Distributed Load-Aware Routing in LEO Satellite Networks

Evangelos Papapetrou
University of Ioannina
Dourouti Campus, 45110, Ioannina, Greece

Fotini-Niovi Pavlidou
Aristotle University of Thessaloniki
Panepistimioupoli, 54124, Thessaloniki, Greece

***Abstract*—In this paper, we propose a lightweight distributed routing algorithm, called Distributed Load-Aware Routing (DLAR). The proposed protocol adopts a distributed approach to handle the complexity of the satellite system and at the same time provides a hop-by-hop mechanism for splitting traffic load in order to alleviate the problem of congestion that occurs near polar regions. The performance of DLAR is assessed through extensive simulations and compared to the performance of centralized routing schemes proposed so far in the literature. Simulation results document and confirm the positive characteristics of the proposed protocol.**

I. Introduction

Satellite systems have always been the prevailing candidate for realizing the vision of anywhere, anytime connectivity on account of their inherent broadcasting nature and large scale coverage. In the past, Low Earth Orbit (LEO) satellite systems [1], have drawn the attention of the telecommunications community since they provide low propagation delay and higher throughput compared to geostationary links. Despite their characteristics, the success of LEO systems in the past decade was limited as a result of their failure to compete with terrestrial infrastructure. However, over the last years the interest in LEO systems has been renewed on the basis that satellite systems should augment the operation of terrestrial networks rather than compete with them [2],[3]. Envisaged application scenarios include the use of LEO systems to integrate world-wide networking infrastructures, provide backbone connectivity to terrestrial systems, etc. Most LEO satellite networks make use of inter-satellite links (ISLs) to provide low-delay connectivity between adjacent satellites. Routing data from source to destination through a LEO network constitutes a daunting challenge. Its characteristics such as its size, the discontinuous operation of ISLs, its high connectivity and its unbalanced loading, pose strict requirements to the design of efficient routing algorithms.

Most of the proposed routing approaches so far involve a centralized scheme for calculating paths for all origin/destination pairs. The common approach in the literature as yet has been to use propagation as well as queuing delay as routing metrics and implement a periodic procedure for collecting such information from each ISL. In [4], [5], [6] centralized routing schemes were proposed, which rely on the Dijkstra shortest path algorithm to compute the optimal path for any pair of satellites. In the case of traffic adaptive routing, frequent network state variations render the induced overhead unacceptable and at the same time the routing accuracy is minimized due to high propagation delays. Therefore, the approach taken in [7], is to implement a stateless distributed routing protocol. The proposed scheme is overhead-free, however, it routes packets based on propagation delay and diverts them to backup paths only when congestion occurs. This reduces the ability of the algorithm to adapt to traffic. Furthermore, the proposed algorithm is applicable only to specific systems (systems with an inclination of 90^o and zero phase shift).

In this paper we propose a *Distributed Load-Aware Routing protocol* (DLAR) for LEO satellite networks that employ ISLs. The protocol adopts a distributed approach for routing in satellite networks. DLAR capitalizes on the system deterministic dynamics to determine a set of available paths to the destination. At the same time, DLAR evaluates local traffic conditions in each hop and splits traffic to the set of paths that lead to the destination in order to minimize delay. The innovation of DLAR is that it manages to implement a traffic adaptive mechanism in a distributed fashion which induces no overhead to the network. Furthermore, DLAR provides a solution even for polar LEO systems that use phase shifting for the placement of satellites in adjacent orbital planes. The performance of the proposed protocol is evaluated through simulation studies, illustrating the positive characteristics of the proposed protocol.

The remainder of the paper is structured as follows. In Section II the system model is presented and the problem formulation is provided. Then, the DLAR protocol is delineated in section III, highlighting its main differences with other routing schemes. Section IV is devoted to the description of the simulation model used for the assessment of the new protocol. In section V simulation results are presented and discussed, while concluding remarks are drawn in section VI.

II. System Model and Problem Formulation

A. System Model

In this study we focus on near polar (or *Walker star*) constellations that employ ISLs (see fig.1a). Polar constellations present several attracting characteristics such as global coverage, deployment of ISLs, etc. In coherence to other studies, we will study the most usual type of constellations in which each satellite is assigned four ISLs: two intra-plane ISLs (namely, links to the adjacent satellites in the same orbital

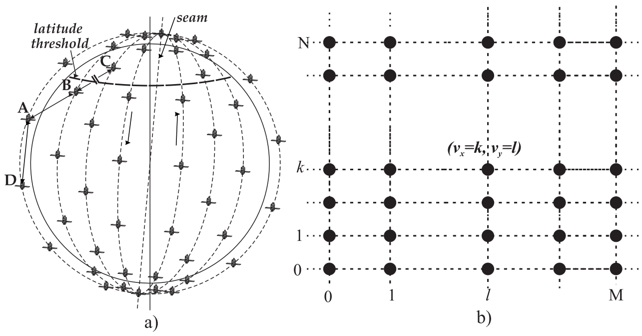

Fig. 1: a)Typical architecture of Walker delta LEO constellation, b)Example of resulting graph of size NM

plane, e.g. link AD in fig.1a) and two inter-plane ISLs (that is, links to the neighboring satellites in the right-hand and left-hand orbital planes, e.g. link AB in fig.1a). While intra-plane ISLs are maintained for the whole satellite period, inter-plane ISLs switch-off (e.g. link BC in fig. 1a) as satellites come close to polar regions (defined by a latitude threshold), due to adverse pointing and tracking conditions. Inter-plane ISLs are reestablished when satellites move to lower latitudes. The network formed by satellites and ISLs is called space segment. This network, on its abstraction level, can be modelled as a graph $G(V, E)$, comprising of a set of nodes V and a set of edges E (figure 1b). It is clear that the size of V is $|V| = NM$, where M is the number of the orbital planes that the system is comprised of and N is the number of satellites per plane. By numbering the orbital planes and the satellites within a plane, we can define a pair of numbers (v_x, v_y), called virtual coordinates, that uniquely identifies a satellite. The concept of virtual coordinates is well-known in the related literature [4]. Clearly, $v_x \in [0, N)$ identifies the position of a satellite within an orbital plane, while $v_y \in [0, M)$ identifies the orbital plane. It must be noted that although in the graph $G(V, E)$, presented in fig. 1b, the relative satellite positions are independent of their actual location over the Earth's surface, the graph itself is not. The set E and therefore the routing decisions depend on the actual location of the communicating satellites with respect to the Earth's surface.

B. Problem Formulation

Based on the system model described in the previous section, an efficient method to route a packet from a source satellite $v^s = (v_x^s, v_y^s)$ to a destination satellite $v^d = (v_x^d, v_y^d)$, is to use a hop-by-hop approach. Possible next hops are decided so that the virtual distance from the destination node reduces. For example, suppose that a packet is in an intermediate node v^i, on its way to the destination node v^d. A neighboring satellite v^j is a possible next hop only if either $|v_x^d - v_x^j| < |v_x^d - v_x^i|$ or $|v_y^d - v_y^j| < |v_y^d - v_y^i|$. This process guarantees that the chosen paths are in the set of the minimum hop paths. This method guarantees that the packet will reach its destination without coming to a deadlock provided that the inclination of the satellite system is 90^o and the phase shift between satellites in adjacent planes is zero (see fig.2-dotted lines indicate ISLs). In that case, deadlocks may be avoided by not forwarding the packet to an intermediate satellite v^i that resides in the polar area while $|v_x^d - v_x^i| > 0$ and $|v_y^d - v_y^i|$ being smaller than the polar region size in hops. In other words, the packet must not enter the polar area while there is still a distance to be covered in the x-direction and its exit from the polar region is not secured. The described method for avoiding deadlocks is simple since it requires only the knowledge of whether the next hop satellite lies within the polar area. However, this method does not apply to realistic satellite systems with inclination less than 90^o and/or non-zero phase shift between satellites in adjacent planes. An example is depicted in fig.3. In that case, even if the packet is forwarded to an intermediate node (node I in fig.3) that is out of the polar region, the packet finally reaches a deadlock, since there is no available interplane ISL. The reason of this deficiency is the fact that the rings formed by satellites connected by interplane ISLs are not parallel to the polar region boundary. Therefore, paths consisting of hops only in the x-direction are not guaranteed to exist if the first satellite in the path is outside the polar region.

Another deficiency of the described procedure is that the actual choice of the next hop is made by evaluating the propagation delay of the candidate ISLs. This results in routing packets through specific ISLs for large periods of time. Therefore congestion may occur. A method for avoiding congestion was proposed in [7], however queueing delay may be considerable high before congestion occurs. Therefore, an efficient mechanism that takes into advantage the loading of ISLs must be devised.

III. Distributed Load-Aware Routing Protocol

Distributed Load-Aware Routing (DLAR) Protocol capitalizes on the system deterministic dynamics to tackle the problem of reaching a deadlock in a realistic satellite system. At the same time it takes into account the loading of ISLs to route packet efficiently. The protocol consists of two phases. In the first phase, all possible next hops are determined while in the second the most suitable one is chosen. Before describing the proposed protocol in detail, we provide some useful definitions.

A. Definitions

Definition 1: **Horizontal Path to Destination**. A path originating at the current satellite v^{cur} that involves only hops in the x-direction and at the same time its size in hops equals $|v_x^{cur} - v_x^d|$, where v^d is the destination satellite. More formally, the horizontal path from v^{cur} to destination v^d is denoted by $P_{HP}^{(cur,d)} = \{v^{cur}, \ldots, v^d\}$ such that $v_y^k - v_y^m = 0$, $\forall v^k, v^m \in P_{HP}^d$.

Definition 2: **Next Horizontal Path to Destination**. The horizontal path to destination originating from the satellite next to the current one (in the same orbital plane towards the destination). More formally, $P_{NHP}^d = P_{HP}^{(i,d)}$ such that $|v_y^d - v_y^{cur}| - |v_y^d - v_y^i| = 1$ and $v_x^{cur} - v_x^i = 0$.

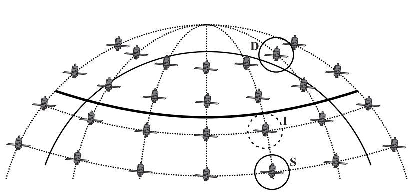

Fig. 2: Communication example in a LEO system with inclination= 90^o and phase shift= 0.

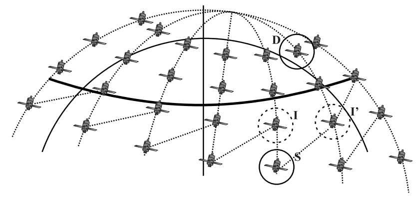

Fig. 3: Communication example in a LEO system with inclination$\neq 90^o$ and phase shift$\neq 0$.

Definition 3: **Last Horizontal Path to Destination**. A path $P^d_{LHP} = P^{(i,d)}_{HP}$, originating from a satellite v^i, in the orbital plane of the current satellite, such that $|v^d_y - v^i_y| = 0$ and $v^{cur}_x - v^i_x = 0$.

Definition 4: **Vertical Path**. A path $P_{VP} = \{v^i, \ldots, v^j\}$ so that $v^k_x - v^m_x = 0$, $\forall v^k, v^m \in P_{VP}$.

B. Determination of forwarding possibilities

DLAR operates on a distributed hop-by-hop basis to determine which of the available neighboring satellites may be used in the source node or in any intermediate node for forwarding a packet. The basic aim of this procedure is to eliminate the possibility of a packet reaching a deadlock, i.e. not reaching the destination due to lack of path availability. In the first step the current satellite v^{cur} elects a neighboring satellite v^n as a possible next hop if $|v^d_x - v^n_x| < |v^d_x - v^{cur}_x|$ or $|v^d_y - v^n_y| < |v^d_y - v^{cur}_y|$, i.e. a satellite that reduces the virtual distance from the destination. In general, more than one satellites may fulfill one of these conditions. However, as discussed in Section II-B, in real life satellite systems, some of the possible next hop satellites, determined by the previous procedure, should be eliminated in order to avoid deadlocks. To this end, DLAR follows the algorithm depicted in fig.4. DLAR starts with checking whether a horizontal path is needed to reach the destination, i.e., if $|v^d_x - v^{cur}_x| \neq 0$. If a horizontal path is not need, the algorithm terminates since vertical paths are always available. In the case that a horizontal path is required, the algorithm examines whether a vertical path is required ($|v^d_y - v^{cur}_y| \neq 0$). If not, this means that the only possibility is to use a horizontal path. Therefore, DLAR takes immediate action to determine whether the horizontal path to destination (P^d_{HP}) is available. This can be done by taking advantage of the deterministic satellite dynamics. In case P^d_{HP} is available, no further action is taken. If it is not, this means that the packet is in a satellite that lies in the polar region (probably because this satellite is the source satellite), so DLAR eliminates all possible next hops except the one that leads out of the polar region. This is done by examining the latitude of the current satellite and whether it is ascending or descending. Coming back to the case that both a horizontal and a vertical path are required to reach the destination, DLAR first checks the availability of the next horizontal path to the destination (P^d_{NHP}). This is a proactive action in order to determine whether the packet approaches a polar area. If P^d_{NHP} is not available, forwarding the packet in the vertical path will result in it entering the polar area. Therefore, before allowing this, DLAR checks the availability of P^d_{HP} and P^d_{LHP}. If both P^d_{HP} and P^d_{LHP} are not available, this means that the packet is already in the polar region and is not foreseen to exit. Therefore, DLAR chooses the link that maintains the initial direction of the packet although this may temporarily increase the virtual distance to the destination. In the case that only P^d_{LHP} is available no action is required since the packet will normally exit the polar region. In the case that P^d_{HP} is available while P^d_{LHP} is not, DLAR eliminates links that lead to vertical paths and the packet is forwarded in the horizontal path.

From the previous description it is clear that DLAR avoids the formation of deadlocks, since packets that are destined to a polar region are entering it only through the appropriate vertical path. Furthermore, packets originating from a satellite within the polar region, are led out by steadily following the appropriate vertical path. Moreover, the loop-free property is provided by allowing only hops that reduce the virtual distance to the destination. This principle is only overridden in order to lead a packet out of a polar region. However, in this case loop formation is not possible since orbital planes do not communicate in polar regions. Finally, the calculations required for determining the availability of horizontal paths are based on the deterministic dynamics of the system. Therefore, such calculations may be performed offline.

C. Next Hop Election

By the procedure described in Section III-B, DLAR manages to determine the set of possible next hops and at the same time guarantees that there is no possibility for a packet to come to a deadlock. The identified next hop satellites correspond to paths with minimum hops count. DLAR takes advantage of the system's capacity for providing many paths with similar delays To complete the forwarding process, DLAR should elect the next hop. To this end, all possible next hops are classified based on the propagation delay of the ISL connecting them to the current satellite. One approach is to choose the next hop associated with the minimum propagation delay (this approach is called DLAR-PR). A second approach is to perform the classification based on the sum of propagation and queuing delay (called DLAR-TR), which represents a traffic adaptive approach. DLAR-TR aims at minimizing delay by routing packets to the less loaded ISL. In this way it manages to split traffic and result in a more uniform loading of the satellite system. The measurement of queueing delay in each ISL may be performed by the Layer 2 mechanisms.

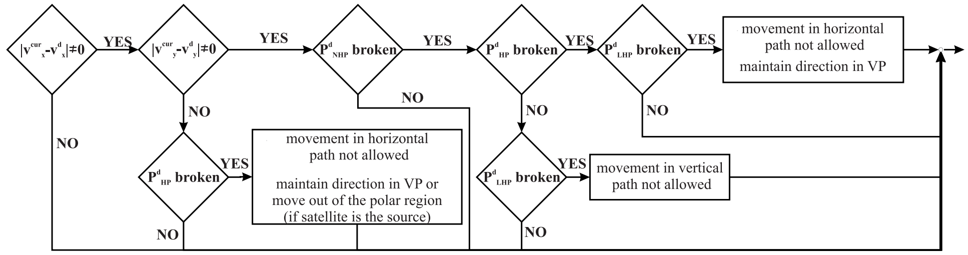

Fig. 4: Flow chart of Distributed Load-Aware Routing

IV. Simulation Framework

To evaluate the performance of the proposed protocol, simulation experiments were conducted to compare its performance with that of two centralized routing schemes proposed so far in the literature, in realistic satellite systems (inclination$\neq 90\%$ and phase shift $\neq 0$). The simulation tool used was coded in the platform of the detailed simulation model of Ns2. Furthermore, the existing code regarding centralized routing was modified to facilitate also a traffic adaptive centralized scheme. Two flavors of DLAR were evaluated, namely DLAR-PR and DLAR-TR. The evaluated protocols were tested in an Iridium-like constellation, where ISLs are switched off when satellites cross the polar regions defined by a latitude threshold ($\pm 60^o$). Moreover 200 terminals were distributed over the six continents according to the *hot spot* scenario described in [8], which is based on the distribution of web servers. An exponential ON/OFF traffic generator is attached to each one of them. The bitrate of each generator was set up to 1200 kbps which produced a heavily loaded system. The simulation parameters are given in Table I. In order to enhance the performance of the periodic centralized schemes, the update interval of routing tables was set to $10\,sec$. In addition to the periodic computation, the calculation of the shortest paths is also triggered whenever a change in the ISL topology occurs.

TABLE I: Simulation parameters

Traffic generator's parameters			
Packet size	1500 bytes	"On" period	0.3 sec
"On"periods bitrate	200 kb/s - 1200 kb/s	"Off" period	0.9 sec
System's parameters			
Up/downlink bwd	15 Mb/s	ISL bwd	10 Mb/s
ISL LL queue size	500 packets	Sim. duration	6050 sec

The centralized routing schemes that we tested were proposed in [5] and [6] are based on the Dijkstra shortest path algorithm and are suitable for realistic satellite systems (inclination$\neq 90\%$ and phase shift $\neq 0$). The link-cost function used for the first protocol is based only on propagation delay (Central-PR), while the second (Central-TR), being a traffic-adaptive scheme, uses as link cost the summation of propagation and queuing delay [5], [6]. The queueing delay is estimated by means of an average value over the update interval. As performance metric we used the mean end-to-end delay, the delivery ratio and the delay jitter. While the first two metrics are the most appropriate for evaluating a routing algorithm, the later represents a useful metric for QoS communications and at the same time provides useful insight to the algorithm operation. Furthermore, comments are provided about the overhead incurred by the compare algorithms. Finally, it should also be noted that the presented results represent average values over 10 independent simulation runs. That number of runs provided 95% confidence intervals of $\pm 5\%$ in the worst case.

V. Protocol Evaluation

The first objective of the performance evaluation is to illustrate the ability of DLAR to correctly deliver packets to their destination. Figure 5(a) presents the delivery ratio for the four evaluated schemes with respect to each terminal's bitrate. DLAR-TR clearly outperforms the other schemes for high bitrates (more than 800 kbps), while for medium and low bitrates the four algorithms are loss-free. For high bitrates, the network becomes congested and queued packets encounter a higher probability of being dropped due to buffer overflow. DLAR-TR minimizes such incidents by efficiently splitting traffic in each hop to multiple paths. Centralized schemes, on the other hand, are more vulnerable to congestion situations since their periodic operation limits their ability to monitor the network state. Therefore, their performance is inferior even to that of DLAR-PR. DLAR-PR performs better than Central-TR because of its ability to capture network state in each hop. It must be noted that the actual delivery ratio difference between all four schemes is not impressive, however it should be noted that the quantitative difference depends on the actual size of the buffers in each satellite. This size introduces a tradeoff in the performance of delivery ratio and end-to-end delay. However, the qualitative difference in the performance of the four schemes is more important.

The comments regarding the performance of the four schemes in terms of delivery ratio, are substantiated by fig. 5(b), where the end-to-end delay for delivering data to their destination is presented with respect to the terminal's bitrate. The high delays when high bitrates are offered, provide evidence of congestion in the network. DLAR-TR by locally evaluating the congestion of each ISL, manages to efficiently split traffic. Furthermore, the capacity of DLAR-TR to route packets efficiently is more intense for high data rates, where its difference from the performances of the other schemes is increased. Actually, DLAR-TR manages a reduction of 51%-53.7% for the bitrate of 1200 kbps while at the same time

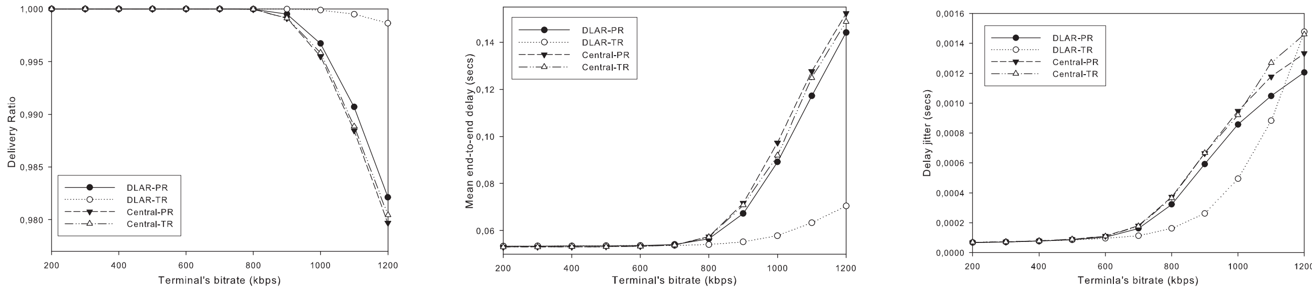

(a) (b) (c)

Fig. 5: a) Delivery ratio, b) Mean delay and c) Delay jitter vs terminal's bitrate

manages to deliver more packets (see fig.5(a)). As far as the centralized schemes are concerned, it is evident that fail to efficiently handle congestion. As a result, long queues are formed and result in increased delays. Surprisingly, Central-TR performs similar to Central-PR although the former utilizes in its link cost the queueing delay. This result can be attributed to the phenomenon of oscillation, which refers to sudden traffic shifts towards the best calculated path in each period. As a result, newly discovered paths become instantly congested. Finally, DLAR-PR presents slightly better performance compered to the centralized schemes. This is a result of the algorithm's strategy to reevaluate locally the network state in each hop.

Fig. 5(c) presents another interesting metric for time-sensitive applications, i.e., the delay jitter. DLAR-TR again performs better than the other schemes in almost the entire range of bitrates. For low and medium data rates, DLAR-TR performes better due to its ability to split traffic, therefore producing stable loading conditions in the ISLs. However, for high bitrates the performance of DLAR-TR degrades. In such congestion conditions, queueing delay in an ISL becomes highly variable. Since, DLAR-TR directly evaluates queuing delay in each ISL, it is more susceptible to its variations. On the other hand, although Central-TR uses an average value of queueing delay, it does not manage to avoid oscillations. At the same time, the schemes that use only propagation delay are more resilient to delay jitter since their routing decision changes slower. Last but not least, it must be noted that the proposed protocol does not incur any routing overhead. This is because on one hand propagation delay can be calculated based on the deterministic system dynamics and on the other, DLAR-TR performs queuing delay measurements only locally by using the Layer-2 capabilities. On the contrary, periodic centralized schemes required the collection of routing information from the entire network and the distribution of routing decisions to each satellite. Both mechanisms involve significant overhead that depends on the network size and on the period of execution. This overhead is difficult to be calculated since this depends on the actual implementation details of collecting and distributing routing information such as where the collection center will be located, what is the mechanism for distributing routing information in the network, etc.

VI. Conclusions

In this paper we proposed and evaluated the performance of a distributed load-aware (DLAR) routing protocol for LEO satellite networks. The proposed protocol was compared to centralized routing schemes, proposed in the literature. Ample simulation results corroborated the superiority of DLAR over those schemes. Specifically, DLAR was shown to attain much lower end-to-end delay and higher delivery ratio while at the same time does not incur any routing overhead. This fact renders it an excellent choice for future LEO satellite networks.

VII. Acknowledgement

This work was carried out in the framework of the EU funded Network of Excellence "SatNEx II" (contract No. IST-027393).

References

[1] A. Jamalipour, *Low Earth Orbital Satellites for Personal Communication Networks*. Boston-London: Artech House, 1998.

[2] B. Evans, M. Werner, E. Lutz, M. Bousquet, G. E. Corazza, G. Maral, R. Rumeau, and E. Ferro, "Integration of Satellite and Terrestrial Systems in Future Multimedia Communications," *IEEE Wireless Communications*, vol. 12, no. 5, pp. 72–80, October 2005.

[3] T. Taleb, N. Kato, and Y. Nemoto, "Recent Trends in IP/NGEO Satellite Communication Systems: Transport, Routing, and Mobility Management Concerns," *IEEE Wireless Communications*, vol. 12, no. 5, pp. 63–69, October 2005.

[4] M. Werner, "A Dynamic Routing Concept for ATM-Based Satellite Personal Communication Networks," *IEEE J. Select. Areas Commun.*, vol. 15, no. 8, pp. 1636–1648, October 1997.

[5] M. Mohorčič, A. Švigelj, G. Kandus, and M. Werner, "Performance evaluation of adaptive routing algorithms in packet-switched intersatellite link networks," *Intern. J. of Satellite Commun.*, vol. 20, pp. 97–120, 2002.

[6] M. Mohorčič, M. Werner, A. Švigelj, and G. Kandus, "Adaptive Routing for Packet-Oriented Intersatellite Link Networks: Performance in Various Traffic Scenarios," *IEEE Trans. Wireless Commun.*, vol. 1, no. 4, pp. 808–818, October 2002.

[7] E. Ekici, I. F. Akyildiz, and M. D. Bender, "A Distributed Routing Algorithm for Datagram Traffic in LEO Satellite Networks," *IEEE/ACM Trans. Networking*, vol. 9, no. 2, pp. 137–147, April 2001.

[8] M. Mohorčič, A. Švigelj, G. Kandus, Y. F. Hu, and R. E. Sheriff, "Demographically weighted traffic flow models for adaptive routing in packet-switched non-geostationary satellite meshed networks," *Computer Networks*, vol. 43, pp. 113–131, 2003.