiatk@rl.af.mil

Abstract

The Quality-of-Service (QoS) in disadvantaged networks is a function of many parameters, including the nature of the physical link over which the disadvantaged network operates. While there exist mechanisms (like specialized versions of TCP) to account for their nature and operate optimally over disadvantaged networks, their operation under adversarial conditions have not been investigated. In our prior work [17], we presented a game theoretic framework to infer the nature of a QoS loss in disadvantaged networks. In this work, we present the translation of the theoretical framework to a satellite based disadvantaged network. We show the feasibility of the game theoretic formulations in satellite networks through simulations in Opnet.

Keywords: Disadvantaged Networks, Game Theory, Multi-Armed K Bandit Problem, QoS, Resource Selection, Satellite Networks.

1. INTRODUCTION

Wireless and satellite networks are examples of disadvantaged networks where the very nature of the physical medium restricts the effective bandwidth. The Quality of Service (QoS) in disadvantaged networks is thus limited by the nature of the physical medium. For example, satellite networks with geosynchronous satellites have a per packet delay of at least 500 ms, that is attributed to the round trip time. Consider a situation where a complex backend system provides some services to the remote system via a satellite network, as shown in Figure 1. The effective QoS between the two systems are also dependent on the nature of the satellite network. A QoS loss to the remote system may be due to several reasons:

- Statistical variations on the satellite network
- Adversarial manipulation of the network (through a Jammer)
- Backend operating under hostile conditions
- Problems on the host of the remote platform

As a practical motivation, consider the case where the remote backend is a Self Regenerative System (SRS) [6]. SRS are complex backend systems composed of various interconnected components that provide a 'self regenerative' capability; thus an SRS under attack could restructure itself in order to continue providing services. An SRS conceptually extends the notion of healing [18] from biological systems into computer systems. The goal [6] of an SRS is to provide cognitive immunity and regenerative capabilities to computer systems. While a number of SRS have been designed (like key distribution schemes [5] and hardware platforms [9]), a standard assumption across all systems is the existence of an enterprise class network with ample bandwidth for the system under consideration (with the exception of mobile networks [15, 14]).


Figure 1: **Complex Backend System Servicing Remote Node**

In our prior work [17], we considered the problem of an SRS operating over a disadvantaged network (such as a satellite network) and presented a game theoretic formulation to infer the nature of a QoS loss over a disadvantaged network. The goal was to infer if a QoS loss is due to:

- adversarial network control or
- statistical variation in the network conditions or
- drop in the service of the SRS due to hostile conditions (DDoS attack) on the backend system

The game theoretic framework is based on the *Multi-Armed K Bandit* [7, 11] problem. This work takes as its point of departure the framework presented in [17]. In this paper, we present how the theoretical framework is translated to a practical implementation for satellite networks. We also present the viability of such a translation by means of OPNET simulations.

1.1. Related Work

The QoS of satellite networks is limited by the physical characteristics of the medium and setup. For example, the round-trip-time of a radio signal for a geosynchronous satellite is at least 500 ms due to the distance of the satellite from the base station. Moreover, the application of standard protocols, like TCP over IP results in very poor performance over satellite networks due to the link layer characteristics. TCP, for example, increases its congestion window size if an acknowledgement is not received "on time". However, the very nature of satellite networks (long propagation delays, large *bandwidth·delay* product and high bit error rates) implies a higher time for acknowledgement, which causes traditional TCP to assume congestion, when there is none. The work [16] gives an overview of the problems with TCP over satellite networks and proposes solutions for the same. Several variants of TCP [10] have been proposed and implemented to account for such deviations from standard behavior. These protocols are tuned so as to adjust for the characteristics of the link layer. For example, the Satellite Transport Protocol (STP) [19] is tuned for satellite networks. Standard TCP is also used in conjunction with Performance Enhancing Proxies (PEPs) [3]. PEPs are hardware devices or software components that are located at the ingress and egress of a satellite network. The objective of the ingress PEP is to send an acknowledgement to the TCP originator and then forward the packets over the satellite link to the egress PEP. Such PEPs are said to split the TCP connection. PEPs can also be application transparent and/or protocol transparent. SaTPEP [20] is an example of a PEP for satellite networks, which employs negative acknowledgements to recover from errors. A common underlying theme of all these works is the consideration of QoS only from the network viewpoint. To the best of our knowledge, our work is the first to consider the problem of inferring the nature of a QoS loss, where not only network conditions, but also an attack on the backend could cause a service loss.

The rest of this paper is organized as follows. Section 2 gives a synopsis of the game theoretic framework. In section 3, we detail the translation of the framework to a satellite network and present preliminary OPNET simulation results for the same. Concluding remarks are presented in section 4.

2. OVERVIEW OF THE *QoS-LI* MODULE

Consider the end-points of the disadvantaged networks shown in Figure 1. Let N1 denote the backend system and N2 denote the remote system. The only means of communication between N1 and N2 is the disadvantaged network: in this case, this is the satcom. The *QoS-LI* module is a component at the end-points (on both N1 and N2). *QoS-LI* assumes that the endpoints have multiple (logical) connections between them, through the same disadvantaged (physical) channel. In this case, the physical channel is the satcom. The QoS is limited by the physical channel: for a geo synchronous satellite, this limitation is expressed in terms of the round-trip time of a packet being at least 500ms. The *QoS-LI* module uses a metric to measure the QoS between N1 and N2; for purposes of simulation, we use the per-packet end-to-end delay, although the metric could be changed to any other meaningful function, based on the nature of the disadvantaged network.

QoS-LI infers the nature of QoS loss between N1 and N2 based on a game theoretic formulation called the *Multi-Armed K Bandit* (abbreviated as *maKb*) problem. The *maKb* problem is a classical tradeoff scenario between the notion of *exploration* and *exploitation*. The setting is as follows: A gambler is a casino has K slot machines to play on. The payoff from each slot machine is not known. The objective of the gambler is to maximize the payoff over a period of N trials. The gambler has two options: try out each of the K slot machines and infer the one with the best payoff distribution (*exploration*), or stop when a slot machine offers an adequate enough payoff and keep playing that machine for the remaining trials (*exploitation*). The payoff of the slot machines may be stochastically distributed or adversarial (deterministically) controlled. Within this scenario, there exist algorithms to converge to the slot machine with the best possible payoff under the different conditions viz., stochastic distribution or adversarial control [13] of slot machine payoffs: the application of these algorithms to the QoS Loss Inference module is described in our prior work [17].

The correlation between the *maKb* game and the loss inference is drawn as follow: each of the slot machines represent multiple links between the nodes N1 and N2. The slot machine's payoff is equivalent to the link payoff, which in our situation is the end-to-end packet delay. The *QoS-LI* module applies the *maKb* algorithms as follows: a *setup stage* is first initiated between the SRS backend and the remote system, and a *detection stage* is switched on if a QoS loss is found to occur.

The setup stage is assumed to occur when the backend system is not under attack and the network is free of adversarial control. The setup stage follows greedy algorithm, in terms of choosing the link with the best payoff, and sets the maximum possible QoS. The greedy strategy and its variants are described in [4]; the work [8] details empirical evaluation that supports the hypothesis that this strategy is the best in terms of achieving greatest payoff. At the end of the setup stage, the nodes N1 and N2 communicate on a dedicated link (which is the outcome of the setup stage). A

default standard of the observed QoS is also set for comparison purposes.

QoS-LI switches to the detection stage if a drop in the effective observed QoS is detected. In the detection stage, the modules multiplex packets over the other links and observe the convergence of the observed QoS. If this convergence occurs within a specified duration of time, the channel may be assumed to have stochastic variations. If the QoS convergence is not observed over the specified time interval, the module switches to an adversarial detection stage, where the packets are multiplexed over the links assuming deterministic (adversarial) control. This algorithm is expected to converge to a QoS (limited by the adversarial manipulation) in a larger period of time, thereby providing the inference that the network is adversarial controlled. If the QoS (still) remains low, it may safely be assumed that the backend is under attack, and hence the effective service quality is low (and will remain so despite any switching strategy).

A complete description of the setup stage and detection stage algorithms is given in [17]. With this background on the *maKb* algorithms and its application to the general QoS loss inference problem, we show its translation to a practical domain, with satcom's as the disadvantaged network.

3. *QOS-LI* APPLICATION ON SATCOM

The *QoS-LI* module proposed in [17] assumes the existence of multiple links between the remote system and the SRS backend. This may physically be true in the case of mobile networks (802.11x) where the last link may operate on different channels, with a suitable receiver at the remote system. However, in a satellite network, the uplink and the downlink are different insofar as their transmission characteristics are concerned; they are both time and frequency multiplexed, with each incoming packet stream transmission being subject to a reservation on the uplink and scheduling on the downlink. The uplink is a multiple access channel, while the downlink is a statistically multiplexed broadcast channel. We derive the motivation for the practical translation of the *QoS-LI* module to satellite networks from the work by Pandya et. al, [12], where the authors proposed a dynamic resource assignment algorithm for effective link layer utilization. This work uses a similar approach in terms of link layer resource allocations for obtaining multiple links on the satcom.

In this work, we consider a packet switched military satellite network (similar to [12]), with a provision for dynamic resource allocation on the uplink and downlinks. Each terminal in a satcom is capable of operating in multiple transmission modes (burst rates). The uplink and downlink are thus characterized by the modulation format (QPSK, BPSK, etc), the coding rate and the data burst rates. This triple is referred to as the *transmission mode* [12]. *QoS-LI* uses multiple transmission modes as the logical equivalent of multiple links over the same physical channel. The testing approach in this work is an absolutist one: we verify that if a consistent transmission mode were used, providing a predetermined (statistically variant) QoS level to the remote system, then a variation in the QoS due to a Jammer (at the SRS end) will be detected by a suitable switch of the transmission mode. With this verification, any resource assignment algorithm (predictable or dynamically assigned) may be used for normal communications: when the QoS drops below a predetermined level, the transmission modes may be switched and an inference on the existence of a Jammer could be made based on the observed QoS.

3.1. Simulation Setup

Figure 2: Single Aggregation Point Simulation Setup

Figure 2 depicts the simulation setup overview. We simulate the *QoS-LI* module with a single aggregation point, i.e., we assume all streams from the remote backend are aggregated at one point (the ingress). It is also possible for satellite networks to have multiple uplink points (when the number of discrete terminals is high), as shown in Figure 3. In the context of this work, figure 3 depicts the situation when multiple remote terminals send back data to the single backend (the reverse situation of figure 2).

The single aggregation point usually transmits the uplink in a time and frequency multiplexed manner, with a scheduling algorithm, to account for different types of streams. For purposes of the simulation, we use a single source traffic distribution pattern (Constant Bit Rate). Other traffic distributions could also be used to detect a QoS drop/pickup after a transmission mode change. A jammer (not shown in figures) is also introduced at the aggregation point (base station) to impede communications to the remote

end. The satellite simulations are performed in OPNET [1], using the wireless module.

Figure 3: Multiple Uplink Points for Single Sink

3.2. Jammer Characteristics

OPNET includes three types of Jammers for simulation purposes. They are:

- single band jammer
- pulsed jammer
- frequency swept jammer

A single band jammer, as the name suggests, transmits signals on a single predefined frequency band. A pulsed jammer is usually used to target a frequency hopping system by targeting its power to a narrow spectrum, thereby hitting a minimal set of frequency hops of the target system. Like the single band jammer, it also transmits on a single fixed frequency band, but masks it with a periodic pulse. A frequency sweep jammer concentrates its power on a particular frequency range, which is constantly shifted in order to cover a wide range of operating frequencies. In this work, we use a single band jammer to impede satellite communications. This provides the best test case for the switch in the transmission mode; a pulsed or a frequency swept jammer would have the same (but limited) effect on the outgoing signal, and its effect on the transmission switch would be perceived only in the overlapping pluses or frequency ranges. We used a single band jammer with a similar frequency range and packet source as the SRS sources.

3.3. OPNET Setup

Figure 4 shows the OPNET setup, with a SRS and a receiver connected only by means of a satellite link. The SRS is modeled as a simple source with a packet multiplexing component that represents the scheduler for traditional satellite uplinks. The satellite receives the SRS signals (operating at the same frequency) and queue's them for processing and transmits them to the receiver. The processing is mod-

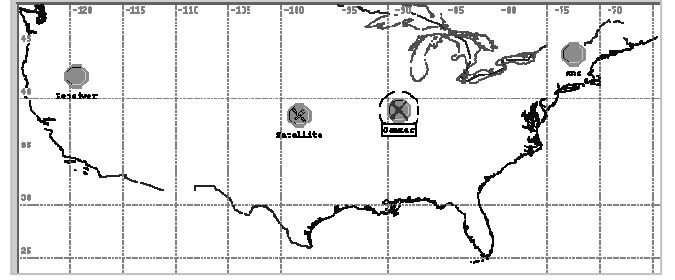eled as a simple delay, with a constant service rate at the queue.

Figure 4: OPNET Simulation Setup

When evaluating resource assignment algorithms, the processing takes form of multiplexing the received streams for broadcast, and in some cases, priority processing. However, in this situation, since we use a single stream, a service delay is appropriate to abstract the normal processing time. The QoS is measured at the receiver, and is currently the end-to-end packet delay in seconds. Note that this delay is at least 250 ms for one trip from the SRS to the remote receiver. The jammer is also similarly set, with frequencies matching the SRS transmitter (which in turn matches the satellite receiver). The jammer is a single band jammer, but with a similar source rate as that of the SRS, instead of the constant rate of 1 per second provided by the default OPNET template.

3.3.1. Simulation Parameters and Assumptions

The transmitting nodes' propagation delay is set to 250 ms; the satellite orbit is set to a geosynchronous orbit (hence no orbital paths are defined in OPNET). The `channel-match` and `closure` properties of the transmitting and receiver pair are set appropriately to ensure that the SRS signal does reach the satellite (and the Jammer, when operational, interferes with the SRS-Satellite link). All terminals are assumed to transmit at full power. Since the `channel-match` and `closure` properties are set to ensure reception, the power variation does not change the simulation results in this context.

The OPNET simulations make the following assumptions with respect to the traffic streams and the Jammer.

1. The traffic from the source to the destination is assumed to be a Constant Bit Rate (CBR) stream. In a typical satellite uplink, multiple sources 'submit' their streams for transmission; a FTP server might submit a file upload, a VOIP/multimedia server would submit a real-time stream. These streams are frequency and time multiplexed, based on a reservation strategy that attempts to deliver real-time traffic with QoS guarantees while maintaining fairness for the non-real-time (FTP)

traffic. Although we do not introduce multiple streams, the simulation results hold true since the QoS metric would include the summation of all streams (which would be affected by a Jammer).

2. We assume that the Jammer cannot predict the transmission mode switching pattern; if an adversary were able to do so, then the Jammer transmission mode could be synced with the satellite transmission, thereby affecting the QoS in a more granular level. In this case, although we would (theoretically) be assured of a optimal QoS convergence [13], the adversary, by suitably manipulating the QoS, could convince the remote end of an attack on the SRS (as opposed to a Jammer operation).

3. The simulation uses only one terminal for testing all the burst rates; in practice, no single terminal will be capable of handling all the burst rates and power requirements. However, a typical satellite ground station will have multiple terminals capable of handling the required burst rates: this is a standard assumption for satellite simulations (unless the simulation involves path fading and/or directional antenna characteristics).

3.3.2. Simulation Results

The uplink and downlink are characterized by the bandwidth, the burst rate and the modulation. As derived from [12], the combination of the typical values for satellite operation are form around 20 transmission modes, each with a different power requirement. The values for the Bandwidth are 4/16/64/256 MHz and the data burst rates are 2/8/32/128 Kb/time slot, where a time slot is 500 ms (the round trip time). For normal operation without a Jammer, with uplink (BPSK, 1Mhz, 1000 kbps), downlink (BPSK, 4MHz, 8000 kbps) with a Constant Bit Rate (CBR) source, the throughput and end-to-end delay curves are shown in Figure 5. The throughput levels off and the end-to-end delay remains constant, around 0.7 seconds. Similarly, the throughput and end-to-end delay curves for Uplink (BPSK, 4Mhz, 8000 kbps), downlink (BPSK, 4MHz, 8000 kbps) with CBR source is shown in Figure 6. As expected, the end-to-end delay is rises linearly since the data burst rates on the uplink and downlink are equal (the processing time on the satellite keeps packets in the service queue). With a limited buffer size, this value would level off soon. With an uplink (BPSK, 4Mhz, 8000 kbps), and downlink (BPSK, 8MHz, 16000 kbps) with CBR source, the curves shown in Figure 7 show a structure similar to Figure 5, but for an improved end-to-end delay (since the downlink bandwidth and burst rates are both higher). The throughput is also seen to approach its asymptotic value around 1 minute and 31 seconds.


Figure 5: Throughput and End-to-End Delay curves: Uplink (BPSK, 1Mhz, 1000 kbps); Downlink (BPSK, 4MHz, 8000 kbps)


Figure 6: Throughput and End-to-End Delay curves: Uplink (BPSK, 4Mhz, 8000 kbps); Downlink (BPSK, 4MHz, 8000 kbps)

Thus, as the transmission mode is changed, the effective QoS (measured in terms of end-to-end delay) is also seen to change. This provides validation that different transmission modes can serve as the logical equivalent of multiple links over the same channel. A Jammer that is successful in lowering the QoS in one transmission mode will not be as effective in another (switched) transmission mode. Figure 8 shows the end-to-end delay and throughput curves with a jammer operating after 210 seconds of SRS/satellite operation. As expected, the end-to-end delay immediately drops and the throughput gradually drops after the jammer starts operating (the jammer operates on the same frequency bandwidth as the SRS-satellite link).

Figure 7: Throughput and End-to-End Delay curves: Uplink (BPSK, 4Mhz, 8000 kbps); Downlink (BPSK, 8MHz, 16000 kbps)


Figure 8: Throughput and End-to-End Delay curves: Jammer operating with exaggerated data rate of 512 kbps

The hypothetical throughput curve for the SRS-satellite link after a transmission mode switch is shown in Figure 9 (this is Figure 8 spliced with figure 7; they are operating in a different transmission modes).


Figure 9: Hypothetical Operation after transmission mode switch

With a transmission mode switch, the single band Jammer is ineffective (in an absolutist sense; a pulsed or frequency sweep jammer would have some effect even after a transmission switch, but the effective observed QoS would

still rise, albeit at a slower pace): in this case, the frequencies do not overlap; in other situations, the data burst rate / modulation also play a role. The jammer effectiveness, unlike the hypothetical situation in Figure 9, may not be completely lowered, but the remote end would perceive a raise in the effective QoS, thereby leading to the inference of a Jammer at the backend (as opposed to an attack on the SRS).

3.3.3. Applicability

In a typical satellite system, there are many other mechanisms in place for ensuring QoS. Dynamic Resource Allocation (DRA) algorithms like the one proposed in [12] change the transmission mode every epoch. These allocation mechanisms can proceed independently of the scheme proposed in this work. Once a drop in the QoS is observed, the remote system can initiate a deterministic (or dynamically evaluated) transmission mode switching strategy to infer the nature of the QoS loss. This argument applies to other link layer optimizations for QoS. The remote system is typically a individual terminal [2] that are marketed by commercial entities. The *QoS-LI* module can be implemented as a software component on the remote system or integrated as a plug-in to existing QoS optimization mechanisms. Depending on the rate of the QoS drop/rise, the inference can be made in real-time, in as low as 10 epochs (5 seconds). The actual time convergence, however, requires further investigation. The existence of PEPs [3] does not affect the *QoS-LI* module; the placement options for the *QoS-LI* module with respect to PEPs have been discussed in our prior work [17].

4. CONCLUSION

This work presents the implementation of a QoS Loss Inference (QoS-LI) module on satellite communication (SATCOM) networks. We presented a translation of theoretical *QoS-LI* module [17] to a satcom, using multiple burst rates and modulation formats as a logical equivalent of multiple links. In the presence of a QoS loss due to statistical variations, it is possible for the remote system and the local SRS backend to engage in an appropriate logical-link switching mode to ensure a return to the previously assured service quality level. In the presence of a Jammer with limited or deterministic frequency hopping, the return to a previously assured QoS level is expected to take relatively more time; in the presence of an attack on the backend network, the end-to-end delay (or any other appropriate metric choice) is not expected to improve. This work verifies through simulations that the logical equivalent of multiple links can be obtained in a satcom by switching between transmission modes. The assumptions made only serve to simplify the simulation setup and do not invalidate the result in a general setting.

However, future work in this area will investigate the usage of metrics other than end-to-end delay to infer QoS loss, when multiple streams of data are sent to the base station. Although the theoretical expectation is still the same with respect to the SRS related traffic, we also intend to simulate the behavior of different applications such as FTP, multimedia services, VOIP, etc., multiplexed over the same satcom link.

Acknowledgements

References

[1] *OPNET Technologies, Inc.: http://www.opnet.com*, 2007.

[2] AOS SkyPipe, *Secure Communications Link Optimization from AOS: http://www.aosusa.com/scrunch.html*, 2004.

[3] J. Border, M. Kojo, J. Griner, G. Montenegro and Z. Shelby, *Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations*, RFC Editor, 2001.

[4] C. J. C. H. Watkins, *Learning from Delayed Rewards*, Cambridge University, 1989.

[5] B. Carlo, D. Paolo, arco, S. Alfredo De and L. Massimiliano, *Design of Self-Healing Key Distribution Schemes*, Des. Codes Cryptography, 32 (2004), pp. 15-44.

[6] DARPA, *Self Regenerating Systems: http://www.darpa.mil/ipto/programs/srs/*, 2004.

[7] H. Robbins, *Some aspects of the sequential design of experiments*, Bulletin American Mathematical Society, 55 (1952), pp. 527–535.

[8] Joannès Vermorel and Mehryar Mohri, *Multi-Armed Bandit Algorithms and Empirical Evaluation*, ECML, 2005.

[9] P. K. Lala and B. K. Kumar, *An Architecture for Self-Healing Digital Systems*, J. Electron. Test., 19 (2003), pp. 523-535.

[10] M. Allman, D. Glover and L. Sanchez, *Enhancing TCP over Satellite Channels using Standard Mechanisms (RFC 2488)*, 1999.

[11] Michel Benaïm and Gerard Ben Arous, *A two armed bandit type problem*, International Journal of Game Theory, 32 (2003), pp. 3.

[12] J. Pandya, A. Narula-Tam, H. Yao and J. Wysocarski, *Network layer performance of a satellite network with dynamic link-layer resource allocation*, International Journal of Satellite Communications and Networking, 25 (2007), pp. 217-235.

[13] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund and Robert E. Schapire, *Gambling in a Rigged Casino: The adversarial multi-armed bandit problem In Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, 1998.

[14] M. Rui and I. Jacek, *Regenerating Nodes for Real-Time Transmissions in Multi-Hop Wireless Networks*, *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks (LCN'04) - Volume 00*, IEEE Computer Society, 2004.

[15] M. Rui and I. Jacek, *Reliable Multipath Routing with Fixed Delays in MANET Using Regenerating Nodes*, *Proceedings of the 28th Annual IEEE International Conference on Local Computer Networks*, IEEE Computer Society, 2003.

[16] S. Oueslati-Boulahi, A. Serrhrouchni, S. Tohme, S. Baier and M. Berrada, *TCP over satellite links: Problems and solutions*, Telecommunication Systems, V13 (2000), pp. 199-212.

[17] S.Vidyaraman, S. Upadhyaya and K. Kwiat, *QoS-LI: QoS Loss Inference in Disadvantaged Networks*, *Proceedings of 2007 IEEE International Symposium on Ubisafe Computing (UbiSafe '07)*, Niagara Falls, Ontario, Canada., 2007.

[18] G. Selvin, E. David and M. Steven, *A biological programming model for self-healing*, *Proceedings of the 2003 ACM workshop on Survivable and self-regenerative systems: in association with 10th ACM Conference on Computer and Communications Security*, ACM Press, Fairfax, VA, 2003.

[19] T. Henderson and R. Katz, *Satellite transport protocol (STP): An SSCOP-based transport protocol for datagram satellite networks*, *Workshop on Satellite-Based Information Services (WOSBIS)*, Budapest, Hungary, 1997.

[20] D. Velenis, D. Kalogeras and S. B. Maglaris, *SaT-PEP: A TCP Performance Enhancing Proxy for Satellite Links*, *Proceedings of the Second International IFIP-TC6 Networking Conference on Networking Technologies, Services, and Protocols*, Springer-Verlag, 2002.