ELB: An Explicit Load Balancing Routing Protocol for Multi-Hop NGEO Satellite Constellations

Tarik Taleb[1,*], Daisuke Mashimo[1,†], Abbas Jamalipour[2], Kazuo Hashimoto[1,*], Yoshiaki Nemoto[1,‡], and Nei Kato[1,†]
[1] Graduate School of Information Sciences
Tohoku University, Japan
*{taleb, kh}@aiet.ecei.tohoku.ac.jp
†{mashimo, kato}@it.ecei.tohoku.ac.jp
‡nemoto@nemoto.ecei.tohoku.ac.jp
[2] School of Electrical & Information Engineering
University of Sydney, Australia
a.jamalipour@ieee.org

Abstract— **Due to geographical and/or climatic constraints, the community of future satellite users will exhibit a significant variance in its density over the Globe. This density variance will yield a scenario where some satellite links are congested while others are underutilized. To ensure an intelligent engineering of traffic over satellite networks, this paper proposes a routing protocol that enables neighboring satellites to explicitly exchange information on their congestion status. A "soon-to-be-congested" satellite requests its neighboring satellites to decrease their data forwarding rates. In response, the neighboring satellites search for less congested paths that do not include the satellite in question and communicate a portion of data, primarily destined to the satellite, via the retrieved paths. By so doing, congestion, and the resulting packet drops, can be avoided. A better distribution of traffic among satellites can be guaranteed as well. The proposed scheme is dubbed "Explicit Load Balancing" (ELB) scheme. A set of simulations is conducted to evaluate the performance of the ELB scheme using the Network Simulator. In terms of Quality of Service, encouraging results are obtained: better traffic distribution, higher throughput, and lower packet drops.**

I. Introduction

To provide global communication with reasonable latency and low terminal power requirements, constellations made of multi Non-Geostationary (NGEO) satellites, known as Low Earth Orbit (LEO) and Medium Earth Orbit (MEO) satellites, have been the focus of several researches in the recent literature [1]. Given the highly dynamic feature of these constellations, along with the high variance in the density of their future users, multi-hop satellite constellations should be prepared to face a challenging scenario where some of their satellites are congested while others are underutilized. Without an efficient routing algorithm, this unfair distribution of network traffic will lead to significant queuing delays and large number of packet drops at the congested satellites.

For an efficient routing over satellite networks, a number of routing protocols have been proposed. Most of these protocols search for the shortest path with the minimum routing cost. Furthermore, the computation of their routing metric does not take into account the total traffic distribution over the entire constellation. As a remedy to the aforementioned issue, this paper proposes an explicit exchange of current congestion status among neighboring satellites. A satellite with high traffic load notifies its neighboring satellites of its current status and requests them to reduce their data forwarding rates. In response, neighboring satellites reduce their transmission rates of traffic originally destined to the "soon-to-be-congested" satellite and search for other alternative paths that do not include the satellite. The proposed scheme is dubbed Explicit Load Balancing (ELB). In the ELB mechanism, satellites use three parameters to indicate their congestion status and to reduce their data transmission rates, respectively. These parameters consist of two queue ratio thresholds and a traffic reduction ratio, respectively. While a simple version of the proposed scheme has been partly investigated in [2], the major improvements presented in this paper consist in a dynamic setting of the system parameters based on an easy to implement mathematical model. Furthermore, in the presented work, we extend our performance evaluation to more general scenarios considering real satellite constellations (e.g. Iridium).

The remainder of this paper is organized as follows. Section II showcases some routing protocols specifically designed for multi-hop NGEO satellite constellations. Section III introduces the ELB scheme with the dynamic settings of its parameters. The performance of ELB is evaluated and compared to other schemes in Section IV. Conclusion is drawn in Section V.

II. Related Work

The intrinsic features of NGEO satellite constellations (e.g. frequent topological variations) have motivated the development of a wide library of new routing protocols. In [2], the authors provide a thorough discussion on the main credits and downfalls of these routing protocols. A common drawback of these protocols is that they are mostly designed for only best-effort light-load traffic. Indeed, they seek for the shortest delay paths without paying much attention to other QoS parameters.

In attempt of supporting QoS in NGEO satellite constellations, several researchers have first considered the integration of satellite networks with the Asynchronous Transfer Mode (ATM) as it originally provides different levels of QoS guarantees [3] [4]. However, the overwhelming growth of applications based on the Internet Protocol (IP) has attracted the attention of satellite operators to IP traffic. A number of solutions have been thus proposed for an efficient QoS routing over IP-based satellites. In [5], a Traffic Class Dependent (TCD) routing algorithm is proposed. In TCD, packets belonging to different traffic classes are provided with different levels of services. The TCD protocol whilst attempts to guarantee QoS for different traffic classes, it assigns a single route for each class. In case a class contains huge traffic data, the chosen path would be heavily overloaded with data packets. Such a

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

scenario ultimately affects the balancing of traffic load over the entire satellite constellation. In [6], a Multi-Protocol Label Switching (MPLS) routing protocol is proposed for NGEO satellite systems. Despite the promising performance of the protocol, issues related to rerouting and maintenance overhead remain unsolved and put the practicality of the protocol in question.

As mentioned earlier, most conventional routing algorithms base their route decision primarily on propagation delays. Such a decision usually yields poor QoS routing in case of heavy loads. This is due to the large contribution of queuing delays in the total communication delay. In this regard, a more appropriate routing cost metric has to be selected. In [7], a Probabilistic Routing Protocol (PRP) is proposed. The PRP scheme uses a cost metric as a function of time and traffic load. The traffic load is assumed to be location homogeneous. The major drawback of the protocol consists in this assumption as it is far away from being realistic. Indeed, newly coming traffic can easily congest the chosen PRP path and leave other resources underutilized. [8] proposes a Minimum Flow Maximum Residual (MFMR) routing protocol where the minimum-hop path with the minimum number of flows is selected. One of the main drawbacks of the protocol consists in the fact that it implies knowledge of the flows over the constellation and does not consider the case where the flows count increases along the selected path. Given the fast movements of satellites, such scenario may occur frequently. This would lead to the congestion of the chosen MFMR paths and ultimately unfavorable performance. In [9], Jianjun *et. al.* propose a Compact Explicit Multi-path Routing (CEMR) algorithm based on a cost metric that involves both propagation and queuing delays. At a given satellite, the queuing delay is predicted by monitoring the number of packets in the outgoing queue of the satellite over a time interval. It is assumed that the network state over each time interval is updated before routing calculation is carried out. While the used cost metric gives a good insight about the queuing delay that may be experienced by a packet at a given satellite, it does not reflect the congestion state of the next hop, nor does it estimate the queuing delay a packet may experience there. It does not reflect the likeliness of packets to be dropped by the downstream hop either. Taking these remarks into account, the research work outlined in this paper aims to develop a routing strategy where packet drops are avoided and traffic burden is efficiently and fairly distributed among all participating satellites.

III. Operational Overview of the ELB Scheme

Before delving into details of the ELB scheme and its parameter settings, we first list the key components used in the scheme. A multi-hop NGEO satellite constellation consists of S satellites uniformly distributed over N orbits, forming a mesh network topology. Each satellite is able to set up a maximum of M links with its neighboring satellites. These links are called Inter Satellite Links (ISLs).

A satellite is considered to be in a Free State (FS) when the queue ratio of its current queue occupancy to the total queue size, Q_r, is inferior to a pre-defined threshold α. Having the queue ratio between the threshold α and another predetermined threshold β, the satellite is considered to be in a Fairly Busy State (FBS). The satellite is considered to be in a Busy State (BS) if its queue ratio exceeds the threshold β.

Assuming that each satellite is aware of its neighboring satellites, satellites mutually and dynamically exchange information on the states of their queue occupancies. Indeed, when a satellite A experiences a state transition from free to fairly busy, it sends a warning message to its neighboring satellites informing them that it is about to get congested. The neighboring satellites are then requested to update their routing tables and start searching for alternate paths that do not include satellite A. When the satellite enters the busy state, it transmits a Busy State Advertisement (BSA) signaling packet requesting the neighboring satellites to reduce their sending rates of traffic destined to satellite A by a ratio χ. The $(1-\chi)$ portion of traffic data will be transmitted via alternate paths retrieved earlier. The BSA signaling packet carries information on the satellite ID and the Traffic Reduction Ratio (TRR) χ. The remainder of this section portrays the setting procedure of the system parameters α, β, and χ.

A. Setting of Queue Ratio Thresholds

The key philosophy behind an optimum setting of β and α is to reflect the packet discarding probability in these two parameters so as to avoid packet drops when a satellite is running under heavy loads. Let I and O denote the total input and output traffic rates at a given satellite, respectively. Let Q_l and $q(t)$ denote its total queue length and its queue occupancy at time t, respectively. Assuming that the input and output traffic rates constant over a short period of time, the elapsing time till a packet drop occurs can be simply expressed as follows.

$$\delta_d = \frac{(Q_l - q(t)) \cdot P_{avg}}{I - O} \tag{1}$$

where P_{avg} is the average packet size. If the satellite is assumed to monitor its queue occupancy every δ interval time, it needs a maximum of $(\delta + d)$ time to notify its neighboring satellites of a possible packet drop, where d indicates the ISL delay. In this case, two scenarios can be envisioned:

- $\delta_d \leq \delta + d$: Packet drops happen before neighboring satellites are notified and adequate measures are taken. In this case, the packet dropping probability is one $(p = 1)$.
- $\delta_d > \delta + d$: In this case, if the satellite keeps receiving and transmitting data at same rates over a number of monitoring intervals, packet drops happen only once during $\left(\frac{\delta_d}{\delta+d}\right)$ times of monitoring operations. The packet dropping probability is thus $\frac{\delta+d}{\delta_d}$.

The packet dropping probability can be expressed as

$$p = Min(1, \frac{\delta + d}{\delta_d}) \tag{2}$$

To reflect the packet dropping probability in the setting of β, we set β to $(1-p)$[1].

$$\beta = 1 - p \tag{3}$$

The rational behind this setting is that when traffic load gets heavy and p gets higher values, β should be set to small values so as the satellite would transit quickly to the busy state and neighboring satellites would be promptly requested to reduce their sending rates to avoid possible congestion and packet drops. In this regard, it should be noted that setting the monitoring interval δ to high values may lead to significant packet drops. Indeed, in case of long monitoring intervals, by the time a satellite monitors its queue length, congestion may have already occurred and packet drops become then inevitable. In such case, the packet dropping probability will be equal to one ($p = 1$). Consequently, β will be always set to zero. As a remedy to this issue, the satellites are assumed to monitor their queues in a real time fashion. Therefore, δ is set to $1ms$ throughout this paper.

An optimum setting of the threshold α is a tradeoff between twofold. First, with small values of α, neighboring satellites can be granted a time long enough to carry out their search for alternative paths before they are asked to detour their traffic from the congested satellite. Second, with high values of α, neighboring satellites are requested to search for alternative paths only when it is necessary, in other words less frequently. This improves the warning accuracy and reduces the overhead that may result in case of frequent searches for alternative paths. Considering these two observations and for the sake of the scheme simplicity, α is set to half of β.

$$\alpha = \frac{\beta}{2} \tag{4}$$

B. Setting of the Traffic Reduction Ratio

The main objective behind the setting of the TRR parameter is to allow satellites to return back to their free state and reside in this state for at least a predetermined period of time θ. Let I_t and I_s denote the total rate of traffic coming from terminals within the coverage area of a satellite and that of traffic coming from neighboring satellites, respectively. When the satellite shifts to the busy state, it requests neighboring satellite to reduce their sending rates. By the time the BSA signaling packet reaches the neighboring satellites, the queue occupancy of the satellite is

$$q(t_{BSA}) = Min\big(Q_l \cdot \beta + \frac{d \cdot (I_s + I_t - O)}{P_{avg}}, Q_l\big) \tag{5}$$

So as that the satellite is ensured a prompt recovery and a residual time in the normal state for at least θ time, the new rate of traffic coming from neighboring satellites, I_s^{new}, should satisfy the following equation

$$(I_s^{new} + I_t) - O = \frac{P_{avg} \cdot \big(q(t_{BSA}) - Q_l \cdot \alpha\big)}{\theta} \tag{6}$$

[1] This setting can be further enhanced using the Exponentially Weighted Moving Average (EWMA). In the conducted simulations, the use of EWMA is not considered.

The TRR parameter can be accordingly computed as

$$\chi = Min\Big(Max(0, \frac{I_s^{new}}{I_s}), 1\Big) \tag{7}$$

By this setting, a long enough recovery time can be granted for satellites before they enter again the busy state and request their neighboring satellites to reduce further their sending rates.

Another important feature that should be also taken into account in the design of the ELB scheme is related to the traffic redistribution cascading issue that may occur to the already-detoured portion of traffic. Indeed, as previously explained, to avoid the congestion of a satellite, neighboring satellites are requested to transmit a portion of traffic data via paths that do not include the satellite in question. At this stage, the system should ensure that the detoured portion of traffic does not experience further detouring along the selected paths till the destination. For this purpose, we use a routing metric that instantly reflects both the one-way propagation delay and the instant queuing delay. This is similar in concept to the idea of CEMR [9]. We assume that routing tables are updated periodically every Δ interval time. At time t, the instant path cost is defined as

$$L_{cost}(t) = T_d + T_B(t) \tag{8}$$

where T_d denotes the one-way propagation delay. $T_B(t)$ denotes a predicted value of the queuing delay at time t and is computed as follows

$$T_B(t) = \frac{1}{\Delta} \times \int_{t-\Delta}^{t} q(i) \times \frac{P_{avg}}{C} \times di \tag{9}$$

where C denotes the ISL capacity.

IV. PERFORMANCE EVALUATION

A. Simulation Setup

In this section, we evaluate the performance of the ELB scheme using the Network Simulator (NS) [10]. We consider the Iridium constellation. The constellation is formed of 66 satellites evenly and uniformly distributed over six orbits. Each satellite maintains four ISLs with its neighboring satellites. Uplinks, downlinks, and ISLs are each given a capacity equal to $25Mbps$ ($C = 25Mbps$). Their delays are set to $20ms$ ($d = 20ms$). With no specific purpose in mind, the average packet size is set to $1KB$ ($P_{avg} = 1KB$). Drop-Tail based buffers, of lengths equal to 200 packets, are used ($Q_l = 200\ pkts$).

For traffic generation, we consider 600 non-persistent On-Off flows. The On/Off periods of the connections are derived from a Pareto distribution with a shape equal to 1.2. The average burst time and the average idle time are set to $200ms$. The source and destination end-terminals are dispersed all over the Earth, divided into six continental regions, following a distribution identical to the traffic distribution used in [11] [12]. The sources send data at constant rates from within the range of $0.8Mbps$ to $1.5Mbps$.

In the performance evaluation, we use the Dijkstra's Shortest Path (DSP) algorithm and CEMR as comparison terms.

While the ELB scheme can be implemented over any routing protocol, we consider two implementations of the scheme; one over DSP and the other over CEMR. Our implementation of CEMR is based on the scheme description in [9]. Similarly to the paper, the routing cost metrics of CEMR and ELB are updated every $1s$ interval of time ($\Delta = 1s$). The performance of the schemes is evaluated in terms of the achieved total throughput and the experienced total packet drops. To investigate how well traffic is distributed over the entire constellation, the following traffic distribution index is used.

$$f = \frac{(\sum_{i=1}^{n} x_i)^2}{n \sum_{i=1}^{n} x_i^2} \quad (10)$$

where n is the number of ISLs and x_i denotes the actual number of packets that traversed the i^{th} ISL. This index ranges from zero to one and indicates how well the traffic is distributed over the constellation. Simulations are all run for $60s$. In the conducted simulations, satellites monitor their current queue occupancy in a real time fashion ($\delta = 1ms$). Finally, unless otherwise specified, the desired time for a satellite to reside in the Free state after a transition to the Busy state θ is set to $200ms$ (e.g. deliberately set to ten times the ISL delay).

B. Simulation Results

To investigate the abilities of the ELB scheme in supporting QoS, we evaluate its performance in terms of the achieved throughput and the total packet drops experienced by the simulated 600 connections. Fig. 1 graphs the total number of packet drops experienced by all the connections during the entire simulation course and that is for different sending rates of the connections. For all the considered bit rates, the implementation of ELB over CEMR shows the best performance as it achieves the lowest packet drop rate. Note also that even the implementation of ELB over DSP avoids more packet drops than the other two routing protocols, DSP and CEMR. This indicates an important feature of the ELB scheme in avoiding packet drops by alleviating congestion at satellites. The good performance of the ELB scheme in avoiding packet drops is manifested also in terms of the high throughput achieved by the ELB scheme. Fig. 2 shows that implementing ELB over CEMR and DSP leads to a remarkable increase in the total achieved throughput compared to the other two schemes, DSP and CEMR.

The ELB scheme results also in a more balanced distribution of traffic over the entire constellation. To illustrate the idea at hand, we plot the traffic distribution index for different values of sending rates in Fig. 3. The figure indicates that the implementation of ELB over DSP significantly outperforms the Dijkstra algorithm. This performance is attributable to the fact that the DSP algorithm bases its routing strategy on only finding paths with the shortest delay. Data is then transmitted over single paths during the entire transmission time. On the other hand, the ELB scheme searches for alternative paths when a satellite is about to get congested. Data is then transmitted over multiple less congested paths. This operation leads intuitively to a better and more efficient distribution of traffic among the constellation links. Compared to the CEMR scheme, whilst the implementation of ELB over CEMR exhibits a better distribution of traffic, the improvement is minimal. The underlying reason beneath this performance consists in the fact the CEMR uses also multiple disjoint paths for transmitting data. This fact makes different links involved in the transmission of data and yields hence a relatively wide distribution of traffic over the entire constellation.

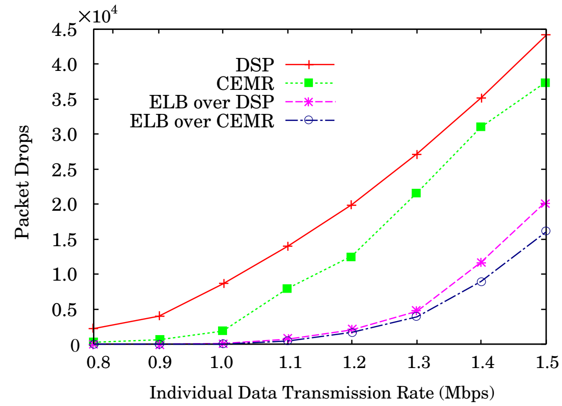

Fig. 1. Packet drops for different individual sending rates

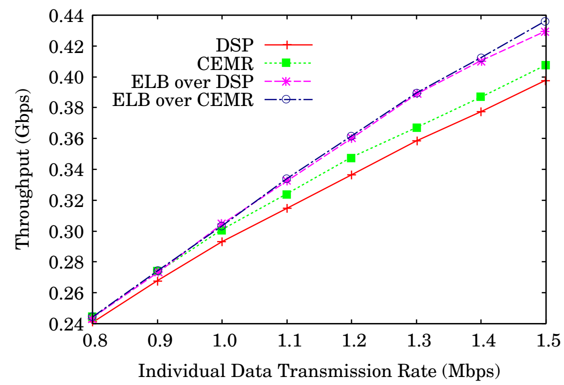

Fig. 2. Total throughput for different individual sending rates

In the ELB scheme, packets belonging to the same flow may have to traverse different paths with different propagation delays. In such case, out-of-order packets and delay jitter may occur. For connectionless protocols (e.g. User Datagram Protocol - UDP), this issue can be easily resolved by buffering capabilities. Indeed, a small buffer at end terminals can ensure coherent reception, remove the jitter added by the network, and recover the original timing relationships between the transmitted data. For applications based on connection-oriented protocols (e.g. Transmission Control Protocol - TCP), a certain level of tolerance in packet disorder or delay jitter should be guaranteed. On the other hand, the ELB scheme sometimes forces packets to traverse more hops than in case of traditional routing algorithms. The ELB scheme may be thus thought of as a scheme that guarantees high throughput and low packet drops, but at the price of higher delays. However, considering the significant queuing delays that may result from congesting satellites, the additional hops that should be traversed by packets can be justified. Furthermore, it would be

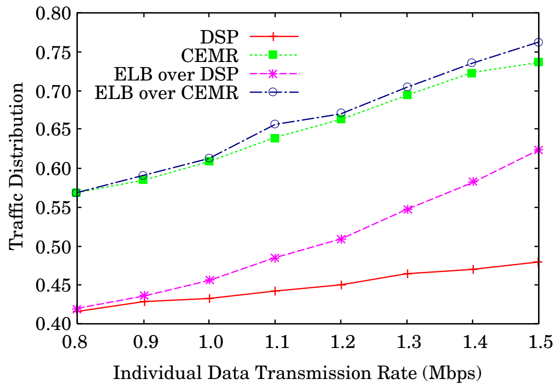

Fig. 3. Traffic distribution index for different individual sending rates

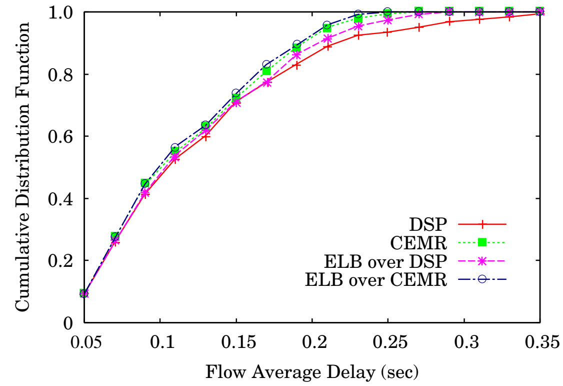

Fig. 4. Cumulative distribution function of flows' average delay (flows' transmission rate = $1.0 Mbps$)

more beneficial for a system to have some packets experience some delay than having them discarded completely, mainly in light of the long time required for their retransmission in environments known for their long propagation delays.

To show the idea with more clarity, we plot the Cumulative Distribution Function (CDF) of the average delay of flows in Fig. 4. Whilst the figure indicates the results when the sending rate of flows is set to $1Mbps$, identical plots were obtained in the case of other sending rates. The figure demonstrates that while some individual flows may experience longer delays than in case of traditional routing schemes, the aggregate performance of the ELB scheme in terms of delay is the best as the CDF plots of "ELB over DSP" and "ELB over CEMR" are higher than the plots of DSP and CEMR, respectively. As previously mentioned, the underlying reason beneath this performance consists in the abilities of ELB to alleviate congestion, to accordingly reduce the queue occupancies of satellites, and to ultimately reduce the queuing delays.

V. Conclusion

In this paper, we proposed an Explicit Load Balancing (ELB) routing strategy that cooperatively involves neighboring satellites and guarantees an efficient distribution of traffic over multi-hop NGEO satellite constellations. In the proposed routing protocol, neighboring satellites explicitly exchange information on their current congestion status. Satellites with queue occupancies exceeding a pre-determined threshold are assumed to get congested soon. To avoid such congestion, along with the resulting packet drops, satellites, running under heavy loads, request their neighboring satellites to reduce their data forwarding rates. In response, the neighboring satellites search for alternative less-congested paths and communicate a predetermined portion of their data via these paths. The working of the proposed routing scheme is based on three metrics. A dynamic setting of these parameters is proposed based on easy-to-implement equations.

A set of simulations is conducted to evaluate the performance of the ELB scheme. Two implementations are considered; one over a recently proposed scheme, CEMR, and the other over the most widely used Dijkstra algorithm. The obtained simulation results elucidate the better performance of the ELB scheme in avoiding congestion, reducing queue lengths, lowering packet drops, and increasing the total throughput while maintaining a more balanced distribution of traffic over the constellation.

References

[1] T. Taleb, N. Kato, and Y. Nemoto, "Recent Trends in IP/NGEO Satellite Communication Systems: Transport, Routing, and Mobility Management", IEEE Wireless Communications Magazine, Vol. 12, No. 5, Oct. 2005, pp. 63-69.

[2] T. Taleb, A. Jamalipour, N. Kato, Y. Nemoto, "IP Traffic Load Distribution in NGEO Broadband Satellite Networks", in Proc. 20^{th} Int. Symposium on Computer & Information Sciences, Istanbul, Turkey, Oct. 2005.

[3] A. Molinaro, F. De Rango, S. Marano, M. Tropea, "A scalable framework for end-to-end QoS assurance in IP-oriented terrestrial-GEO satellite networks", IEEE Commun. Mag., Vol. 43, No. 4, Apr. 2005, pp. 130-137.

[4] J. Chen, A. Jamalipour, "An adaptive path routing scheme for satellite IP networks", Int. J. Commun. Sys., Vol. 16, No. 1, Feb. 2003, pp. 5-21.

[5] A. Svigelj, M. Mohorcic, G. Kandus, A. Kos, M. Pustišek, J. Bešter, "Routing in ISL networks considering empirical IP traffic", IEEE J. Select. Areas in Commun., Vol. 22, No. 2, Feb. 2004, pp. 261-272.

[6] A. Donner, M. Berioli, M. Werner, "MPLS-based satellite constellation networks", IEEE J. Select. Areas in Commun., Vol. 22, No. 3, Apr. 2004, pp. 438-448.

[7] H. Uzunalioglu, "Probabilistic routing protocol for low earth orbit satellite networks", in Proc. of IEEE Int. Conf. Commun. 1998. Atlanta, GA, USA, Jun. 1998.

[8] R. Kucukates, C. Ersoy, "High performance routing in a LEO satellite network", in Proc. 8^{th} IEEE Int. Sym. on Comp. Commun. (ISCC 2003), Washington, DC, USA, Jun. 2003.

[9] B. Jianjun, L. Xicheng, L. Zexin, P. Wei, "Compact explicit multi-path routing for LEO satellite networks", in Proc. of 2005 IEEE Workshop on High Performance Switching and Routing, Hong Kong, P.R. China, May. 2005.

[10] UCB/LBNL/VINT: Network Simulator - ns (version 2). http://www.isi.edu/nsnam/ns/

[11] M. Mohorcic, M. Werner, A. Svigelj, and G. Kandus, "Adaptive Routing for Packet-Oriented Intersatellite Link Networks: Performance in Various Traffic Scenarios", IEEE Trans. Wireless Commun., Vol. 1, No. 4, Oct. 2002, pp. 808-818.

[12] J. Hutcheson and M. Laurin, "Network Flexibility of the Iridium Global Mobile Satellite System", in Proc. 4^{th} Int. Mobile Satellite Conference, Ottawa, Canada, Jun. 1995.

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.