

POWER

Technical Report 2020-13

Title: **LEO satellite constellations to offload optical terrestrial networks in placement of popular content in 5G edge nodes**

Authors: Alexis A. Dowhuszko, Juan A. Fraire, Musbah Shaat, Ana Pérez-Neira

Report Number: 2020-13

ERC Project: Power to the People. Verified.

ERC Project ID: 695614

Funded Under: H2020-EU.1.1. – EXCELLENT SCIENCE

Host Institution: Universität des Saarlandes, Dependable Systems and Software
Saarland Informatics Campus

Published In: ICTON 2020

This report contains an author-generated version of a publication in ICTON 2020.

Please cite this publication as follows:

Alexis A. Dowhuszko, Juan A. Fraire, Musbah Shaat, Ana Pérez-Neira.

LEO satellite constellations to offload optical terrestrial networks in placement of popular content in 5G edge nodes.

22nd International Conference on Transparent Optical Networks, ICTON 2020, Bari, Italy, July 19-23, 2020. IEEE 2020, ISBN 978-1-7281-8423-4: 1-6.



POWER TO THE PEOPLE.
VERIFIED.



LEO satellite constellations to offload optical terrestrial networks in placement of popular content in 5G edge nodes

Alexis A. Dowhuszko¹, Juan Fraire^{2,3}, Musbah Shaat¹, and Ana Pérez-Neira^{1,4}

¹Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Barcelona, Spain

²CONICET - Universidad Nacional de Córdoba (UNC), Córdoba, Argentina

³Saarland University, Saarland Informatics Campus, Saarbrücken, Germany

⁴Dept. Signal Theory and Commun., Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

Tel: (+34) 93 645 2900. Email: alexis.dowhuszko@cttc.es; juanfraire@unc.edu.ar;

musbah.shaat@cttc.es, ana.perez@cttc.es

ABSTRACT

Telecom operators are constantly searching for solutions to deal with notable growths on the data demand. One of these options consists in predicting the popularity of content to-be-requested by the subscribers on a daily basis, and store them in advance on the 5G edge nodes during the off-peak hours. To perform this placement, point-to-point optical terrestrial links are conventionally used, transmitting the same content in parallel from the operators' servers to the 5G edge nodes. One simple way to reduce this placing time consists in offloading the terrestrial (optical) network, by multi-casting over satellite (radio) the most popular content. This way, top-ranked files are transmitted only once by the satellite node, and can be received by all 5G edge nodes in coverage range. To implement the satellite segment of such hybrid backhaul solution, two options are considered: A small-sized constellation of low-cost LEO satellites with sporadic contacts with ground points, and a single high-cost GEO satellite that enables continuous pan-European connectivity regardless the location. Simulation results show that the placing time can be notably reduced using the proposed hybrid terrestrial-satellite backhauling. More precisely, the use of few LEO satellites can provide a placing time reduction that is similar to the one with GEO, even though the implementation and operational costs can be much lower with LEOs than with a GEO.

Keywords: Free Space Optical, LEO constellation, GEO satellite, Correlated Zipf-distribution, 5G edge caching.

1. INTRODUCTION

With the proliferation of on-demand video streaming services such as Youtube and Netflix, mobile operators must find solutions to deal with the much higher data demand of customers that are not willing to pay much more for their subscription fee. One of these solutions has been the construction of a Content Delivery Network (CDN) owned by the operator, where the most popular content *to-be-requested* by subscribers is first predicted on a daily basis, and then stored in advance on the corresponding 5G edge nodes during the night time [1]. For this purpose, point-to-point wired (optical fiber) or wireless (Free-Space Optical) links are usually deployed to connect the operators' servers with the 5G edge nodes. However, due to the multi-hop physical topology that terrestrial backhaul links have, the same piece of content must be transmitted in parallel to each 5G edge node that requests it, scaling notably the placing time that is needed even when the data rate that each individual optical link can support is high. One simple way to tackle this problem consists in offloading the terrestrial optical network, by multi-casting the most popular content to the edge nodes in the same geographical area [2].

Geostationary (GEO) satellites have been around since 1960s, and have been traditionally used to provide *continuous* communication services taking advantage of their fixed position on the sky. On the other hand, after the commercial failure that Low Earth Orbit (LEO) satellite ventures experienced in the 1990s, the use of these constellations for broadband *bursty* data connectivity has gained momentum again lately, motivated by more flexible RF regulations, cheaper launch costs, and reduced capital and operational expenditures that LEO satellites have. Only recently, the integration of satellite and terrestrial communication technologies has been seriously considered as part of 5G [3], and interesting papers have been published in this area [2]. For example, the authors of [4] studied the benefits of using a GEO satellite to feed several network caches at the same time, whereas the optimal placement of content on-board LEO satellites was considered in [5] to minimize the delay. Nevertheless, most of these reported results depend notably on the architecture that was implemented, and addressed mainly the dimensioning of the CDN to attain *e.g.* a given cache hit probability [4].

In this paper, we assume that the indexes of files to be requested by 5G edge nodes in each geographical region are known in advance, and we focus on the resource allocation strategy that the hybrid satellite (radio) and terrestrial (optical) backhaul network should apply to minimize the placing time. The study that is presented here complements the one published in [6], where a cluster of 5G edge nodes equipped with GEO and Free Space Optical (FSO) links was considered, assuming that all edge nodes belonged to the same geographical region (*i.e.*, same Zipf-popularity distribution) [7]. In this paper, a pan-European coverage of the CDN is assumed, in which different geographical areas with correlated Zipf-popularity distributions should be defined. Apart from characterizing the time reduction that is achieved when placing (part of) the content using a GEO satellite, we also study the performance gain when using LEO constellations with a baseline scheduling algorithm. Note that in the new proposed approach, the LEO satellites act as data mules (*i.e.*, storing and forwarding data), whereas the GEO satellite works in a bent-pipe mode (*i.e.*, amplifying and forwarding the message).

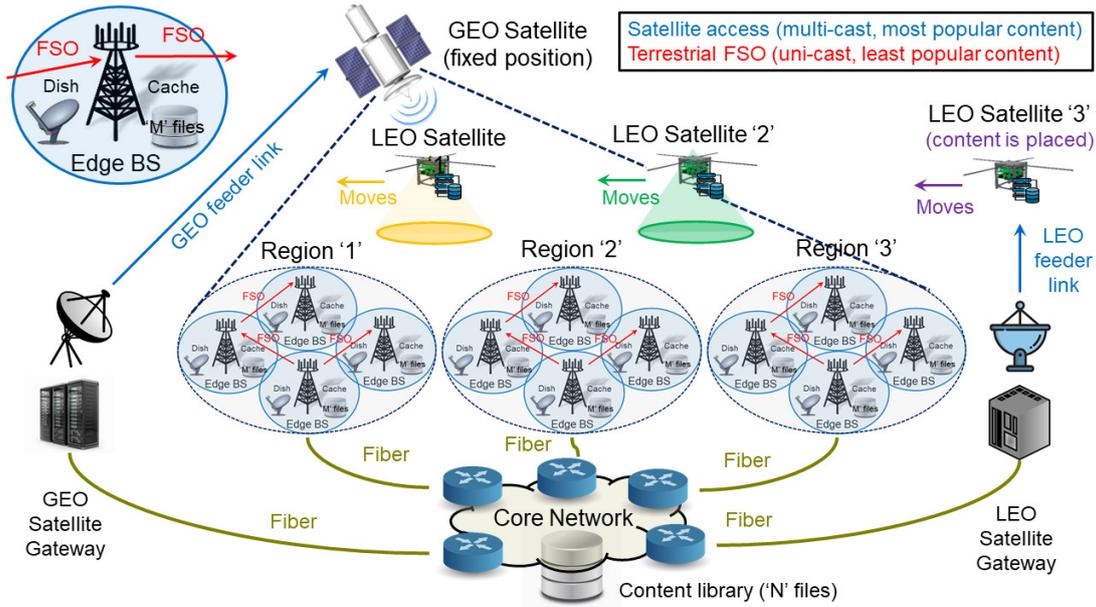


Figure 1: Hybrid terrestrial (FSO) and satellite (GEO/LEO) backhaul network for content delivery into 5G edge. In this example, there are $G = 3$ geographical regions, $L = 3$ LEO satellites, and $K = 4$ edge BSs per region.

The rest of the paper is organized as follows: Section 2 gives an overview of the hybrid backhaul network, including the statistical model to determine the number of 5G edge nodes requesting each piece of content per region, the details of the GEO/LEO satellite orbits and link budgets, and the baseline scheduling algorithm for LEOs. Section 3 shows the simulation results, with emphasis on the placing time improvement when using both terrestrial (FSO) and satellite (GEO/LEO) technologies. Finally, conclusions are drawn in Section 4.

2. SYSTEM MODEL

The simplified system model of the hybrid backhaul network that combines both optical terrestrial (multi-hop) and GEO/LEO satellite (multi-cast) wireless links for 5G edge cache filling is illustrated in Fig. 1. We assume that the whole pan-European coverage area is divided into G non-overlapping geographical regions which do not necessarily overlap with the territory of different countries. Each of these regions, whose index $g \in \mathcal{G}$ can be also mapped into coordinates (k, l) as shown in Fig. 2a, have on average K edge Base Stations (BSs) whose associated users will request the delivery of content taken from a universal library with N elements; from them, the M most popular will be pro-actively stored in each 5G edge node. Each edge BS is equipped with (at least) a pair of FSO transceivers, which connect the edge node with the nearest neighbors, and through them to the operator's server in a multi-hop fashion. In addition, each edge BS has also a satellite dish to receive the content that the GEO satellite, or the LEO constellation with L satellites, multi-cast to all the BS in range (GEO satellite is always visible, whereas the visibility of a given LEO satellite varies with the time of the day).

Finally, the use of a pre-defined off-line caching algorithm is assumed, which is able to determine *a priori* the indexes of files to be placed in each edge node. Moreover, it is considered that the resource manager of the core network has access to this information before the placing period starts, to identify the most convenient way to deliver the content using either a GEO/LEO satellite multi-cast or multiple FSO uni-cast transmissions.

2.1 Popularity Distribution

The probability that a given file with index $n \in \mathcal{N}$ is randomly requested by a user associated to an edge BS is assumed to follow the Zipf's law [7]. That is, popularity of file n in region g is approximated as

$$p_g[n] = \frac{\Omega_g}{n^{\alpha_g}} \quad n = 1, \dots, N, \quad \Omega_g = \left(\sum_{n=1}^N \frac{1}{n^{\alpha_g}} \right)^{-1} \quad g = 1, \dots, G, \quad (1)$$

where α_g controls the tail of the Zipf-distribution. Without loss of generality, we assume $\alpha_g = 0.7$ for all $g \in \mathcal{G}$.

We assume that the European territory is divided into $G = 35$ regions of approximate size $550 \text{ km} \times 550 \text{ km}$ ¹, as illustrated in Fig. 2a. Then, we need to generate the Zipf-distributions of these regions jointly, such that the popularity of a given file is not equal in all regions, but it is somehow correlated. That is, if an item is ranked high in a given region, then it is also expected to be very popular in a neighboring region, and vice versa. Moreover, to make this approach consistent, we assume that the farther the regions under analysis are, the more unlikely is that they share the same content in the ranking of most popular files (*e.g.*, in the top-100 ranking) [8].

Let us assume that ρ_g is the coefficient that controls the correlated rank and popularity of files in different regions. Then, the proposed algorithm to obtain the correlated Zipf-distributions is summarized as follows:

¹For the sake of precision, the size of each region is 5-deg latitude by 4.29-deg longitude, which is actually closer to $550 \text{ km} \times 320 \text{ km}$.



(a) Indexes of regions to model popularity distribution.

LEO Space Segment		GEO Space Segment	
Satellites	1 to 8	Satellite	1
Formation	Ring	Longitude point	10 deg
Altitude	600 km	Inclination	0 deg
Arg. Perigee	0 deg	Altitude	35788 km
Inclination	98 deg	Ground Segment	
RAAN	0 deg		
True Anomaly	0 to 360 deg	Regions(5x7)Lat	35 to 60 deg
Access cone	120 deg	Regions(5x7)Lon	-10 to 20 deg

	LEO	GEO
Frequency	15 GHz (Ku-band)	40 GHz (Ka-band)
Beamwidth	90 deg	6 deg
Gain	6 dB	29.5 dB
Trans. Power	30 dBW	
Modulation	BPSK	
Target BER	<1.0e-9	
Data rate	300 Mb/sec	50 Mb/sec
Bandwidth	600 MHz	100 MHz

(b) Parameters of the LEO and GEO space segment.

Figure 2: Ground and satellite configuration parameters of the hybrid backhaul network for 5G edge caching.

- 1) **Independent Gaussian Random Matrix (G)**. Aim is to generate a matrix containing independent and identically distributed (i.i.d.) samples $G_{(k,l)}$ taken from a zero-mean unit-variance Gaussian distribution.
- 2) **Correlated Gaussian Random Matrix (\tilde{G})**. The correlated Gaussian samples are then derived as follows:

$$\tilde{G}_{(k,l)} = \begin{cases} G_{(k,l)} & \text{for } k = 0 \text{ and } l = 0, \\ \rho_g \tilde{G}_{(k-1,l)} + \sqrt{1 - \rho_g^2} G_{(k,l)} & \text{for } k \geq 1 \text{ and } l = 0, \\ \rho_g \tilde{G}_{(1,l-1)} + \sqrt{1 - \rho_g^2} G_{(1,l)} & \text{for } k = 0 \text{ and } l \geq 1, \\ \rho_g [\tilde{G}_{(k,l-1)} + \tilde{G}_{(k-1,l)}] - \rho_g^2 \tilde{G}_{(k-1,l-1)} + (1 - \rho_g^2) G_{(k,l)} & \text{otherwise.} \end{cases} \quad (2)$$

The mean and variance of the correlated Random Variables (RVs) derived in (2) are given by $E\{\tilde{G}_{(k,l)}\} = 0$ and $\text{Var}\{\tilde{G}_{(k,l)}\} = 1$ regardless of the values of k and l , whereas the correlation coefficient attains the form

$$\text{Corr}\{\tilde{G}_{(k_1,l_1)}, \tilde{G}_{(k_2,l_2)}\} = \frac{E\{\tilde{G}_{(k_1,l_1)} \tilde{G}_{(k_2,l_2)}\} - E\{\tilde{G}_{(k_1,l_1)}\}E\{\tilde{G}_{(k_2,l_2)}\}}{\sqrt{\text{Var}\{\tilde{G}_{(k_1,l_1)}\}} \sqrt{\text{Var}\{\tilde{G}_{(k_2,l_2)}\}}} = \rho_g^{|k_1-k_2|+|l_1-l_2|} \quad \forall k_1, k_2; l_1, l_2. \quad (3)$$

- 3) **Correlated Uniform Random Matrix (\tilde{U})**. The correlated Gaussian samples obtained in the previous step are now mapped into uniformly distributed random samples in the interval $[0, 1]$ using the transformation

$$\tilde{U}_{(k,l)} = \Phi[\tilde{G}_{(k,l)}] \quad \forall k, l, \quad (4)$$

where $\Phi(x)$ is the Cumulative Distribution Function (CDF) of a standard Gaussian RV X .

- 4) **Rank and popularity of files in each region ($r_g[n]; p_g[n]$)**. Let us assume that N samples from RV $\tilde{U}_{(k,l)}$ are obtained for each region $g \in \mathcal{G}$, which can be mapped into coordinates (k, l) using Fig. 2a. Then, after ordering these samples, the rank of a given file with global popularity index n is determined according to

$$r_g[n] < r_g[m] \iff \tilde{U}_{(k,l)}[n] > \tilde{U}_{(k,l)}[m] \quad n = 1 \dots, N; m = 1, \dots, N; n \neq m, \quad (5)$$

whereas the popularity is obtained by replacing index n by $r_g[n]$ in (1). That is, $p_g[n] = \Omega_g / (r_g[n])^{\alpha_g}$.

Based on the analysis presented in [6], the average number of edge BS requesting file n in region g is

$$B_g[n] = \sum_{k=0}^K k P_g[k], \quad P_g[k] = \binom{K}{k} \exp[-\lambda_g(K-k)] [1 - \exp(-\lambda_g)]^k, \quad \lambda_g = p_g[n] M, \quad (6)$$

where K is the average number of edge BS per region and M is the average number of files requested per edge BS. Since the first element of the sum in the left-hand side of (6) equals zero, we have that

$$\begin{aligned} B_g[n] &= K [1 - \exp(-\lambda_g)] \sum_{k=1}^K \binom{K-1}{k-1} [\exp(-\lambda_g)]^{K-k} [1 - \exp(-\lambda_g)]^{k-1} \\ &= K [1 - \exp(-\lambda_g)] \sum_{l=0}^{K-1} \binom{K-1}{l} [\exp(-\lambda_g)]^{(K-1)-l} [1 - \exp(-\lambda_g)]^l \\ &= K [1 - \exp(-\lambda_g)] [\exp(-\lambda_g) + 1 - \exp(-\lambda_g)]^{K-1} = K [1 - \exp(-\lambda_g)], \end{aligned} \quad (7)$$

where the latter expression results when applying the inverse of the well-known binomial theorem.

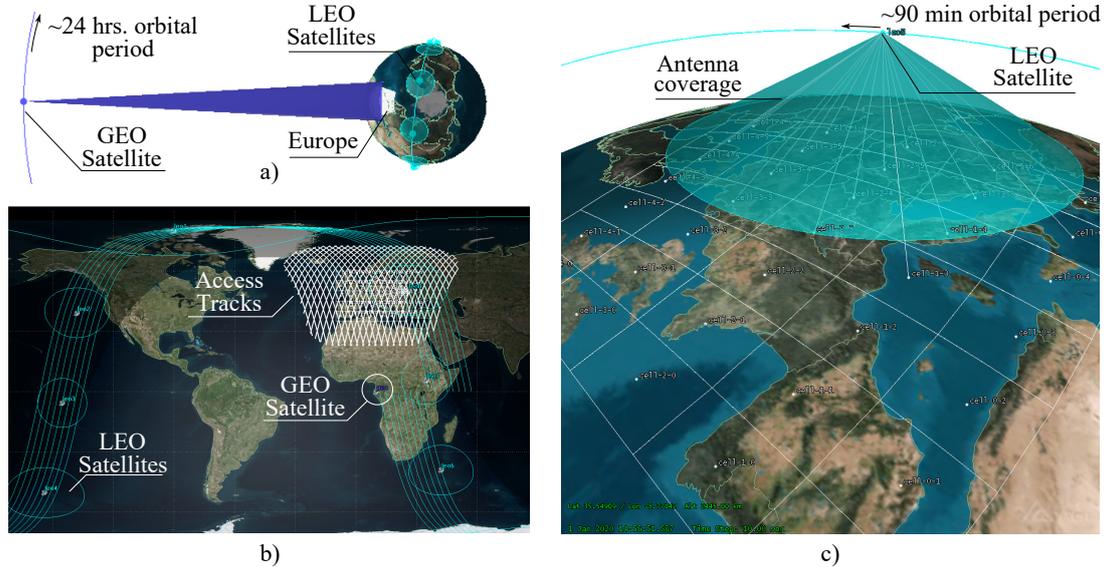


Figure 3: Orbital configurations for the satellite constellations under analysis. (a) GEO satellite remains fixed (blue beam), whereas LEO satellites move (cyan beam). (b) LEO satellites fly in cluster (along-track), passing over European territory twice a day. (c) Duration of LEO contact opportunities depends on the orbit dynamics.

2.2 Orbital configuration of GEO satellite and LEO satellite constellation

Two orbital configurations are evaluated: LEO with L satellites and GEO with single satellite. Figure 2b summarizes the key parameters of both space and ground segments. Orbital propagation is computed for 24-hour window, starting on 1st Jan. 2020 at 00:00:00 GMT, using the Systems Tool Kit (STK) developed by AGI [9].

- 1) **GEO satellite:** Placed at 35 788 km, its position remains fixed when observed from the Earth, enabling pan-European coverage continuously. In particular, the 5×7 grid of regions – in which the continent is divided – can be served by a single antenna/spot-beam, multi-casting in *bent-pipe* mode the most popular content to-be-placed on all the edge BSs deployed in the whole area. Note that a similar GEO scenario was already addressed in [6], but with emphasis on the hybrid placement of content in a single geographical region. In this work, we consider a single GEO satellite placed over 10-deg longitude and 0-deg inclination.
- 2) **LEO satellites:** Placed at about 600 km, each LEO satellite orbits the Earth in about 90 min. In contrast to GEO, the LEO satellite is not continuously connected to the ground station and, due to that, acts as a *data mule*. Moreover, the LEO satellite *passes* over a territory offering communication opportunities for no more than 10 minutes, 4 times-a-day. However, the reduced range in LEO links enables a better channel margin than in GEO. To overcome the limited contact time, LEO satellites are typically arranged into constellations. This way, when a LEO satellite starts to hide below the horizon, a subsequent LEO satellite should become visible to continue serving the target area. As expected, the more satellites are in the constellation, the more continuous will be the content placing service. However, persistent connectivity from LEO requires constellations of several dozens of satellites, such as in Iridium and Starlink.

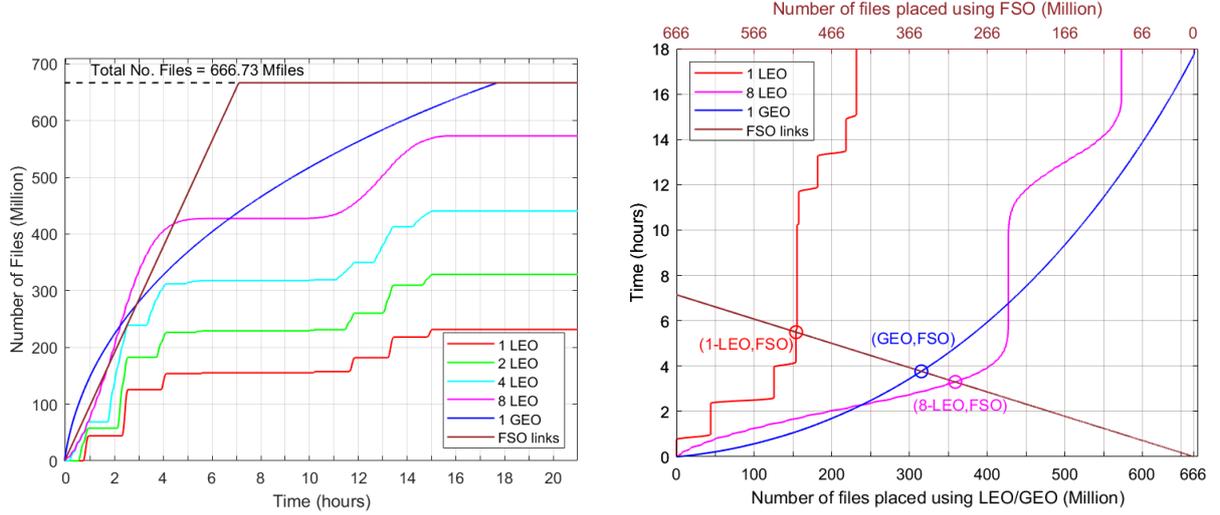
The aim of this paper is to show that a small-sized LEO constellation can provide reliable placement of content in 5G edge nodes, but demanding a lower implementation complexity/cost when compared to GEO. In particular, since the placing of content does not require continuous connectivity, we claim that few LEO satellites can act as data mules, storing popular files on-board and distributing them when flying over the target regions. For this purpose, we evaluate a *ring* formation, where LEO satellites fly on the same orbital track, as illustrated in Fig. 3b. Our study explores ring constellations with $L = 1, 2, 4, 8$ LEO satellites, in order to determine the scalability of placing time performance with respect to the size of the LEO constellation.

2.3 Link budget analysis of both GEO and LEO satellites

The key parameters of the satellite radio links are summarized in Fig. 2b. In particular, appropriate carrier frequencies in the Ku- and Ka-band are selected for the LEO and GEO links, respectively [10]. Also, due to the satellite-to-ground range, the antennas and resulting channel gains exhibit different properties. While the antenna on GEO provides pan-European coverage with a 6-deg beamwidth (29.5 dB main-lobe gain), the LEO satellites are equipped with a 90-deg beamwidth antenna (6 dB main-lobe gain), mimicking the so-called iso-flux antenna pattern [11]. Resulting data rates are $R_{leo} = 300$ Mbps and $R_{geo} = 50$ Mbps, when target BER is 10^{-9} . Separation between LEO satellites was selected to have (at most) one LEO satellite visible per regions. This way, the effect of inter-satellite interference coming from the same LEO constellation can be neglected.

2.4 Scheduling approach for delivery of content with LEO satellites

Since the GEO satellite can cover the whole European territory with a single antenna (see Fig. 3a), the scheduling approach to be used in GEO is straightforward: That is, files are ranked according to their global popularity, and most popular files are transmitted first. On the other hand, though the beamwidth of LEO antennas enables



(a) Number of files placed using different technologies.

(b) Integration of satellite and terrestrial communications.

Figure 4: Performance of the different transmission technologies when used to cache content in 5G edge nodes. Ground segment: $G = 35$ regions, $K = 23\,656$ BS/region, $M = 802.25$ files/BS, $R_{\text{fso}} = 1$ Gbps. Space segment: $R_{\text{geo}} = 50$ Mbps, $R_{\text{leo}} = 300$ Mbps. Popularity distribution. $N = 10\,000$ files (server), $\alpha_g = 0.7$ and $\rho_g = 0.9$.

a wider (120-deg) field-of-visibility, its lower orbit enables to cover only part of European territory in each pass (see Fig. 3c). Accordingly, a scheduling algorithm is needed to select the content to multi-cast in each pass.

The definition of the optimal scheduling approach that maximizes the number of files that can be placed in each LEO pass is not trivial. This is because the coverage of a LEO satellite is limited and, in each pass, different regions with different popularity rankings of content are visible [12]. Nevertheless, we now propose the following simplifying assumptions to identify a feasible upper bound for the LEO placing time performance:

- The 24-hour time window is divided into instances with index $i \in \mathcal{I}$. In each of these instances – with duration t_i sec. – the geographical regions $g \in \mathcal{G}$ that are visible by the same LEO satellite do not change.
- The resource manager of the hybrid network estimates in advance the number of edge BS requesting each file per region ($B_g[n]$), and makes sure that each LEO satellite has on-board the content to-be-delivered.
- The number of files to be broadcast by LEO satellite $l \in \mathcal{L}_i$ is given by $M_i = \lfloor (t_i R_{\text{leo}}) / (8S) \rfloor$, where $\lfloor x \rfloor$ is the greatest integer less than or equal to x , and S is the average size of each file in MB.
- Let $\mathcal{G}_{l,i}$ be the regions in which LEO satellite $l \in \mathcal{L}_i$ is visible. Then, the files that are transmitted verify:

$$n_{l,i}^* = \arg \max_{n \in \mathcal{N}} \sum_{g \in \mathcal{G}_{l,i}} \max \{B_{g,i}[n]\}, \quad \text{where } B_{g,1}[n] = B_g[n] \forall n \text{ and } B_{g,i+1}[n_{l,i}^*] = 0. \quad (8)$$

The same process is repeated M_i times for all time instances $i \in \mathcal{I}$.

- Files that cannot be placed within the 24 hours with LEO are downloaded using the terrestrial FSO links.

3. PERFORMANCE EVALUATION

The results presented here are based on the system model depicted in Fig. 1. The average number of edge BSs that are deployed per European region (K) is approximated by combining the territorial typology statistics reported in [13] for different European countries, and the average number of BSs deployed per km² for different territorial typologies in the UK [14]. This way, we assume an average density of mobile cells² of 0.0782 BS/km², which implies a coverage of 12.79 km² per cell, each of them with a 2-km average coverage radius.

The end-to-end data rate of the terrestrial backhaul network is bounded by the longest haul of the FSO multi-hop link. Without loss of generality, we assume that mean separation between adjacent cell sites is $d_{\text{fso}} = 4000$ m, which enables to transmit up to $R_{\text{fso}} = 1$ Gbps in case of clear sky conditions [6]. Note that this raw data rate is shared per cluster of 100 BSs, which are proposed to have a tree topology with one gateway connection to the core network. Finally, the data rate of the GEO and LEO satellite links are $R_{\text{geo}} = 50$ Mbps and $R_{\text{leo}} = 300$ Mbps, respectively, regardless of the weather conditions and the geographical location that each of the edge BSs takes.

Figure 4a shows the number of files of size $S = 40$ MB that can be placed on the different edge BS for different times. According to the previously presented analysis, $K = 23\,656$ edge BS are deployed on average per geographical region, requesting the placement of $M = 805.25$ files/BS belonging to a (universal) library with $N = 10\,000$ items. Then, when the $G = 35$ regions are considered, we have about 666.73×10^6 files to be placed in the whole European territory on a daily basis. As expected, the number of files that can be placed with LEO in the 24-hour window increases with the size of the constellation L , but this growth is not linear since the gain of multi-cast transmission decreases as the popularity of files to be placed decreases. Similarly, when comparing the GEO and LEO satellite technologies, it is possible to see that the wide pan-European coverage of the GEO satellite enables a better performance in long term, as a file needs to be transmitted only once

²For simplicity, we neglect the effect that large sea areas have on this calculation. Note that more precise studies should actually consider the territorial typologies of the countries found in each region g , complicating the analysis without changing the main claims of this paper.

regardless the shape of the aggregated popularity distribution (*i.e.*, the α parameter for the Zipf-distribution for Europe). In contrast, for LEO satellites, the narrower coverage that they have due to their reduced visibility enables a good performance only in the initial phase of placement, where the most popular content is distributed among adjacent regions with strong correlation in their top-ranked files to-be-placed. Note that the placing time of point-to-point FSO links grows linearly with time, as no multi-cast gain is reaped in this situation.

Finally, Fig. 4b illustrates the minimum delivery time that is feasible when both FSO terrestrial uni-cast and LEO/GEO satellite multi-cast transmissions are combined in the backhaul network to perform the file placing. In this figure, simulation results are presented for 1-LEO, 8-LEO, 1-GEO, and FSO. Satellite technologies start placing the most popular content in the geographical regions that are visible (lower horizontal axis), whereas the FSO link can be used to place the least popular content (upper horizontal axis). According to this figure, the placing time using the terrestrial FSO links is 7.2 hours, and can be reduced to 5.5 hours, 3.8 hours and 3.3 hours when 1-LEO, 1-GEO and 8-LEO are used. It is interesting to note that, even with the moderate LEO constellation sizes under analysis, where backhaul connectivity cannot be guaranteed 100% of the time, the performance gain that can be obtained when delivering content in non-real time is at least as good as the one provided by a much more expensive GEO satellite, which has a much more wider geographical coverage.

4. CONCLUSIONS

In this paper, we studied the reduction in placing time that is feasible when combining GEO/LEO satellite (radio) and terrestrial (optical) technologies to build the hybrid backhaul of a 5G network with edge caching. We considered that the whole European territory is divided into regions, which have popularity distributions of content that are correlated according to their geographical separation. Baseline orbital configurations, link budgets and scheduling algorithms were proposed to determine the order in which files should be transmitted using the different technologies (GEO, LEO and FSO), enabling the assessment of the required placing time for different hybrid solutions (GEO-FSO and LEO-FSO). As expected, notable performance gains were observed when combining satellite (LEO/GEO) and terrestrial (FSO) technologies to place content in the 5G edge. Moreover, even with small-sized LEO constellations, the time reduction that was achieved was similar to the one obtained with much more expensive GEO. Note that the integration of LEO satellites into 5G is aligned with the current launch schedule of more than 17000 LEOs, which will span the so-called “*new space*” mega-constellations.

ACKNOWLEDGMENTS

This work has received funding from the Ministry of Science, Innovation, and Universities of Spain under Project TERESA-TEC2017-90093-C3-1-R (AEI/FEDER, UE), the Catalan government under grant 2017-SGR-01479, the ERC Advanced Grant 695614 (POWVER), and the DFG Grant 389792660 as part of TRR 248.

REFERENCES

- [1] E. Bastug, M. Bennis, and M. Debbah, “Living on the edge: The role of proactive caching in 5G wireless networks,” *IEEE Commun. Mag.*, vol. 52, no. 8, pp. 82–89, 2014.
- [2] P. Wang, J. Zhang, X. Zhang, Z. Yan, B. Evans, and W. Wang, “Convergence of satellite and terrestrial networks: A comprehensive survey,” *IEEE Access*, vol. 8, pp. 5550–5588, Aug. 2020.
- [3] M. Luglio, S. Romano, C. Roseti, and F. Zampognaro, “Service delivery models for converged satellite-terrestrial 5G network deployment: A satellite-assisted CDN use-case,” *IEEE Network*, vol. 33, no. 1, pp. 142–150, Feb. 2019.
- [4] T. Vu, Y. Poirier, S. Chatzinotas, N. Maturo, J. Grotz, and F. Roelens, “Modeling and implementation of 5G edge caching over satellite,” *Int. Journal Satellite Commun. Netw.*, pp. 1–12, Jan. 2020.
- [5] S. Liu, X. Hu, Y. Wang, G. Cui, and W. Wang, “Distributed caching based on matching game in LEO satellite constellation networks,” *IEEE Commun. Letters*, vol. 22, no. 2, pp. 300–303, Feb. 2018.
- [6] A. Dowhuszko, M. Shaat, and A. Pérez-Neira, “Integration of optical and satellite communication technologies to improve the cache filling time in future 5G edge networks,” in *Int. Conf. Transparent Optical Networks*, July 2019, pp. 1–5.
- [7] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, “Web caching and Zipf-like distributions: Evidence and implications,” in *Proc. IEEE Int. Conf. Computer Commun.*, vol. 1, Mar. 1999, pp. 126–134.
- [8] D. Romani, “A popularity-based approach for the design of mobile content delivery networks,” Master’s thesis, Università degli studi di Padova, Oct 2013.
- [9] AGI, “Systems Tool kit (STK),” Apr. 2020, url: <http://www.agi.com/STK>.
- [10] J. Farserotu and R. Prasad, “A survey of future broadband multimedia satellite systems, issues and trends,” *IEEE Commun. Mag.*, vol. 38, no. 6, pp. 128–133, June 2000.
- [11] G. Minatti, S. Maci, P. De Vita, A. Freni, and M. Sabbadini, “A circularly-polarized isoflux antenna based on anisotropic metasurface,” *IEEE Trans. Antennas and Prop.*, vol. 60, no. 11, pp. 4998–5009, Nov. 2012.
- [12] J. Fraire, P. Madoery, and J. Finochietto, “Traffic-aware contact plan design for disruption-tolerant space sensor networks,” *Elsevier Ad-Hoc Networks*, vol. 47, no. C, p. 41–52, Sept. 2016.
- [13] Eurostats, “Territorial typologies,” Jan. 2019. [Online]. Available: https://ec.europa.eu/eurostat/statistics-explained/index.php/Territorial_typologies
- [14] E. Oughton and Z. Frias, “The cost, coverage and rollout implications of 5G infrastructure in Britain,” *Telecommunications Policy*, vol. 42, no. 8, pp. 636–652, Sept. 2018.