

On-Demand Routing in LEO Satellite Systems

Stylios Karapantazis
Aristotle University of Thessaloniki
Panepistimioupoli, 54124
Thessaloniki, Greece
Email: skarap@auth.gr

Evangelos Papapetrou
University of Ioannina
Dourouti Campus, 45110
Ioannina, Greece
Email: epap@cs.uoi.gr

Fotini-Niovi Pavlidou
Aristotle University of Thessaloniki
Panepistimioupoli, 54124
Thessaloniki, Greece
Email: niovi@auth.gr

Abstract—In this paper, a location-assisted on-demand routing (LAOR) protocol for low earth orbit (LEO) satellite systems is proposed and evaluated. This protocol can be viewed as a variant of the well-known ad-hoc on-demand distance vector (AODV) routing protocol, tailored to the requirements imposed by the characteristics of LEO satellite systems' topology. The LAOR protocol is assessed for different link-cost metrics and compared to centralized routing protocols proposed in the literature thus far. Simulation studies further document and confirm the positive characteristics of the proposed protocol.

I. INTRODUCTION

The past decade witnessed a growing interest in non-geostationary (non-GEO) satellite systems that orbit the earth at altitudes between 500 km and 1500 km, which are termed low earth orbit (LEO) satellite systems, on account of the appealing features that these systems are endowed with, such as low propagation delay and global coverage [1]. Nevertheless, after the limited commercial success of the two operating narrowband LEO satellite networks, which mainly provide telephony service, the interest in non-GEO satellite systems has dwindled away. Notwithstanding this decline in interest, the current trend toward the migration to all IP-based service opens new opportunities to this kind of systems. In this context, LEO constellations can be employed to unify far-flung groups of people and provide high quality telecommunication services to areas where there is no substantial terrestrial infrastructure. Moreover, LEO satellite networks can play a pivotal role in the evolving telecommunications infrastructure, providing ubiquitous access to the Internet and supporting multimedia services at low transmission cost.

However, as in any wireless or wired telecommunications system, routing data from the source all the way to the destination constitutes a daunting challenge in LEO satellite networks as well. All the studies in this field focused on constellations with inter-satellite links (ISLs). Further, most of them take advantage of the deterministic dynamics of this network topology to divide the system period into a number of time intervals during which the topology of the network is fixed.

In [2], [3], [4], [5] centralized routing schemes were proposed, which relied on the Dijkstra shortest path algorithm to compute the optimal path for any pair of satellites. The study in [6] treated the issues of both topological design and routing in LEO satellite systems using a finite state automaton (FSA)

approach. A modified Flow Deviation algorithm was proposed in [7]. The main function of the Flow Deviation algorithm is to split traffic into many different paths. A distributed routing protocol that routes each IP packet independently based on propagation delay and induces no signaling overhead at all was assessed in [8].

The common approach in the literature as yet has been to periodically collect and evaluate propagation and queuing delays in each ISL. Nonetheless, albeit update intervals of about 30 secs (adopted in the literature thus far) provide satisfactory adaptation to propagation delay variations, when adaptivity to traffic related metrics (i.e., queuing delay) is required, periodic updates may result in high overhead. In this paper we propose a *location-assisted on-demand routing* (LAOR) protocol for LEO satellite IP networks that employ ISLs. The protocol introduces the well-known concept of on-demand routing to satellite networks. LAOR can be viewed as a variant of the *ad-hoc on-demand distance vector* (AODV) routing algorithm [9], tailored to the requirements imposed by the characteristics of LEO networks' topology. The prime aim of the LAOR protocol is to minimize end-to-end delay and delay jitter, while keeping signaling overhead to a minimum. Toward this end, the path discovery process is invoked independently for each individual connection. However, our algorithm capitalizes upon the deterministic characteristics of the LEO mesh architecture in order to restrict the number of satellites that are flooded with route request messages. The performance of the proposed protocol is assessed for different link-cost metrics and compared to the performance of proposed centralized routing protocols. Simulation studies further document and confirm the positive characteristics of the proposed protocol.

The remainder of the paper is structured as follows. The LAOR protocol is delineated in section II. Section III is devoted to the description of the simulation model used for the assessment of the new protocol. In section IV simulation results are presented and discussed, while concluding remarks are drawn in section V.

II. LAOR PROTOCOL DESCRIPTION

Hereinafter we consider a polar (*Walker star*) LEO constellation, wherein each satellite is assigned four ISLs; two intra-plane ISLs (that is, links to the adjacent satellites in the same orbital plane) and two inter-plane ISLs (namely to neighboring satellites in the right-hand and left-hand orbital planes). It

should be pointed out that intra-plane ISLs are permanent, whereas inter-plane ISLs are usually switched off as satellites pass over the poles due to adverse pointing, acquisition and tracking conditions. We also consider that there do not exist links that interconnect satellites in counter-rotating orbits, as is the case with the Iridium system [2].

On its abstraction level, the aforementioned network topology can be regarded as a mesh grid like the one depicted in Fig. 1. The fact that most of the proposed routing algorithms compute the shortest paths periodically, thus being unable to capture variations in network state in the meanwhile, motivated us to develop an on-demand routing protocol that will calculate independently the shortest path for each individual connection, thus making it responsive to the current network state. Such an approach, however, can induce significant signaling overhead. To overcome this effect, we propose the *location-assisted on-demand routing* (LAOR) protocol which takes advantage of the predictable network topology in order to diminish the number of satellites that will be involved in the path discovery process. The proposed algorithm makes use of some of the mechanisms that the AODV protocol employs. In particular, we have borrowed the concept of the *node's sequence number* from the AODV protocol which is used in the forward and the reverse path formation.

From now on we assume that each satellite maintains a counter, called the *node's sequence number* as well as a queue that contains all the packets waiting for a path to be set up, which is called *LAOR queue* hereinafter. Furthermore, we consider that each satellite should maintain a route entry *RT* for each destination of interest. Each route entry contains the following information:

- Destination satellite:** it identifies the destination satellite.
- Destination sequence number:** the last heard sequence number of the destination satellite.
- Next hop:** the satellite that a packet should be sent to.
- Path expiration time:** the time instant at which the route entry, and thus the path, will become invalid.
- Timestamp:** the time instant at which the route entry was recorded.
- Path cost:** the cost of the path.

In addition to the aforementioned route entry, which is used to forward data packets, each satellite also maintains a route entry $R_{eq}T$ where it stores information regarding *route requests* (RREQs) originated by that satellite. The latter kind of entry comprises four fields:

- Source satellite:** the satellite that sent the RREQ.
- Source sequence number:** the last heard sequence number of the source satellite.
- Next hop:** the satellite that a route reply message should be sent to.
- Timestamp:** the time instant at which the route entry was recorded.
- Path cost:** the cost of the path from the source to this node.

The rationale behind the use of this kind of entry is that the shortest forward and reverse paths between two satellites are not always the same when the path cost takes account of queuing delay, since the loading of each link is not the same in both directions.

For the sake of presentation, the LAOR algorithm is broken down into three processes:

- *Restricting network topology process*




Fig. 1: A mesh grid and the concept of restricted query area

- *Path discovery process*
- *Route entry management*

Before setting out the delineation of the algorithm, we should lay out the notations that will be used in the rest of the paper. Supposing that the system consists of N orbital planes that each one comprises M satellites, each satellite is uniquely defined by the pair of virtual coordinates (x, y) , where $0 \leq y \leq N - 1$ and $0 \leq x \leq M - 1$.

A. Restricting Network Topology Process

Let us assume that a ground terminal served by the j^{th} satellite in the i^{th} orbital plane (hereafter referred to as the source satellite) wants to communicate with a terminal that is covered by the l^{th} satellite in the k^{th} orbital plane (from now on referred to as the destination satellite) as illustrated in Fig. 1. As soon as the up-down links (UDLs) are set up and the first packet arrives at the source satellite, the latter places the packet in its *LAOR queue* and thereupon invokes the *restricting network topology process*.

Assuming that the destination satellite is known at that time, the source satellite compares its coordinates to the ones of the destination satellite. The outcome of this comparison will allow the source satellite to estimate whether *Route requests* (RREQ) must be sent on its right or on its left. However, the signaling overhead can still be considerable. To further reduce overhead the proposed algorithm employs another mechanism, which is actually the workhorse of the LAOR protocol. This mechanism is the adaption of the area to which RREQs will be sent according to the virtual coordinates of the source and destination satellites, hence the name *location-assisted on-demand routing* that was given to the protocol.

Let us turn again to the example depicted in Fig. 1 and let x_{min} , x_{max} , y_{min} and y_{max} denote the boundaries of the *restricted query area*. The source satellite calculates them as follows:

$$\begin{aligned} x_{min} &= \min\{i, k\}, & x_{max} &= \max\{i, k\} \\ y_{min} &= \min\{j, l\} - 1, & y_{max} &= \max\{j, l\} + 1 \end{aligned}$$

The *restricted query area* is represented by the dash-dotted square in Fig. 1. One may wonder why we opted for looser limits on the “*y-axis*”. The rationale behind this choice will be given through an example. Consider that $\Delta y = |l - j| \leq 1$ and that one of the two satellites or a group of intermediary satellites

is very close to one of the pole regions. That probably translates in some broken inter-plane ISLs, thus the establishment of a path between the source and the destination satellite cannot be assured. Thereby, looser limits guarantee that a path will be set up. Moreover, when the link-cost metric takes account of queuing delay, looser limits may result in the establishment of a shorter path.

B. Path Discovery Process

As soon as the *restricted query area* is formed, the *path discovery process* is initiated. The source satellite generates a RREQ packet regardless of whether it has a route entry for the destination satellite or not. This type of packet comprises eight fields which are described below.

Type identifier: it indicates that this is a RREQ packet.

Source satellite: the address of the source satellite.

Source sequence number: the sequence number of the source satellite.

Destination satellite: the address of the destination satellite.

Destination sequence number: the last obtained sequence number of the destination satellite. If the source satellite does not have a route entry for that satellite, then this field is set to zero.

Path cost: the cost of the path (i.e., either the propagation delay or the sum of the propagation and queuing delays) from the source satellite to the current satellite.

Path expiration time: the time instant at which the aforementioned path will become invalid due to the switching off of an ISL.

Timestamp: the time instant at which the RREQ packet was sent out by the source satellite.

The IP header is also attached to this packet and a replica of this packet is sent to those neighboring satellites whose virtual coordinates (x, y) satisfy the following criteria:

$$x_{min} \leq x \leq x_{max} \quad (1)$$

$$y_{min} \leq y \leq y_{max} \quad (2)$$

Now let us assume that an intermediate satellite has received a RREQ. First it will check if it has already received a RREQ from that source with the same *sequence number*. To this end, it checks if there exists an $R_{eq}T$ entry with the same sequence number and then, the *path cost* fields of the $R_{eq}T$ entry and the RREQ packet are compared. The RREQ is dropped only if *path cost* of former is smaller than the one of the latter, otherwise the following steps are taken.

The intermediate satellite first has to see if it has a *fresh RT* entry for the destination satellite. Hereinafter, an RT entry is considered to be *fresh* only if it has been recorded after the time instant contained in the *timestamp* field and its *destination sequence number* is greater than the *destination sequence number* of RREQ packet. Supposing that there exists a *fresh* route entry, that satellite replies to RREQ with a *reply packet* (RREP). This packet will arrive at the source satellite through the path that has been established by the RREQ packet, namely using the $R_{eq}T$ entries. However, for the sake of presentation let us assume that the intermediate satellite cannot satisfy the RREQ. In that case it records an $R_{eq}T$ entry for the source satellite (or it updates the existing one) and sends a replica of the RREQ packet to the satellites whose virtual coordinates, which are denoted as (x_s, y_s) , fulfill the criteria imposed by inequalities (1) and (2) and the following inequalities:

$$x_s \geq x_c, \quad \text{if } x_{src} \leq x_{dst} \quad (3)$$

$$x_s \leq x_c, \quad \text{if } x_{src} > x_{dst} \quad (4)$$

where x_c denotes the χ -coordinate of the current intermediate satellite, whereas x_{src} and x_{dst} denote the χ -coordinates of the source and destination satellites respectively. These equations ensure that RREQs will always be forwarded toward the direction of the destination satellite. In addition to these checks, the intermediate satellite also avoids sending a RREQ to the satellite from which it received the RREQ or forwarding a RREQ obtained from a path that will become broken very shortly. The last constraint can be expressed as

$$\text{path expiration time} \leq \text{current time} + \text{expiration threshold}$$

where *expiration threshold* is a parameter of the algorithm and determines how *fresh* a path is.

Now let us assume that the destination satellite has received a RREQ. It should be stressed that if the link-cost metric is based on the end-to-end delay, that is the sum of propagation and queuing delays, the first RREQ that will be received will correspond to the shortest path. Nevertheless, if the propagation delay is selected to be the link-cost metric, then it is evident that the first packet that will arrive at the destination satellite will not necessarily correspond to the shortest path. The latter case is more complicated, thus we will base our description on it hereafter.

If the destination satellite has already an $R_{eq}T$ entry for the source satellite with the sequence number indicated by the respective field of the packet, then it will compare the *path cost* field of this entry to the one of the RREQ packet. If the latter is greater, the packet is dropped, otherwise the entry is updated with the details of the RREQ packet, the *sequence number* of the destination satellite is incremented accordingly and a RREP packet is generated and sent to the source satellite through the path established by the received RREQ. As in the case of RREQs, the IP header is attached to RREPs as well. The fields of the RREP packet are similar to the ones of the RREQ packet.

Type identifier: this field indicates that this is a RREP packet.

Source satellite: it contains the address of the source satellite.

Destination satellite: it contains the address of the destination satellite.

Destination sequence number: it contains the sequence number of the destination satellite.

Path cost: the cost of the path from the source to the destination.

Path expiration time: it communicates the path's expiration time instant to the source satellite and is set equal to the respective field of the RREQ packet.

Timestamp: equal to the respective field of the RREQ packet.

As soon as the first RREP packet arrives at the source satellite, all the queued packets as well as any subsequent packets are sent to the destination satellite through this newly established path. The concept of *sequence number* has been reckoned to provide the loop-free property required in any multi-hop network [9]. Thus, the proposed technique guarantees the establishment of loop-free routes.

C. Route Entry Management

What is yet to be specified is the management of routing tables. As previously noted, each satellite maintains a route entry RT for each destination of interest. Each route entry contains a field called *route expiration time* that indicates the time instant at which the path will become invalid. However, a new path should be established prior to the expiration of the previous one. To this end, the source initiates the *path discovery process* at time $route\ expiration\ time - expiration\ threshold$. A new path should also be computed if one of the end terminals is handed over to another satellite. A new path established after the admission of a new connection or after a handover will also be of benefit to ongoing connections between terminals served by the same pair of satellites.

III. SIMULATION MODEL

The experiments conducted in this work aim at evaluating the performance of the proposed routing scheme as well as comparing it with three centralized routing schemes proposed so far in the literature. For this purpose, the simulation tool used was coded in the platform of the detailed simulation model of Ns2 [10]. However, as far as routing is concerned, only a centralized implementation of the Dijkstra algorithm, taking into account propagation delay, exists. The results to be presented were obtained by implementing the LAOR protocol and modifying the centralized implementation of Ns2 to take account of queueing delay as well. To the best of the author's knowledge, this is the first time that Ns2 is used to simulate and compare different routing protocols for LEO satellite systems. Two versions of the LAOR protocol were evaluated. One that takes only the propagation delay into account to find the shortest path and another that takes account of the total end-to-end delay, that is the sum of propagation and queueing delays. The five protocols were tested in an Iridium-like constellation, where ISLs are switched off when satellites cross the polar regions defined by a latitude threshold ($\pm 60^\circ$). Moreover 200 terminals were distributed over the six continents according to the *hot spot* scenario described in [11], which is based on the distribution of web servers. An exponential ON/OFF traffic generator is attached to each one of them. The simulation parameters are given in Table I.

TABLE I: Simulation parameters

Traffic generator's parameters	
Packet size	1500 bytes
"On" period	0.3 sec
"Off" period	0.9 sec
bitrate during "On" periods	200 kb/s - 1200 kb/s
System's parameters	
Up/downlink bandwidth	15 Mb/s
ISL bandwidth	10 Mb/s
ISL LL queue size	500 packets
LAOR queue size	3000 packets
Simulation duration	6050 sec

In order to enhance the performance of proactive protocols 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. Moreover, as far as LAOR is concerned, we set the parameter *expiration threshold* to 0.5 *sec* in order to ensure that no packets will be in-flight when the path becomes invalid.

A. Link Cost Metrics

The centralized routing schemes that we tested were proposed in [4], [5] and are based on the Dijkstra shortest path algorithm. The link-cost function used for these protocols is based on a mixture of propagation and queueing delay:

$$Link\ cost = WF_{PD} \times PD + WF_{QD} \times QD \quad (5)$$

where PD and QD denote the propagation delay and queueing delay respectively, while WF_{PD} and WF_{QD} represent the weighting factors given to them. The first of these three protocols is based only on propagation delay, thus $WF_{PD} = 1$ and $WF_{QD} = 0$. The two other schemes take account of queueing delay. In [4], [5] different combinations of WF_{PD} and WF_{QD} were studied. For the sake of fairness we set $WF_{PD} = 1$ and $WF_{QD} = 1$ because this combination yields the best performance. Furthermore, any other combination is superficial and does not represent accurately the network state. The difference between the two protocols that take into account the queueing delay, lies in the method the latter is estimated. In the first algorithm, queueing delay is taken into consideration by means of an average value over the update interval. The second one uses an *exponential forgetting function* (EFF).

To derive a feasible methodology for estimating queueing delay we used the ability of Layer-2 protocols to provide statistics related to ISL utilization. Then, according to the $M/M/1$ queueing model [12], the mean number of packets in the queue M_{PQ} can be estimated by

$$M_{PQ} = \frac{u_{ISL}}{1 - u_{ISL}} \quad (6)$$

where u_{ISL} is the mean link utilization provided by LL. When an EFF function is used, link utilization is calculated by:

$$u_{ISL}^k = link_state + decay \cdot (u_{ISL}^{k-1} - link_state) \quad (7)$$

where u_{ISL}^k and u_{ISL}^{k-1} are the estimated link utilizations when the k^{th} and $(k-1)^{th}$ packets, in the current interval, arrive at the interface queue. Link state is 0 if there is no packet either in the interface queue or under transmission, otherwise it is set to 1. Decay is the *forgetting rate* and is defined as

$$decay = e^{-\Delta t_k} \quad (8)$$

where Δt_k denotes the time interval between the arrivals of packets k and $k-1$ in the queue. Then, the queueing delay can be estimated by the following expression:

$$Queueing\ delay = M_{PQ} \frac{Packet\ size}{ISL\ bandwidth} \quad (9)$$

where the fraction denotes the packet transmission delay.

IV. PERFORMANCE RESULTS

For the sake of clarity, the five examined schemes are presented below.

Central.-PD: the centralized routing scheme when the link-cost metric is propagation delay.

Central.-AV: the centralized routing scheme when the link-cost metric is the sum of propagation delay and average queuing delay.

Central.-EXP: the centralized routing scheme when the link-cost metric is the sum of propagation delay and the queuing delay perceived by the exponential forgetting function.

LAOR-PD: the LAOR protocol when the link-cost metric is propagation delay.

LAOR-TD: the LAOR protocol when the link-cost metric is the sum of propagation and queuing delay.

It should also be noted that the presented results represent average values over 10 independent simulation runs. That number of runs provided 99% confidence intervals of $\pm 4\%$ in the worst case.

Our first objective is to show that LAOR can accurately deliver packets to their destination regardless of traffic intensity. Table II shows the delivery ratios for all the examined schemes and for different values of the terminal's bitrate. For all moderate bitrates all the examined routing schemes manage to successfully deliver generated packets to their destination. Nonetheless, when bitrate increases to 900 *kb/s* and above, some links become congested and therefore, packets are dropped. Note that *LAOR-TD* outperforms the rest of schemes, the performance of *LAOR-PD* is very close though. Concerning the centralized routing schemes, their poor performance compared to LAOR can be attributed to the limited ability of the periodic mechanism to effectively evaluate the network state. Therefore, we expect that these schemes will also perform poorly in terms of end-to-end delay.

Fig. 2, which depicts mean end-to-end delay versus terminal's bitrate, substantiates our expectations. It is evident that the LAOR protocol constitutes a significant improvement on centralized schemes. Specifically, for bitrates higher than 1000 *kb/s* mean end-to-end delay is dropped down by half. Its striking results are ascribed to its ability to capture traffic variations as well as changes in the propagation delay of ISLs, which stems from the fact that the shortest path discovery process is based on instantaneous measurements. Moreover, simulation results reveal that there exist no differences in the performance of the centralized routing protocols, albeit one would expect the schemes that take account of queuing delay to outperform the one that relies only on propagation delay. The reason for this poor performance is that as soon as a new path is computed for a pair of satellites, all communications between

TABLE II: Delivery Ratio

Bitrate (kb/s)	Central.-PD	Central.-AV	Central.-EXP	LAOR-PD	LAOR-TD
200-700	1	1	1	1	1
800	0.999967	0.999962	0.999967	0.999990	0.999992
900	0.999115	0.999122	0.999162	0.999904	0.999986
1000	0.995494	0.995826	0.995662	0.999346	0.999932
1100	0.988451	0.988860	0.988717	0.998021	0.999731
1200	0.979689	0.980435	0.980061	0.995544	0.999019

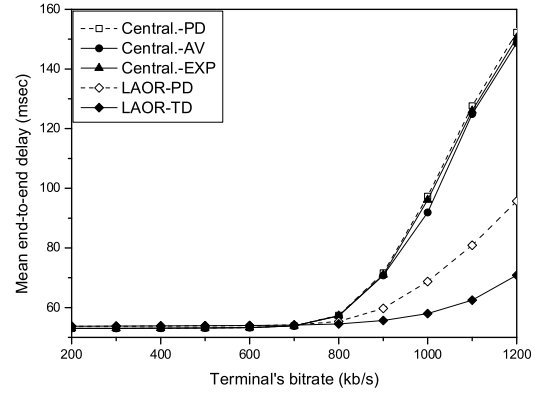


Fig. 2: Mean end-to-end delay vs terminal's bitrate

terminals served by these two satellites will be forced to use this new path, thus making it congested. This phenomenon repeats in every update cycle and is known as oscillation. It is more intense in connectionless communications, such as IP forwarding, since data flows quickly adapt to path switching.

Fig. 3 illustrates another performance metric that is of great interest, i.e., mean delay jitter. From this figure it becomes clear that LAOR outperforms centralized schemes. An amelioration in network performance is witnessed even for moderate bitrates. Not only do these results prove that LAOR succeeds in finding the less congested path, but that successive calculated paths do not differ significantly in terms of end-to-end delay as well. This is an indication that LAOR effectively splits traffic to different paths. As expected, when the path discovery process is based on propagation delay, the performance slightly degrades. At the other extreme, centralized protocols suffer from the effects of oscillations. The protocol that relies only on propagation delay achieves a better performance than the other two protocols since oscillations are not so intense in this case. On the contrary, oscillations are mostly the result of queuing delay variations. Furthermore, in order to minimize total delay, Central-EXP and Central-AV result in paths that consist of a greater number of hops. As we have already explained, these paths become congested thereupon their establishment, thus variations in end-to-end delay are slightly higher.

Last but not least, in order for our study to be complete

TABLE III: Normalized Overhead (%)

Bitrate (kb/s)	LAOR-PD		LAOR-TD	
	Pkt. Overh.	Byte Overh.	Pkt. Overh.	Byte Overh.
200	0.9593626	0.0325039	0.9486526	0.032148
300	0.6499853	0.0220204	0.6388396	0.0216494
400	0.4921545	0.0166723	0.4808541	0.0162958
500	0.3982082	0.0134892	0.3874941	0.0131321
600	0.335055	0.0113487	0.3237726	0.0109727
700	0.2890788	0.009791	0.2774186	0.0094018
800	0.2556861	0.0086039	0.2441891	0.0082758
900	0.2306593	0.0078107	0.2163541	0.0073324
1000	0.2115723	0.0071633	0.1949158	0.0066058
1100	0.1986213	0.0067237	0.1788703	0.0060621
1200	0.1878121	0.0063566	0.1653855	0.0056052

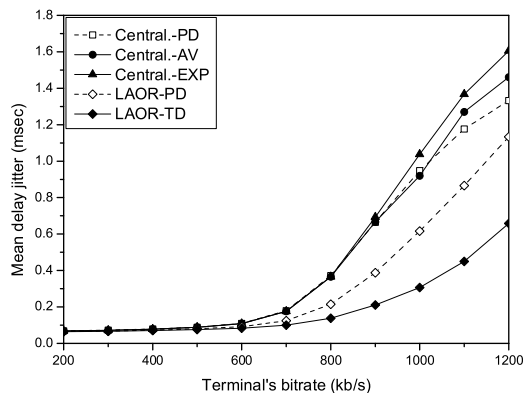


Fig. 3: Mean delay jitter vs terminal's bitrate

we should also evaluate the signaling overhead that LAOR induces to the system. Table III illustrates both the normalized packet overhead and the normalized byte overhead. Apparently, the overhead of this protocol is negligible regardless of the employed link-cost metric. Byte overhead appears to be much lower than packet overhead since RREQs and RREPs are much smaller than data packets. This kind of packets comprise a few fields (bytes) appended to the IP header. Even for the lowest bitrate, which represents the worst-case scenario in terms of signaling overhead, the normalized byte overhead is much less than 0.1%, whereas packet overhead is smaller than 1%. It should be noted that the two centralized protocols that take account of queuing delay have not been assessed because their performance in terms of overhead depends on the location of the central node that undertakes to collect the required information from all satellites, compute the shortest paths, and disseminate the routing tables, as well as on the route table update interval. Nevertheless, supposing that a ground station is shouldered this daunting responsibility, we can come up with a rough estimate of the normalized packet overhead. Let us assume that the algorithm is only periodically implemented every 10 sec. If we take into consideration that each satellite generates a packet with information about the queuing delay in its ISL queues, then we are led to the conclusion that every 10 sec 65 packets are generated (we do not consider the packet generated by the satellite that is directly connected to the designated ground station). As soon as the ground terminal computes the shortest paths it should distribute the routing tables to the satellites, thus 65 other packets are inserted into the network. This translates in 130 packets every 10 sec. Hence, the total amount of packets generated during the simulated time interval is $130 \cdot (6050/10) = 96800$. The average number of delivered data packets for bitrate 200kb/s is around $4.7 \cdot 10^6$, therefore, the normalized packet overhead is 0.02%. However, it is reasonable to believe that the size of those packets would be greater than the one of RREQs and RREPs, since the former would have to contain much more information. For example, packets containing the results of route calculation should carry 65 paths, one for each destination. On this account, the byte overhead of centralized routing schemes is expected to be

greater than the byte overhead of LAOR. Finally, as regards *Central.-PD*, it induces no overhead since it capitalizes on the deterministic dynamics of LEO networks. Nevertheless, this advantage is nullified by its poor performance.

V. CONCLUSIONS

In this paper we proposed and evaluated the performance of a location-assisted on-demand routing protocol for LEO satellite IP networks. In order to diminish the signaling overhead induced to the system, a mechanism has been proposed that aims to reduce the number of satellites that will take part in the aforementioned procedure. The proposed protocol was compared to centralized routing schemes that compute the shortest paths periodically based on the Dijkstra algorithm. Ample simulation results corroborated the superiority of LAOR over those schemes. Specifically, LAOR was shown to attain much lower end-to-end delay, delay-jitter and higher delivery ratio with negligible overhead. This fact renders it an excellent choice for future LEO satellite IP networks.

ACKNOWLEDGMENT

Mr. Stylianos Karapantazis thanks the Bodossaki Foundation for supporting his PhD studies.

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] M. Werner, C. Delucchi, H. Vögel, G. Maral, and J.-J. D. Ridder, "ATM-Based Routing in LEO/MEO Satellite Networks with Intersatellite Links," *IEEE J. Select. Areas Commun.*, vol. 15, no. 1, pp. 69–82, January 1997.
- [3] 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.
- [4] M. Mohorčić, 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.
- [5] M. Mohorčić, 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.
- [6] H. S. Chang, B. W. Kim, C. G. Lee, S. L. Min, Y. Choi, H. S. Yang, D. N. Kim, and C. S. Kim, "FSA-Based Link Assignment and Routing in Low-Earth Orbit Satellite Networks," *IEEE Trans. Veh. Technol.*, vol. 47, no. 3, pp. 1037–1048, August 1998.
- [7] E. Papapetrou and F.-N. Pavlidou, "A Proposal of Optimal Routing Techniques for non-GEO Satellite Systems," *International J. of Wireless Inform. Networks*, vol. 8, no. 2, pp. 75–83, April 2001.
- [8] 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.
- [9] C. E. Perkins and S. R. D. Elizabeth M. Royer, "Ad hoc on-demand distance vector (AODV) routing," RFC 3561, July 2003. [Online]. Available: <http://www.ietf.org/rfc/rfc3561.txt>
- [10] K. Fall and K. Varadhan, "The ns manual," VINT Project, Univ. California, Berkeley, CA, 2001. [Online]. Available: <http://www.isi.edu/nsnam/ns/ns-documentation.html>
- [11] M. Mohorčić, 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.
- [12] D. Bertsekas and R. Gallager, *Data networks (2nd ed.)*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1992.