

On the Impact of Intrinsic Delay Variation Sources on Iridium LEO Constellation

Amal Boubaker^{1,2}[0000–0002–2226–1004], Emmanuel Chaput¹, Nicolas Kuhn³[0000–0001–6671–3490], Jean-Baptiste Dupé³[0000–0001–9973–698X], Renaud Sallantin⁴[0000–0003–3421–6005], Cédric Baudoin⁴, and André-Luc Beylot¹[0000–0002–9049–4723]

¹ IRIT Lab, University of Toulouse, Toulouse, France

² TéSA, Telecommunications for Space and Aeronautics, Toulouse, France

³ Centre National d’Etudes Spatiales (CNES), Toulouse, France

⁴ Thales Alenia Space (TAS), Toulouse, France

Abstract. The recent decades have seen an increasing interest in Medium Earth Orbit and Low Earth Orbit satellite constellations. However, there is little information on the delay variation characteristics of these systems and the resulting impact on high layer protocols. To fill this gap, this paper simulates a constellation that exhibits the same delay characteristics as the already deployed Iridium but considers closer bandwidths to constellation projects’.

We identify five major sources of delay variation in polar satellite constellations with different occurrence rates: elevation, intra-orbital handover, inter-orbital handover, orbital **seam** handover and Inter-Satellite Link changes. We simulate file transfers of different sizes to assess the impact of each of these delay variations on the file transfer.

We conclude that the orbital **seam** is the less frequent source of delay and induces a larger impact on a small file transfers: the orbital **seam**, which occurs at most three times during 24 hours, induces a 66% increase of the time needed to transmit a small file. Inter-orbital and intra-orbital handovers occur less often and reduce the throughput by approximately $\sim 8\%$ for both low and high throughput configurations. The other sources of delay variations have a negligible impact on small file transfers, and long file transfers are not impacted much by the delay variations.

Keywords: Satellite constellations · Iridium constellation · CUBIC TCP · Handovers

1 Introduction

Nowadays, satellite constellations have reemerged, due to the need for worldwide high-speed internet coverage that terrestrial solutions fail to deliver. Medium Earth Orbits (MEO) and Low Earth Orbits (LEO) constellations could efficiently complete their Geostationary Orbits (GEO) counterpart. Many satellite constellation projects are competing with Iridium LEO constellation. While Iridium

offers low-speed data communications, constellation projects aim at increasing the throughput provided to the end user (*e.g.*, for collective terminal services).

We provide in this paper a large overview of the constellation projects. They can mainly be divided into two main families: the systems that exploit Inter-Satellite Links (ISLs) (such as Telesat or Starlink) and those who do not (such as OneWeb). Indeed, introducing ISLs enhances the coverage, reduces latency and limits the size of the ground segment. However, this can result in delay variation, which needs to be analysed and whose impact on transport layer needs to be assessed.

The aim of the paper is thus to determine to which extent delay variations caused by the intrinsic characteristics of the satellite constellation topology would affect the performances of CUBIC TCP algorithm, taking Iridium as a representative topology. The contributions of this paper are the following:

- we identify five major sources of delay variation in polar satellite constellations with different ranges of occurrence: elevation, intra-orbital handover, inter-orbital handover, orbital **seam** handover and ISLs changes;
- we measure that the orbital **seam** is the less frequent source of delay and induces the larger impact on a small file transfer;
- we measure that other sources of delay variations have a negligible impact on small file transfers and low impact on long file transfers.

It is worth pointing out that some studies on the impact of the satellite constellations on TCP behaviour have been carried out in the past [25,28]. However, the decomposition of the delay variations into different sources and the impact of each of them on a TCP-based transfer is missing. Moreover, considering some recent evolutions of the TCP stack and the increased throughput offered by the satellite constellations, the conclusions of these papers may be reconsidered.

2 State of the Art of satellite constellations projects

In this section, we present an overview of operating, in-development and satellite constellations, the motivation behind the choice of the Iridium constellation and the different sources of delay variations within the constellation.

2.1 Constellation projects

The most important information about operating, in-development and upcoming satellite constellations that could come to fruition is outlined in the Table 1. Such systems were proposed mainly for telecommunication purposes, the particular application LEO or MEO constellations we are interested in.

So as we can deduce from the Table 1, by 2020 we would be able to see MEO, LEO to VLEO satellite constellations, from different companies. They are expected to deploy their fleet, mostly with ISLs thus offering more connectivity and mainly in the Ka band hence more usable frequencies and better satellite

beams directivity. They should be then able to provide access to broadband services.

Among the existing satellite constellations presented in the Table, we have decided to focus on Iridium's. It is a LEO near-polar constellation, already deployed and has 4 ISLs, which we expect could be used as a representative constellation and would be generalized for the similar in-deployment and future projects.

We will now describe Iridium constellations, focusing on the properties that could jeopardise TCP performance.

2.2 Details on Iridium satellite constellation

In the Figure 1, a scheme of the Iridium topology on an unprojected map, we can distinguish the 6 orbital planes, the 11 satellites per plane and 4 ISLs per satellite.




Fig. 1: A plot of Iridium's topology [27]

An interesting aspect of polar or near-polar satellite constellations is the **seam** that we can see in the Figure 2, which is the gap between the last plane of 'ascending' satellites (moving towards north pole) and the counter-rotating (or 'descending' moving away from north pole) satellites. [29] Thus the **seam** is where ascending and descending planes pass each other. [5] Typically, Iridium's satellites maintain 4 ISLs: 2 intra-plane (in black) in the Figure 1 that are always maintained and 2 inter-plane intermittent links (in blue). However, the inter-plane links are deactivated close to the poles due to high-speed rotating satellites and on the **seam** because of counter-rotating overlapping satellites at high speed.

Table 1: Recent satellite constellations [20,19,1,11,9,10,23,13,14,15,21,22,24,?,17,18,16,8,6,7,12,3,4,2]

Operating satellite constellations						
Constellation	Operational in	Orbit	ISL Band	Type of service	# of satellites	# of spots per satellite
- Iridium (Thales Alenia Space, Orbital ATK)	2018	LEO (near-polar 780 km)	4 - Ka (inter-satellite) - L (uplink and downlink)	Low-speed data communications	- 66 - 6 orbital plans - 11 /plane	- 48 Mobile Satellite Service (MSS) beams +2 (feeder links) - steerable
- O3B 1 st Generation (Thales Alenia Space)	2014	MEO (8000 km)	— Ka	Trunking (cruise ship)	20	- 10 beams /sat - 12 steerable antennas
- Globalstar 2 nd Generation (Thales Alenia Space)	2010-2013	LEO (1414 km)	— L-S-C	Satellite phone and low-speed data communications	24	- 16 spot beams /sat - fixed spot beam
In-development satellite constellations						
Constellation	Operational in	Orbit	ISL Band	Type of service	# of satellites	# of spots per satellite
- O3Bn (O3B 2 nd Generation) (Boeing)	2021	MEO (8000 km)	None Ka	Trunking (cruise ship)	7	4000+ beams /sat (30000 spot beams)
- OneWeb (OneWeb, Airbus, Virgin)	2020	LEO (1200 km)	None - Ku (sat user link) - Ka (sat gateway link)	Broadband	- 720 - 18 near-polar orbital planes	- 16 user beams - non-steerable
Upcoming satellite constellations						
Constellation	Operational in	Orbit	ISL Band	Type of service	# of satellites	# of spots per satellite
- Starlink (SpaceX)	2024	2 sub-constellations : - VLEO (335,9-345,6 km) - LEO (1110-1325 km)	— V	Broadband	- VLEO : 7518 sats - LEO : 4425 sats	- Steerable (phased array)
- Boeing (Boeing Satellite)	2022	LEO	None V	Broadband	1396-2956	—
- LeoSat (Thales Alenia Space)	2022	LEO (1400 km)	4 Ka	Fiber optic-cable-like	78-108	12 steerable spot beams
- Theia (Theia Holdings A, Inc.)	2022	MEO (sun-synchronous 800 km)	— Ka	EBSS (Earth Exploration Satellite Service)	120	—
- Telesat Canada (Manufacturer To be defined)	2021	- LEO (polar 1000 km and inclined 1248 km)	max 4 Ka	Fiber optic-cable-like	- 117-512 (120 by 2021) - 6 polar planes: 12+ /plane - 5 inclined planes: 9+ /plane	45 (16+ steerable user beams & 2 steerable gateway beam)
- ViaSat (Viasat Inc.)	2020	MEO (polar 8200 km)	— Ka	Broadband	24 (Viasat 1 & 2) 8 planes	—




Fig. 2: Position of the orbital **seam** in polar or near-polar satellite constellations, particularly in Iridium and end terminals positions regarding the **seam**

The more ISLs there are, the more candidate routes are available, which introduces delay variations within the constellation. These delay variations have different sources that we will study further in the Sub-section 2.3.

Table 2: Iridium simulation parameters

Parameter	Value
Altitude	780 km
Planes	6
Satellites per plane	11
Inclination (deg)	86.4
Inter-plane separation (deg)	31.6
Seam separation (deg)	22
Intra-plane phasing	yes
Inter-plane phasing	yes
ISLs per satellite	4
Cross- seam ISLs	no

Satellites on both edges of the **seam** maintain two intra-plane links and only one inter-plane link. However, cross-**seam** links are a possible alternative to maintain the links through the **seam**, but for most of the polar or near-polar satellite constellations cross-seam links are turned off due to very large Doppler shifts and hand-offs have to happen more frequently and rapidly which might not allow the satellites to have a synchronized update on the state of the constellation and introduces further delay variation due to the variety and diversity of available paths within the constellation. In Iridium, cross-**seam** links are deactivated.

Table 2 sums up the main characteristics of the Iridium constellation.

2.3 On the delay variations in Iridium

LEO constellations are known to have lower propagation delay than the MEO or GEO ones. However, they are subject to higher delay variations due to the

movement of satellites with respect to the ground terminal and the constellation's topology itself. This can be broken down into several factors [26]:

1. **Elevation variation:** When the satellite is moving with respect to the ground terminal, the elevation angle of the satellite varies which results in a variation of the ground terminal-satellite slant-range. Thus, resulting in the variation of the propagation delay. This variation could be calculated using a simplified scheme in the Figure 3a Where :
 - O is the center of earth,
 - U is the position of the terminal on earth,
 - Sat_{min} is the position of the satellite in orbit for a minimum distance with the user on earth and d_{min} is the corresponding distance,
 - Sat_{max} is the position of the satellite in orbit for a maximum distance with the user on earth and d_{max} is the corresponding distance,
 - $\alpha = 8.2^\circ$ is the elevation angle of the maximum distance position for Iridium,
 - and $\alpha = 90^\circ$ is the elevation angle of the minimum distance position,
 - $R_e=6371$ km is the earth's radius,
 - $h=780$ km the altitude of the satellite,

Thus, $d_{min}=h=780$ km where the satellite's nadir point coincides with the user's position and with some trigonometric formulae we get $d_{max}=2463$ km. Hence, the propagation delay varies between [2, 6 ms; 8, 2 ms].
2. **Intra-orbital handover delay :** When the satellite drops below the elevation mask of the terminal, the connection is consequently handed over to the following satellite that could be in the same plane and that meets the criterion. For fixed communicating end terminals, this phenomenon occurs every ~ 10 minutes for Iridium for it has 6 orbital planes and 11 satellites per orbital plane. This results in frequent delay variations.
3. **Inter-orbital handover delay :** The rotation of earth on its axis or the movement of the ground terminal along the longitude make that the coverage of the ground terminal is handed over from the current covering satellite to another one in the adjacent orbital plane. This handover results in a route change and therefore in a delay variation. For fixed communicating end terminals, the inter-orbital handover happens typically every ~ 2 hours for Iridium.
4. **Seam handover delay :** A particularity of the polar or near-polar satellite constellations is that the satellites in the last and first orbital planes do not have any links with each one another. When satellites on both parts of the **seam** are sought, route changes occur *e.g.* typically the traffic is rerouted to the satellites that are over the pole. The separation by the **seam** of two fixed communicating end terminals happens at most three times and at least twice over 24 hours, depending on their position with respect to the **seam**. As for the duration of this phase, it depends on the longitudinal separation of the communicating end terminals.
5. **ISLs changes delay :** For polar or near-polar satellite constellations, inter-orbital links are deactivated at the poles because of high-speed rotating

satellites that cross one another. Which dictates rerouting accordingly and thus results in delay changes.




Fig. 3: Satellite elevation delay and different delay variations transitions

In the Figure 3b, we can see the different transitions of 4 out of the 5 sources of delay variations with respect to earth's motion and satellite constellation's topology, for a fixed terminal denoted by a human figure here. The fifth one (ISLs changes) can not be depicted here.

Most of these causes in the satellite constellation have a detrimental impact on the delay variation. In the Figure 4, the profile of the propagation delay of packets generated from a Constant Bit Rate (CBR) source is shown. The source is a terminal on a boat in the Atlantic Ocean to the destination is in London geographically. They are close enough to have a detailed look on the profile of the Iridium satellite constellation over 24 hours based on the delay as a cost metric for route computation. We can clearly see the effect of delay variations in the constellation. The elevation, the intra-orbital handoffs, the inter-plane handoffs and the **seam** delay variations, respectively noted by 1, 2, 3 and 4 on the Figure 4 (and in the text above). Indeed, we can see that the observed values confirm what has been said earlier. Some delay variation factors are frequent but with a low magnitude, while others are less frequent but with more important magnitude. Ergo, TCP should be dealing with each of these different sources of delay variation differently. However, this figure can not show exactly the ISLs changes that are due to the satellites crossing one another in the poles. This delay variation is a result of the hop count change due to routing changes.

3 Experiment test architecture

This section details the characteristics of the experiments that are driven to evaluate the performance of TCP when it faces the delay variations that occur in Iridium.




Fig. 4: One-way delay evolution in Iridium. Main factors of delay variation are numbered following Figure 3

3.1 Details on the TCP stack and the application layer data

- TCP stack: In this document, servers exploit CUBIC TCP (default congestion control algorithm in Linux and Windows systems) and clients use SACK with a 3-*Max SACK Blocks*.
- Application layer data: The impact of TCP performance in this context on the end user experience depends on the application that is carried out. In this paper, we consider file transfers of various sizes:
 - 9 kB file: this short file would fit in a 10-packet Initial Window;
 - 15 MB file: this larger file would let TCP get out of slow start;
 - unlimited-bulk file: this unlimited file would let us assess TCP behaviour in congestion avoidance phase.

These file sizes have been chosen in order to analyze the behaviour of TCP in various phases to better understand the impact of delay variations on its algorithm.

For this purpose, we will focus on three delay-sensitive metrics:

- file transfer time for a limited file size because from a *QoE* point of view it is the metric for which delay variation over the satellite constellation has a quantified impact from a user's perspective,
- *cwnd*
- and the instantaneous received throughput for unlimited file size, since it gives a more detailed and refined idea on the behaviour of TCP. In some cases, we also observe other relevant metrics, such as the congestion window or packet sequence number evolution.

3.2 Choice of Iridium gateway

Assessing the performance when both the terminals' and the gateways' locations vary can make the analysis quite complex. Since several gateways are available,

we had to appraise the relevance of choosing one satellite gateway for the rest of the paper.

We have run 50000 simulations. In each run, the terminal is given a random position and downloads at a random time a 9 kB file or a 15 MB file from one of the 7 Iridium gateways.

Figure 5 represents the transfer time for each of the 9 kB and the 15 MB files and for each Iridium gateway. The results illustrate that apart from maximum values that exhibit slight differences, the distribution of the transfer time is quite the same, whatever the position of the gateway.




Fig. 5: Transfer time for random starting time and terminal positions

We concluded from the results presented in this section that the choice of the Iridium gateway does not matter much. The rest of this paper considers Iridium's gateway at Hawaii.

3.3 Summary of the simulation characteristics

The Table 3 sums up the basic characteristics used for all the experiments presented in this paper. Other parameters (simulation start time, file sizes or end terminals positions) will be detailed separately in each following section. Other parameters, related to Iridium, can be found in Table 2.

We define two set of bandwidths (low throughput and high throughput, where there is a factor of 80). Both configurations have been considered in the evaluations and a subset of the results is presented in this paper. In general, the rationale behind this choice is that we want to determine the impact of a constellation with ISLs on TCP and not restrict our conclusions to the sole Iridium.

Table 3: Iridium simulation parameters

Parameter	Value
Tool	NS-2.34
Iridium gateway	Hawaii
Model	Error-free
Cost metric	Hop count
Queue size	BDP
TCP	Sender: CUBIC - Receiver: SACK
TCP IW	10 packets
Up/downlink bandwidth	Low throughput : 1.5 Mbps
	High throughput : 120 Mbps
ISLs bandwidth	Low throughput : 25 Mbps
	High throughput : 2 Gbps

4 On the impact of the seam

This section is thus dedicated to the analysis of the impact of the **seam** handover delay on a file transmission between two terminals. It also considers the low throughput configuration: results for the high throughput exhibited the same trend.

The **seam** handover delay (number 4) results in important delay variations. This is illustrated in Figure 4. One objective of this section is to evaluate to what extent the **seam** handover delay is much of an issue for various file sizes.

4.1 On the impact of the seam on 9 kB flows

We have run several simulations of a 9 kB file transfer between two terminals: one is the Iridium gateway at Hawaii, which is considered to be the server. The other terminal, the client, is placed in 17 different cities that we chose scattered around the globe to be representative enough to conclude on the results. The starting time of the simulations is chosen randomly within two of the intervals: when the end terminals are not separated by the **seam** and when they are separated by the **seam**. The metric visualized in these simulations is the transfer time of the file, application layer point of view.

Figure 6 represents the time needed to transfer 9 kB for terminals located in different cities, whether they are separated by the **seam** or not. We can clearly see that the mean transfer time when the end terminals are separated by the **seam** is much more important than when they are not separated by the **seam**. For instance, if we take worst and best cases *i.e.* the cities for which the **seam** has had the least and the most impact on the mean transfer time *e.g.* for New Delhi it goes from 124ms when the terminals are not separated by the **seam** to 144ms when the terminals are separated by the **seam** which makes a 15.6% increase. However, for Los Angeles the mean transfer time goes from 90 ms when

the terminals are not separated by the **seam** to 179 ms when the terminals are separated by the **seam** which makes a 98 % increase.




Fig. 6: Transfer time simulations of a 9 kB file




Fig. 7: Data transmission of a 9 kB file when **seam** handover occurs during transmission (4th use case of the Figure 4)

In order to better understand how the **seam** impacts the 9 kB file transfer, we present in Figure 7 different events for a source terminal on a boat in the Atlantic Ocean communicating to a destination end terminal in London. There are out-of-order ACKs which are due to the fact that the ACKs sent earlier have been through a longer route whereas the ones sent later have been through a shorter one. In the Figure 8, we can see that ACKs of packets n°0 and n°1 have been through a 1-hop path, right before the two end terminals being separated by the **seam**. When the **seam** separated the end terminals, the ACKs of packets n°2-3 took an 11-hop path. And the ACKs of packets n°4 .. 9 took an 10-hop path.




Fig. 8: Paths taken by ACKs of all the packets during a transmission of a 9 kB file when **seam** handover occurs during transmission (4th use case of the Figure 4) [27]

As opposed to the results presented in Figure 5, the results presented in Figure 6 exhibit an important variation. This can be explained by the fact that the terminals have random positions and the gateway has a fixed one in the Section 3.2 while the terminals and the gateway have fixed positions in this section.

This section has measured that, for a given satellite gateway and whatever the position of the satellite terminal, the **seam** has an important impact on small files transfers.

4.2 On the impact of the seam on 15 MB flows

This section focuses on the impact of the **seam** handover delay for the transmission 15 MB files. We have run several simulations of a 15 MB file transfer, between two terminals: one is the Iridium gateway at Hawaii, which is considered to be the server, and the other terminal, the client, is placed in 17 different

cities. The starting time of the simulations is chosen randomly within two of the intervals: when the end terminals are not separated by the **seam** and when they are separated by the **seam**. The metric visualized in these simulations is the transfer time of the file.

Figure 9 represents the time needed to transfer 15 MB for terminals located in different cities, whether they are separated by the **seam** or not. There does not seem to be any blatant difference between the values of the transfer time when the end terminals are separated by the **seam** and when they are not-separated by the **seam**.




Fig. 9: Transfer time simulations of a 15 MB file

4.3 Impact of the seam handover delay and file size

Section 4.1 illustrated that the **seam** handover delay increases the transmission time of a 9 kB in a non negligible manner. However, Section 4.2 illustrated that this impact can be neglected when the file weighs 15 MB. This section aims at determining to what extent the **seam** handover delay impacts a file transfer, depending on the file size.

Table 4 gathers the Mean Transfer Time (MTT) for a file transfer between one terminal located in Los Angeles and another located in Hawaii. $MTT_{no\ seam}$ means that the file transfer between the terminals is not affected by the **seam** handover delay while it is the case for the MTT_{seam} . This table confirms the conclusions that have been proposed in the beginning of this section.

We have run several simulations with different file sizes transferred from a server which is located near the Iridium gateway at Hawaii to a terminal client

Table 4: Mean Transfer Time (MTT) for a file transfer

File size Parameter	9 kB file	15 MB file
$MTT_{no\ seam}$	90.35 ms	89.67 s
MTT_{seam}	179.43 ms	91.02 s
$MTT_{no\ seam}/MTT_{seam}$ (%)	50.35	98.52

which is placed in 17 different cities. Table 5 presents the mean percentage values of the ratio of the $MTT_{no\ seam}$ and MTT_{seam} .

Table 5: Impact of the **seam** vs File sizes

File size (kB) Parameter	25	50	100	500	1000	5000	10000
$mean(MTT_{no\ seam}/MTT_{seam})(\%)$	59.38	59.35	69.79	90.71	96.78	99.97	98.55

Considering the values in the Table 5, we can conclude that starting 100 kB the **seam** has much less effect on file transfer.

5 On the impact of the delay variations due to different sources other than the seam

In this section, we will illustrate the impact of the different sources of delay variations of Iridium other than the **seam** on a TCP connection. We will also present the results for both low and high throughput configurations to illustrate the impact of the delay variation sources according to the values presented earlier in the Table 3.

5.1 Objective

We have seen that the variation of the delay in LEO constellations involve five phenomena that are unequally frequent, as detailed in the section 2.3. After studying in detail the effect of the **seam**, we now want to see closely the effect of the pendulous behaviour of the other three causes of delay variation on TCP that we could simulate *i.e.* elevation variation denoted by 1, intra-orbital handover delay denoted by 2 and inter-orbital handover delay denoted by 3 on the Figure 4.

5.2 Details of the simulation

We have run a simulation for a non-stop CUBIC TCP source which is an end terminal on a boat in the Atlantic Ocean and a receiving terminal in London.

The same couple of end terminals that we used in the section 2.3, particularly the Figure 4. The results shown here are during the steady phase *i.e.* after the slow-start during 51 min19s starting 12h21 min. The metrics visualized in this simulation are the *cwnd* and the instantaneous receiving throughput with a 1 second granularity.

The simulation set presented above are valid for both low and high throughput cases.

5.3 Results for the low throughput configuration




Fig. 10: Impact of different delay sources on *cwnd* and Receiving throughput for low throughput configuration

On the Figure 10, we have gathered the 3 different use cases of delay variation stated in the Subsection 2.3.

In the Sub-figure 10a, we can see that the elevation variation delay does not affect the performance metrics of TCP *i.e.* *cwnd* and the instantaneous receiving throughput. As explained in the Subsection 2.3, the delay varies smoothly between [2.6 ms; 8.2 ms].

If we focus on the intra-orbital handovers (from the 1st to the 2nd use cases in Sub-figure 10b), there is not much impact on the throughput even if it impacts the *cwnd* evolution. Since the delay is increasing, then previously sent packets and *acks* take less time to be received which induces the sender in incrementing the *cwnd*. However, in the reverse transition (from the 2nd to the 1st in the Sub-figure 10c), previously sent packets and ACKs take more time than the ones freshly sent since the delay is decreasing, which results in a decreasing of the *cwnd*, thus the throughput drops from 1.47 Mbps (given that the bottleneck is fixed to 1.5 Mbps) to 1.36 Mbps which is a 7.48% decrease.

If we focus on the inter-handover delay (for the transition from the 1st to the 3rd use case in Sub-figure 10d), the delay variation has an impact on the throughput as it goes from 1.47 Mbps to 1.35 Mbps which makes a 8.16% decrease. In addition to that, the *cwnd* and the throughput decrease and then pick up again twice after. We keep the same explanation as previously (the transition from the 1st to the 2nd use cases). Yet, the impact is more noticeable here, since the inter-orbital handover induces more delay than the intra-orbital delay, as seen in the Sub-section 2.3. From inter-handover to elevation in the Sub-figure 10e, we keep the same explanation as previously (the transition from the 2nd to the 1st use cases). Nonetheless, the impact is more considerable here, since the inter-orbital handover induces more delay than the intra-orbital delay, as seen in the Sub-section 2.3. The throughput decreases from 1.47 Mbps to 1.35 Mbps which makes a 8.16% decrease. This is similar to the impact of the intra-orbital (7.48% decrease in the value of the throughput).

5.4 Results for the high throughput configuration

On the Figure 11, we have also gathered the 3 different use cases of delay variation stated in the Subsection 2.3.

When compared to the low throughput counterpart, we can see that in the Sub-figure 11a, the elevation variation delay does not affect the performance metrics of TCP (*i.e.* *cwnd*) and the instantaneous receiving throughput for the high throughput configuration either. In these cases, the conclusions from the low throughput cases apply.

However, for the intra-orbital handovers (transition from the 1st to the 2nd use-cases, in Sub-figure 11b), we notice an impact on both the throughput and the *cwnd*. The throughput dropped from 118.19 Mbps (given that the bottleneck is fixed to 120 Mbps) to 116.69 Mbps which is a 1.27% decrease. From the transition from the 1st to the 2nd use-cases, in Sub-figure 11c, the throughput drops from 118.18 Mbps to 111.98 Mbps which is a 5.25% decrease.

For the inter-handover transition, (transition from the 1st to the 3rd use case in Sub-figure 11d), we keep the same explanation as for the intra-orbital handover (transition from the 1st to the 2nd use cases). The impact is slightly




Fig. 11: Impact of different delay sources on *cwnd* and Receiving throughput for high throughput configuration

more noticeable here, since the inter-orbital handover induces more delay than the intra-orbital delay, as seen in the Sub-section 2.3. The 1 to 3 transition has an impact on the throughput as it goes from 118.18 Mbps to 113.81 Mbps which makes a 3.7% decrease. In addition to that, the *cwnd* and the throughput decrease and then pick up again several times after. For the inter-handover transition (transition from the 3rd to the 1st use case in Sub-figure 11e), we can notice that the first glitch corresponds to an elevation to intra-orbital handover (the transition from the 1st use case: the elevation, to the 2nd, as in the Sub-figure 11b). As for the second glitch, we keep the same explanation as previously (the transition from the 2nd to the 1st use cases). Nonetheless, the impact is more considerable here, since the inter-orbital handover induces more delay than the intra-orbital delay, as seen in the Sub-section 2.3. The throughput decreases from 118.19 Mbps to 110.88 Mbps which makes a 6.19% decrease.

5.5 Discussion

The three studied sources of delay variations, elevation variation, intra-orbital handover delay and inter-orbital handover delay as in the Figure 4, have different impacts on a TCP connection that we can see through the *cwnd* and the instantaneous receiving throughput. The elevation handover occurs continuously, whilst for fixed communicating end terminals intra-orbital handovers occur every ~ 10 minutes and inter-orbital handovers happen every ~ 2 hours.

As stated in the Sub-sections 5.3 and 5.3, the elevation variation delay does not have an impact on the communication. The available goodput is fully exploited by TCP. TCP *cwnd* evolution is the same as CUBIC is expected to have over a non-variable delay link.

Nevertheless, the elevation to intra-orbital handover results in a 1.27% decrease in the instantaneous throughput for a high throughput configuration. Where, the intra-orbital to elevation handover results in a 7.48% decrease in the instantaneous throughput for a low throughput configuration and a 5.25% decrease for high throughput one.

Whereas, the elevation to inter-orbital handover can result in a 8.16% decrease for a low throughput configuration. As for the high throughput configuration, the elevation to inter-orbital handover results in a drop of 3.7% in the instantaneous throughput. While the inter-orbital to elevation handover results in a decrease of 6.19% in the instantaneous throughput.

The reduced throughput can be explained through the evolution of the *cwnd* of TCP that exhibits different average values for various operating points. This may be due to the different delays of the end-to-end paths that are exploited.

We notice that for the high throughput configuration the impact of the different delay variation is generally smaller. There may be an exception with the elevation to intra-orbital handover.

6 Conclusion

The recent decades have seen a boom in MEO/LEO satellite constellations for broadband purposes. That is why we examined the recent and deployed Iridium LEO satellite in this paper. In our work we focused on the impact of the constellation's topology dynamic on TCP from delay variation point of view. We have seen that there are 5 major sources of delay variation in polar satellite constellations: elevation, intra-orbital handover, inter-orbital handover, orbital **seam** handover and ISLs changes. All of our tests were conducted for only one single flow at a time to evaluate the impact of the intrinsic characteristics, *e.g.* delay variations induced by the satellite constellation topology on a TCP connection. The consequences of the constellation's behaviour on congestion deserves another in-depth study. We have seen that, using a recent TCP algorithm such as CUBIC TCP, during the orbital **seam**, which happens three times at most and at least twice during 24 hours, the transfer time suffers a significant increase for small flows (66.6% mean), whereas for large flows it is marginal (1.47% mean increase).

As for the other sources of delay variation, we have seen that, through the *cwnd* and the instantaneous receiving throughput, the inter-orbital handover, happening typically every ~ 2 hours, resulted in 8.16% decrease in the throughput from 1.47 Mbps to 1.35 Mbps. While in the high throughput configuration, the elevation to inter-orbital handover resulted in a 3.7% decrease in the throughput from 118.18 Mbps to 113.81 Mbps. And the reverse handover resulted in a 6.19% drop in the throughput from 118.19 Mbps to 110.88 Mbps. Whilst the intra-orbital handover, occurring every ~ 10 minutes, resulted in 7.48% decrease in the throughput from 1.47 Mbps to 1.36 Mbps, for a low throughput configuration. Whereas for a high throughput configuration, the elevation to intra-orbital handover resulted in a 1.27% decrease of the throughput from 118.19 Mbps to 116.69 Mbps, and the reverse transition resulted in a 5.25% drop from 118.18 Mbps to 111.98 Mbps. Unlike the elevation delay, which occurs continuously, that did not result in any change on the studied metrics for both of the low and high throughput configurations. These results could be extended to satellite constellations with similar characteristics.

The results that are shown in this paper illustrate that CUBIC TCP is not really impacted by the delay variations in a generic LEO satellite constellation. However, more recent TCP stack include reordering management based on delay variations. Future works include considering more complex stacks to confirm these trends and further confirm that TCP splitting mechanisms may not be needed for LEO constellations.

References

1. Fcc ngso constellations. <https://www.fcc.gov/document/cut-established-additional-ngso-satellite-applications>, 2017
2. Globalstar 1st generation. http://space.skyrocket.de/doc_sdat/globalstar-2.htm, [Online; accessed: 2018-03-12]
3. Globalstar 2nd generation. <https://fr.slideshare.net/SambitShreeman/iridium-globalstar-ico-satellite-system>, [Online; accessed: 2018-03-12]
4. Globalstar 2nd generation. <https://en.wikipedia.org/wiki/Globalstar>, [Online; accessed: 2018-03-12]
5. Introduction to satellite constellations. <https://savi.sourceforge.io/about/loyd-wood-isu-summer-06-constellations-talk.pdf>, [Online; accessed: 2019-04-15]
6. Iridium. https://en.wikipedia.org/wiki/Iridium_satellite_constellation#Next-generation_constellation, [Online; accessed: 2018-03-06]
7. Iridium. http://www.argo.ucsd.edu/sat_comm_AST13.pdf, [Online; accessed: 2018-03-06]
8. Iridium fcc. http://licensing.fcc.gov/myibfs/download.do?attachment_key=1031348, 2013
9. Leosat. <http://leosat.com/technology/>, [Online; accessed: 2018-03-06]
10. Leosat. <http://leosat.com/media/1114/leosat-technical-overview.pdf>, [Online; accessed: 2018-03-06]
11. Leosat fcc. http://licensing.fcc.gov/myibfs/download.do?attachment_key=1158225, 2017

12. O3b 1st generation. <http://spaceflight101.com/ses-orders-super-powered-satellites-from-boeing-to-expand-o3b-broadband-services/>, [Online; accessed: 2018-03-06]
13. O3b 2nd generation. https://en.wikipedia.org/wiki/O3b_Networks, [Online; accessed: 2018-03-06]
14. O3b 2nd generation. http://space.skyrocket.de/doc_sdat/o3b-21.htm, [Online; accessed: 2018-03-06]
15. O3b 2nd generation. <https://www.ses.com/networks/o3b-mpower>, [Online; accessed: 2018-03-06]
16. Oneweb. https://fr.wikipedia.org/wiki/OneWeb#Caract%C3%A9ristiques_techniques, [Online; accessed: 2018-03-06]
17. Oneweb fcc. https://licensing.fcc.gov/myibfs/download.do?attachment_key=1134939, 2013
18. Oneweb fcc fact sheet. https://transition.fcc.gov/Daily_Releases/Daily_Business/2017/db0601/DOC-345159A1.pdf, 2017
19. Satellite constellation. https://en.wikipedia.org/wiki/Satellite_constellation, [Online; accessed: 2018-03-06]
20. SpaceX fcc. https://licensing.fcc.gov/myibfs/download.do?attachment_key=1190019, 2017
21. Telesat. <https://www.telesat.com/services/leo>, [Online; accessed: 2018-03-12]
22. Telesat. <http://www.mit.edu/~portillo/files/Comparison-LEO-IAC-2018-slides.pdf>, [Online; accessed: 2019-03-21]
23. Theia. <http://www.parabolicarc.com/tag/theia-holdings/>, [Online; accessed: 2018-03-12]
24. Viasat. <https://www.viasat.com/news/going-global>, [Online; accessed: 2018-03-12]
25. Chotikapong, Y., Cruickshank, H., Sun, Z.: Evaluation of tcp and internet traffic via low earth orbit satellites. *IEEE Personal Communications* **8**(3), 28–34 (2001)
26. Goyal, R., Kota, S., Jain, R., Fahmy, S., Vandalore, B., Kallaus, J.: Analysis and simulation of delay and buffer requirements of satellite-atm networks for tcp/ip traffic. *arXiv preprint cs/9809052* (1998)
27. Henderson, T.R., Wood, L.: ns-2 satellite plot scripts (2000), <https://savi.sourceforge.io/sat-plot-scripts/>
28. Subramanian, S., Sivakumar, S., Phillips, W.J., Robertson, W.: Investigating tcp performance issues in satellite networks. In: 3rd Annual Communication Networks and Services Research Conference (CNSR'05). pp. 327–332. IEEE (2005)
29. Wood, L.: Satellite Constellation Networks, pp. 13–34 (04 2003). https://doi.org/10.1007/978-1-4615-0431-3_2