Received: 12 February 2018 | Revised: 26 May 2018 | Accepted: 6 June 2018
DOI: 10.1002/sat.1260

RESEARCH ARTICLE

WILEY

QoE enhancements on satellite networks through the use of caches

Adrien Thibaud[1,2] | **Julien Fasson**[1,2] | **Fabrice Arnal**[3] | **David Pradas**[4] | **Emmanuel Dubois**[5] | **Emmanuel Chaput**[1,2]

[1] TéSA, Toulouse, France
[2] ENSEEIHT/IRIT/University of Toulouse, Toulouse, France
[3] Thales Alenia Space, Toulouse, France
[4] Viveris Technologies, Toulouse, France
[5] Centre National d'Etudes Spatiales (CNES), Toulouse, France

Correspondence
Adrien Thibaud, TéSA, France.
Email: adrien.thibaud@tesa.prd.fr

Summary

A high throughput does not necessarily translate to a good Quality of Experience, especially in a satellite context. The round trip time, for instance, also has a tremendous impact on the reactivity of applications and thus on the Quality of Experience. Content delivery networks are massively used for over-the-top services in terrestrial network: They reduce the load of network and the delay as they draw the content closer to the end user. In a satellite system, the content delivery network presents a good opportunity for enhancing the end users' Quality of Experience and can change the conventional use of performance-enhancing proxies. This paper investigates the satellite as access link for home networks or a backhaul link for small cells for a 5G perspective.

We analysed the impact of caching on both gateway side and satellite terminal side for 2 on-trend services: web browsing and adaptive video streaming (dynamic adaptive streaming over HTTP). The main contribution is an evaluation of transparent caching through a satellite platform. The caching policy is out of the scope of this paper. One large part of the testbed is based on an open-source platform, OpenSAND that emulates the satellite system. To confirm the results, some real experiments have been conducted on a commercial satellite link.

As expected, the transparent caching at the satellite terminal side can increase the Quality of Experience to its upper border as long as the content is available in the cache. For the satellite gateway cache, the performances exceed the expectations. Although the application experiences the satellite delay in this case, the traversal time of different Internet service provider networks also delays the delivery of content. Then it may have a greater impact on reactivity than the satellite itself. Through careful analysis of the different results, we noticed some issues. Transparent caching is unable to cache encrypted or dynamic content. Moreover, a misuse of caching can provoke bad behaviour of dynamic adaptive streaming over HTTP mechanisms and severely decrease the Quality of Experience. We designed a solution that in addition to solving the issue, alleviates the storage space of satellite terminal caches.

KEYWORDS

caching, CDN, DASH, QoE, satellite

1 | INTRODUCTION

Recent communication improvements mainly target increasing bandwidth. However, high bitrate does not necessarily translate to a good Quality of Experience (QoE) for end users. Indeed, the Quality of Service (QoS) criteria that impact the QoE may differ according to the applications. For instance, a high latency impairs the QoE of web browsing users. Satellite networks effectively illustrate this principle. These networks offer major advantages: They can feed a very large number of users, propose a high bandwidth, and have low deployment costs. But they also have one significant disadvantage: Due to their location, users suffer a latency of around 500 ms (we focus on geostationary satellites in our study). This latency has a harmful impact in term of QoE. Satellite end users may be disappointed by their so-called high-speed Internet, when they experience a lack of responsiveness.

The competition between the increasing number of over-the-top services has actually led to an escalation towards the best QoE. For these services, a massive use of content delivery networks (CDN) is unavoidable. Content delivery networks solutions mainly rely on caching the content as close as possible to the end user and redirecting users' requests.[1]

Taking into account CDN capacities in the field of communications satellites (satcom) is a necessary step to improve users' QoE. As we will see, this deployment can match the competitiveness of terrestrial networks. Therefore, we propose in this paper to deploy a CDN in a satellite network and to study its impact on the experience of the end users. We set the study in a backhaul context where the satellite may feed home networks (HNs) or 5G small cells. Although we discuss numerous questions about the location of caches, the targeted applications, the type of stored content or the caching policy, we choose to focus on a proof of concept through measures based on an open-source testbed (OpenSAND) and a real satellite system. The results highlight an increase in QoE for both web browsing and video streaming. Nevertheless, the transparent caching for video streaming over HTTP, ie, dynamic adaptive streaming over HTTP (DASH), may induce poor performances in some cases. Thus, we propose a solution to improve the transparent caching and address cache misuses. The contribution of this paper are summarised as follows:

- We study the impact of the use of CDN on the QoE of satellite end users.
- We run our evaluation on an emulation and a real satellite testbed to emphasise our conclusions.
- We study several applications that reflect the current Internet usages (web browsing, video streaming and peer to peer [P2P]).
- We successfully solve a misuse of caching detected during the evaluation.

The remaining sections of this paper are organised as follows. In Section 2, we provide some background information and related work. Then Section 3 discusses the different solutions for deploying a CDN in a satellite context, and Section 4 presents the testbeds and the use cases that have been selected for this study. In Section 5, we evaluate the performances of the use cases. Finally, Section 6 concludes the paper and addresses areas for future work.

2 | RELATED WORK

Content delivery networks are commonly deployed in the Internet with major players like Akamai or Amazon. It consists in placing caches with replicated data at the edge of the network and redirecting users' requests to these caches.[1] Content delivery networks improve data accessibility and in turn user satisfaction.

The QoE represents the user's satisfaction for a particular service. The criteria depend on the application and are linked to the QoS. ITU-T recommends using the page load time as a QoE measurement for web browsing application.[2] ITU-T[3] and Fiedler et al[4] propose a relationship between QoE and QoS for these applications. The solution links the mean option score (MOS) to packet loss ratio, page load time, or delivery bandwidth. The authors explain that a user experiences a perfect QoE when the page is loaded in about 0.1 s or less, an acceptable QoE when the delay is about 1 s and loss of attention when the delay is above 10 s. They also claim that a user will be more understanding and patient if they expect a prolonged session time (for instance, with the use of geostationary satellites). Khorsandroo et al[5] define a relationship between MOS and packet loss ratio for video streaming application. However, for adaptive streaming, multiple encodings are available, and each encoded video is cut into chunks. In this case, the QoE depends on the quality of the received chunk and its download time. If this delay is too long, it may cause the playback to interrupt. This behaviour is commonly referred to as *stalling*.

Some studies deal with the QoE in a satellite system. For instance, Bisio et al[6] propose a compromise between transmitted power and QoE through a rate allocation algorithm of the satellite bandwidth: It tries to minimise the loss of QoE while maximising the gain of transmitted power. Although the performance-enhancing proxies (PEP) present several issues,[7] they are the main commercially used enhancer of QoE in satcoms. They split transport control protocol (TCP) connections in order to isolate the satellite part and accelerate them.[8] Performance-enhancing proxies also provide HTTP prefetching and many more tuning options. Nevertheless, their issues (complexity, security, and mobility) have led the satellite community to look for other solutions. One kind of solution consists in changing TCP.[9,10] But some of these modifications must be adopted by the whole Internet.[11] Another kind of solution is caching. Wu et al[12] propose a solution to minimise both the uplink and downlink bandwidth consumption. This solution consists of 3 parts: a combined utilisation of caches in the satellite and in the ground stations, an aggregation of the requests at the satellite level, and an optimised caching strategy. Their results show great improvement in satellite bandwidth consumption, but on-board caching is not yet mature. Brinton et al[13] use the satellite as an overlay on an existing terrestrial CDN. A central station aggregates the terrestrial requests to populate the cache of ground stations through the satellite point-to-multipoint communication. They propose a caching strategy that favours small and popular content.

Both papers mentioned propose new caching strategies. They focused on the multicast ability of satcoms and on bandwidth gain. However, the user's satisfaction and the location of caches are never discussed. Our work differs from the previous approaches because we focus on the QoE for real applications through emulations and a real satellite link. Moreover, our caches are directly part of the satcom, without interaction with another CDN.

3 | DEPLOYING A CDN IN A SATELLITE SYSTEM

A first step to enhance the QoE of end users in a satellite context consists in defining the targeted applications of a CDN. In fact, satcoms may assume different positions in a network, rendering many services, more or less linked to final users. Moreover, there are different types of CDN, with various caching policies and a broad offer of services. Because this work is part of a wider context, this section aims to outline this study.

First, to improve the experience of end users with CDN, we will focus on the application and thus the role of the satellite system. Then, we will discuss the caching functionalities and their placements. Eventually, we will conclude with a scenario for our study.

3.1 | Satellite role

The satellite system may be located in different places in the network. Whereas satcoms as backbone links could generally be considered as a pipe where end users' applications, and their QoE, are masked, their role as an access or backhaul network has the potential to be far from the pipe concept.

Years ago, satcoms were limited to niche applications, ie, expensive interconnections of remote areas and military usages. However, a real effort in standardisation based on digital video broadcasting, new codes and innovative modulations[14] have brought a lot of mainstream offers, especially for Internet access (Tooway, Viasat, Europasat, etc).

In an access context, the satellite directly feeds the HN of end users via a satellite terminal (ST) that provides functionalities similar to those of a home router. Satellite operators may be the Internet service provider (ISP) directly or may offer the connectivity to third-party ISPs. Moreover, the emergence of 5G offers a real opportunity to satcoms. Ultradense networks need to feed numerous small cells. The satellites are expected to offer part of these backhaul capacities.[15]

Apart from the mobility, these 2 kinds of scenario propose many similarities in terms of applications and network architectures. First, it can be considered as backhauling of small networks in both cases. They also share the same end users, with the same behaviours: They increasingly use smartphones and HTTP for web browsing, video streaming, or to retrieve data. Finally, with 5G, the HN might become a small cell like any other.

For this study, we have selected this specific use case: The satellite is placed as a backhaul link for small networks, ie, HNs or small cells. This context is motivated by good prospects for the future.

3.2 | Caching

Here, we will discuss where our caches are located and provide information about the caching policy used.

There are 2 reasonable places where a caching function could be deployed in a satellite network: at the gateway (Gw) side and at the ST side. Satellite payloads are not yet ready to deploy such on-board intelligence, especially in a geostationary context. The location of the caches conditions the opportunities available.

A Gw cache may provide an important storage space, and its location makes it possible to feed all the STs of the system via the multicast ability of the satellite.

An ST cache clearly makes it possible to reduce the delay because it avoids a satellite jump and is in close proximity to the end users. It could be an Internet box of a HN or an access point of a small cell. Therefore, an ST cache could offer the best acceleration for user applications. It could also save satellite resources, which is beneficial for the satcom operator.

But the most important questions for CDN are what should be cached and how to cache it?

The CDN can be independent from service providers and transparent to them. It can be directly managed by the satellite operator or the ISP. This solution offers flexibility but, in return, some kinds of data that are too dynamic or encrypted cannot be cached. Otherwise, the CDN can be an extension of the service provider or a CDN provider, but it only targets contents linked to the provider. A last solution consists in the interconnection with an existing CDN or several CDNs. The caches become an extension of another CDN, which may include the agreement for secured data. However, the interconnection is quite a difficult task due to the lack of standardisation in the domain, despite the CDNi initiative.[16]

Caches cannot contain the whole Internet content; they require a caching policy. It could simply be opportunistic caching: data (web pages, video chunks, etc) are cached as they pass through the satellite link. For instance, the ST cache has few end users available, especially in a HN;

thus, there is no need to store data that will probably not be watched a second time. However, this opportunity may be a real advantage for popular content at the Gw side. Another option consists in prefetching the data. Prefetching with a satellite system can be proactive: to download firmware or a film for future use. It can also be opportunistic, for instance, when an ST downloads a video on demand, other ST caches can store it if their network end users have common interests. This can also be used in a 5G scenario where the users are mobile.

For the rest of this paper, we will not focus on these points because this study addresses the advantages of CDN in a satellite context rather than cache management. It implies that the contents considered will be available in our caches except when explicitly stated otherwise. In some cases (that will be described later), only some of the contents considered will be available in our cache. We consider these cases to be consistent with the fact that our ST cache may have limited storage space.

3.3 | Scenario

The global scenario is described in Figure 1.

Because this work proposes a proof of concept, we examined a transparent caching solution (ST and Gw caches). For the HN, we only focused on 1 end user because the results are the same for any other user. We have studied 3 kinds of applications: web browsing, video streaming, and P2P file sharing.

These applications represent classic uses of Internet access for end users. For web browsing, we have selected different types of websites and measured the page load time.

Video streaming is based on DASH.[17] Dynamic adaptive streaming over HTTP proposes multiple encoding definitions for each chunk of a video, that we referred to as "qualities." The DASH player begins with the lowest quality. By estimating the available bitrate, based on the download time of a chunk, it adjusts its quality and requests the next corresponding chunk. We have measured the delay to obtain each chunk and their quality. Stalling may occur if the delay is prolonged. In this case, the QoE is considered unacceptable.

For the P2P application, we have used BitTorrent to download an open-source operating system (considered popular). We have measured the download time of the file.

4 | TESTBED AND USE CASES

In this study, we analyse the benefits of the CDN on the QoE of previously described applications. In this paper, we only present the use case and the results of web browsing and video streaming (convincing results for P2P application have also been produced).

In the first subsection, we introduce the 2 distinct satellite platforms we used for our tests. Then, in the next 2 subsections, we present the scenarios we used for the web browsing and the video streaming evaluations.

4.1 | The satellite platform

We conducted our tests on 2 separate satellite platforms: an open-source emulated platform named OpenSAND and a commercial Internet access by satellite, called Tooway on Eutelsat KA-SAT.

OpenSAND[18] is a user-friendly and efficient tool to emulate geostationary satellite communication systems. It provides a suitable and simple framework for performance evaluations and validations of new network and access mechanisms. OpenSAND allows us to deploy real applications on real equipment, providing a great demonstrator. We used OpenSAND to emulate a satellite, a Gw, and an ST (see Figure 1). The Gw is then connected to the Internet backbone, which as we will see, has an impact on the caching performances. We configure the satellite system to have

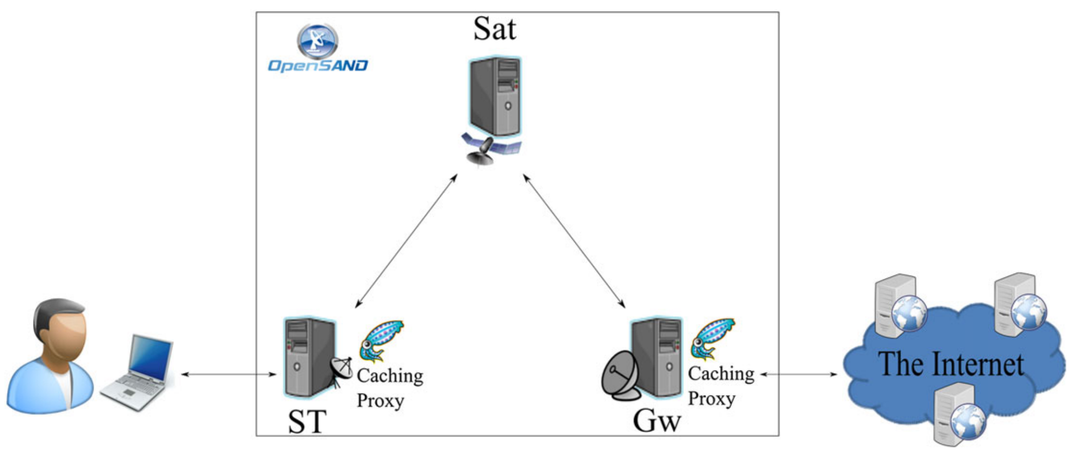

FIGURE 1 Satellite emulation testbed. Abbreviations: Gw, gateway; ST, satellite terminal [Colour figure can be viewed at wileyonlinelibrary.com]

a continuous rate assignment of 5 Mb/s for the return link and a rate of 20 Mb/s on the forward link. No PEP is used on the transport layer to accelerate traffic.

For the Tooway offer, the characteristics are 22 Mb/s of download speed; 6 Mb/s of upload speed; and 10 Gb of data on a daily basis. This commercial flat fee includes a TCP acceleration through PEP, so it allows us to compare the performances with and without a PEP. The client only has access to the satellite terminal. Thus, we are only able to manage a cache at the ST side. Tooway is only used for the web browsing application (for validation purposes), because we want to confirm our results on a real satellite system.

4.2 | Web browsing

For this application, the W3C[19] presents an application programming interface (API) able to measure the different events that occur during the page load time. Figure 2 summarises the main events that the API provides:

- requestStart, which corresponds to the emission of the first HTTP request to the server (HTTP GET);
- responseEnd, which corresponds to the end of the reception of the response of the server (response of the first HTTP request); and
- loadEventEnd, which corresponds to the time where the page is fully loaded on the user's browser.

In terms of QoE, the user experiences the loadEventEnd. This measure is consistent with the state of the art (see Section 2).

The network initialisation includes the domain name system (DNS) resolution, the potential redirection, and the TCP 3-way handshake. In our case, these events are independent from caching (DNS cache can be used, but it is not studied here). So, we chose requestStart as our origin for the evaluations.

On the Internet, the websites have their own characteristics. We chose to run our test on 3 websites that we think accurately represent the Internet as a whole. We used a static website with a few images, another static website featuring HTTPS elements and a dynamic website that proposes a lot of images. The detailed description is available in Table 1.

For each website, we measured the page load time when there is no cache (as a reference), with the cache on the Gw and with the cache on the ST. The software Squid[20] is used to cache the content (see Figure 1). We ran our simulations a sufficient number of times (30 times) to generate a meaningful confidence interval.

4.3 | Video streaming

For this application, we used the DASH-JS player.[17] Multiple qualities of a video are available. The client estimates the available rate and asks for the next chunk of the corresponding quality. But these kinds of algorithms are optimised for terrestrial networks. The estimation of the available rate is based on the download time of the previous chunk, which includes the round trip time. Due to the GEO satellite delay, it underestimates its bandwidth. It may induce a poor quality for satellite users. In this scenario, we show the benefits of CDN and the drawbacks in case of misuse.

We have used the Big Buck Bunny video,[21] which was sliced into 150 chunks of 4 s each. Four seconds was considered a good trade-off on satellite networks: The chunks are short enough to quickly adapt to network variations and sizable enough to provide a measurement of the download time. The available qualities are listed in Table 2.

The lower qualities are useful for mobile users, for example. They may have a low connection and most of all they have small screens and do not require the highest qualities to achieve a good QoE. However, a high quality will not affect their QoE provided the chunks arrive in a reasonable amount of time. In the worst case scenario, it will waste some network bandwidth. We measured the delay to get every chunk as well as the quality of the chunks downloaded by the user. The client has a buffer of 10 chunks. When the buffer is empty, the video display stops. If the time to retrieve the chunks is greater than 4 s, the buffer will run out and, at some point, stalling will occur. The parameters are a constraint of the

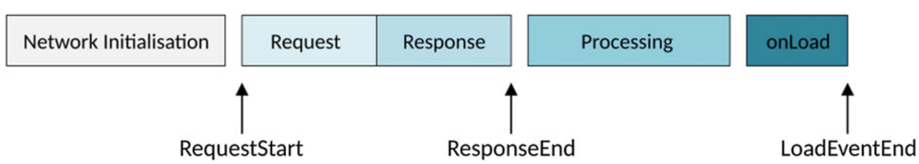

FIGURE 2 Page load time [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 1 Website description

Website	Characteristics
http://opensand.org	Static page, hosted in France, no javascript, 2 images. Size: 396 kB
http://www.hereistoday.com	Static page, hosted in the US, javascript, HTTPS elements (twitter, facebook). Size: 399 kB
http://www.viveris.com	Hosted in France, javascript, lots of images, elements from Google Analytics. Size: 1000kB

TABLE 2 Qualities available

Quality	Associated Rate
Q1	200 kb/s.
Q2	500 kb/s.
Q3	1000 kb/s.
Q4	2000 kb/s.
Q5	3000 kb/s.

DASH client that the user downloads on the server. Users do not have control of these parameters that are designed to perform well on usual networks (nonsatellite networks).

For this application, we do not use the QoE measurement presented in the state of the art because it is not appropriate for adaptive video streaming (see Section 2). We use a combination of criteria to estimate the QoE. The first criterion is the quality of the video: The better the quality of the video, the better the QoE of the user will be. However, for those using mobile devices, which have smaller screens, the QoE will be at its maximum with middle qualities, and low qualities will represent a decent QoE. The second criterion is stalling (its duration and its frequency). If the video freezes frequently, even for a short time, we assume that the QoE will decrease significantly. Similarly, if the duration of the stalling is too long, the QoE will be downgraded. The third criterion is the frequency of video quality adaptation (switching). Indeed, the user could get bored if the video quality is switched too frequently. Stalling and switching issues are both directly linked to the retrieving duration of chunks. We will see in our simulations that the switching issue always leads to stalling. Thus, we only consider this last criterion. A final criterion could be the start-up delay of the video playback, but it is not relevant in our case because the user already expect higher delay on a satellite connection (as we explained in Section 2). For these reasons, we make this supposition: in case of one-time stalling (and assuming it does not last too long), the QoE of the user will not be downgraded; however, if stalling occurs multiple times, the impact will be highly detrimental. The QoE could have been specified in MOS. In a future work, we plan to map our QoE criteria to MOS values though a student pool.

We propose multiple subscenarios in order to analyse the impact of CDN in a satcom environment in QoE terms. In our first scenario S1, we will show the results when there is no cache (as a reference), a cache on the Gw and eventually a cache on the ST. In our second scenario S2, the qualities are distributed between the caches according to Table 3. As explained in Section 3, the storage space of the ST cache may be limited. To take this into account, we only prefetched some qualities in the ST cache. The distributions of qualities are chosen to show all the major impacts of a CDN on the users' QoE.

Finally, in our third scenario S3, we propose a solution that solves the misuses detected in the previous subscenario.

We deployed a DASH server in Toulouse with an upload rate of 1.4 Mb/s and a delay from the server to the testbed site of around 100 ms. As we will see, the users are able to get the best quality available with the help of CDN, even if the upload rate is not sufficient.

5 | PERFORMANCE EVALUATION

This section presents the results of our experimentations on web browsing and adaptive video streaming.

5.1 | Web browsing

The results of the web browsing evaluation on the OpenSAND platform are presented in Figure 3A and consist of the measure of the 3 events defined in Section 4.2. Numerical results are summarised in Table 4.

Without the cache (solid curve), the page load time is very high (up to 30 s for the *viveris* website). This is due to the low reactivity of the satellite connection. The users experience a very bad QoE in this case.

TABLE 3 Use case Scenario 2

	Qualities in Gw	Qualities in ST
S2a	all	Q1 ; Q5
S2b	all	Q1 ; Q4
S2c	all	Q5
S2d	none	Q1 ; Q5
S2e	none	Q5
S2f	none	Q1 ; Q4

Abbreviations: Gw, gateway; ST, satellite terminal.

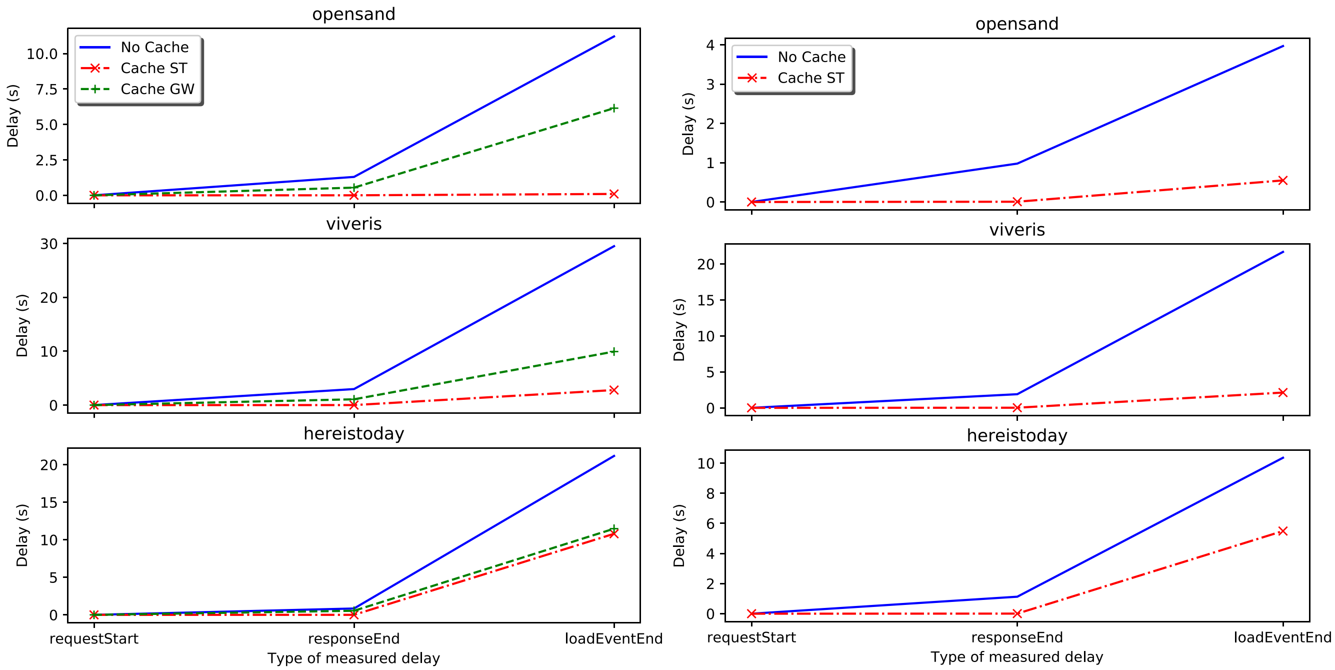

(A) Web browsing with the OpenSAND platform (B) Web browsing with the KA-SAT platform

FIGURE 3 Web browsing results. Abbreviations: Gw, gateway; ST, satellite terminal [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 4 Web browsing results

Plateform	Website	Scenario	Mean loadEventEnd, ms	Standard Deviation loadEventEnd, ms
OpenSAND	opensand	No Cache	11 210	1490
		Cache Gw	6149	106
		Cache ST	97	4
	viveris	No Cache	29 502	2446
		Cache Gw	9935	194
		Cache ST	2793	154
	hereistoday	No Cache	21 169	2552
		Cache Gw	11 477	408
		Cache ST	10 800	565
KA-SAT	opensand	No Cache	3969	695
		Cache ST	549	51
	viveris	No Cache	21 678	636
		Cache ST	2125	197
	hereistoday	No Cache	10 361	145
		Cache ST	5492	340

Abbreviations: Gw, gateway; ST, satellite terminal.

Except for the *hereistoday* website, the results with the cache on the ST (dotted curve) are excellent (only a few seconds for the *viveris* website and nearly instantly for the *opensand* website) and show the best case of utilisation for the CDN. Those with the cache on the Gw (dashed curve) show great enhancements too, around 50% gain. Surprisingly, the use of a CDN on the Gw is already a huge benefit in terms of QoE for the user. In fact, it only suffers the satellite delay and not the backbone delay.

For the *hereistoday* website, the results with the cache on the Gw and those with the cache on the ST are similar and show an important increase in the page load time between the respondEnd event and the loadEventEnd event. The reason is the presence of HTTPS elements, which cannot be cached without the explicit authorisation of the owner of the content. However, those results are still better (around 50%) than those without a cache because the end user still benefits from the caching of HTTP elements.

The results of the web browsing evaluation on the KA-SAT platform are presented in Figure 3B. As explained in Section 4, the cache is only at the ST side.

The results have the same shape as those obtained on the OpenSAND platform. The previous conclusions are then confirmed. Moreover, we can observe an enhancement on the global results, even the results without cache. Part of this improvement is induced by the PEP acceleration on this commercial system. One extra conclusion can be highlighted with these results: Even when a PEP is used, caching in the satellite system is beneficial to the QoE.

5.2 | Video streaming

Numerical results are summarised in Table 5. Figure 4 presents the results of S1.

As expected, without a cache (solid curve), the user is not able to get the quality in line with its bandwidth (approximately 1.4 Mb/s here, which is the maximum for the upload rate of our DASH server). The delay (Figure 4A) to retrieve a chunk is simply too long. Thus the algorithm considers the download rate as insufficient for better quality (Figure 4B). Some delays, between the 80th and the 100th fragments, are longer than the duration of a chunk itself (4 s). Even with the buffer of 10 chunks, a user may experience some stalling, which in turn has a negative impact on the QoE.

The use case with a cache on the Gw (the dashed curve) highlights the first noteworthy result. After a convergence phase (due to the slow start of TCP on the satellite link), the user is able to obtain the best quality directly from the Gw cache. The delay to get each chunk is always below 4 s; thus, it ensures that no image will freeze.

Finally, the results when the ST is caching the content are excellent. As expected, the DASH client appraises the bandwidth as very good and asks for the best quality. In fact, the time to retrieve any fragment is negligible.

To conclude this first scenario, we see the inefficiency of the DASH algorithm on a satellite network (case without CDN). The use of a cache on the Gw side allows the user to experience the best QoE after an initialisation phase (no stalling and optimal quality). Finally, with the ST cache, the QoE is optimal from the first chunk, and no satellite resources are used. This first scenario serves as a reference: We now have a calibration of the performances of a cache depending on its location. We will use it in order to analyse more conceivable scenarios where the ST cache is limited by the storage capacity due to its role as an Internet box.

The first results of S2 (a, b, and c) are presented in Figure 5 (the delays are shown in a and the quality in b). In this scenario, the Gw cache has all the qualities but the ST cache has only a few in order to save storage space. If the user requests a chunk that is available on the ST cache, the ST will send it. If the chunk is not available on the ST cache, the request will reach the Gw, which stores all the available quality in its cache.

The curve of scenario S2a (dotted curve) equates to the one with the cache on the ST from scenario 1. Indeed, it always gets chunks from the ST, the first one with Q1 quality, and the others of Q5 quality. It is a first significant savings on the storage capacity of the cache.

TABLE 5 Video Streaming results

			Mean Quality, kb/s	Standard Deviation Quality, kb/s	Mean Delay, ms	Standard Deviation Delay, ms
S1	No Cache		290	136	2291	929
	Cache Gw		2858	7	1381	296
	Cache ST		2992	10	113	8
S2	a	Gw(all) - ST(Q1;Q5)	2992	9	112	8
	b	Gw(all) - ST(Q1;Q4)	2980	3	1546	451
	c	Gw(all) - ST(Q5)	2812	28	226	41
	d	ST(Q1;Q5)	2992	9	112	8
	e	ST(Q5)	306	124	2276	832
	f	ST(Q1;Q4)	1531	979	11 757	9334
S3	ST(Q1;Q5) with proxy MPD		1998	5	91	3

Abbreviations: Gw, gateway; MPD, media presentation description; ST, satellite terminal.

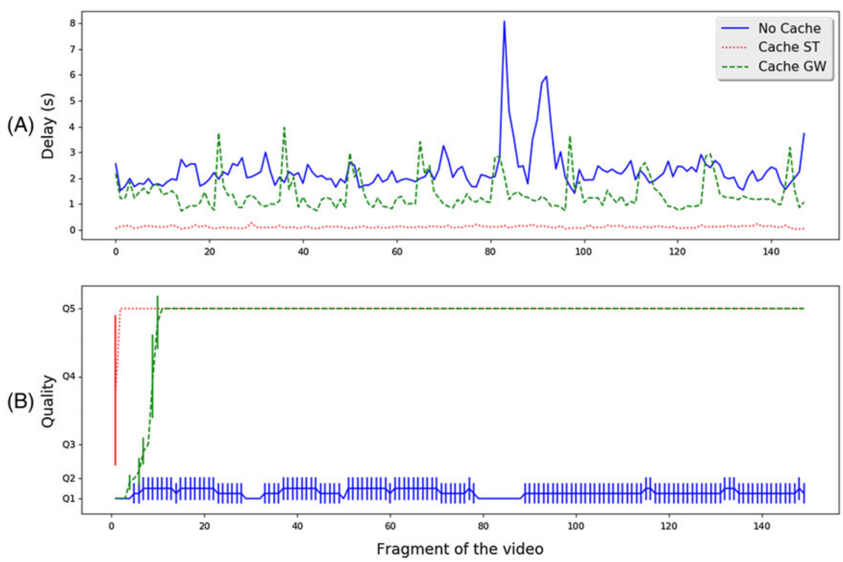

FIGURE 4 Video streaming—Scenario 1. Abbreviations: Gw, gateway; ST, satellite terminal [Colour figure can be viewed at wileyonlinelibrary.com]

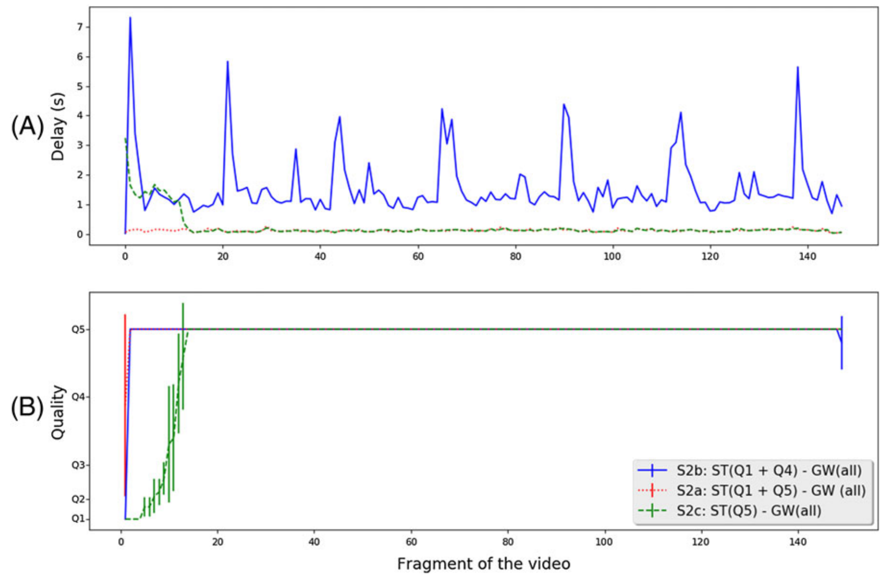

FIGURE 5 Video streaming—Scenario 2a, b, and c. Abbreviations: Gw, gateway; ST, satellite terminal [Colour figure can be viewed at wileyonlinelibrary.com]

The scenario S2b (solid curve) shows the first misuse of CDN. In this case, the DASH client will evaluate a high rate from the first chunk because it is obtained from the ST cache. Then, the client asks for the best quality, Q5. However, it is not available on the ST, so it must get it from the Gw cache. As we have seen with scenario 1, the Gw is able to supply the user with chunks of the best quality without causing any freezing images. The QoE of the user is not downgraded. Nevertheless, from the satellite operator point of view, caching does not improve its bandwidth usage, nor does it reduce the user flat fee. Furthermore, the cache on the ST possesses all the chunks for the quality Q4, which is never used (waste of storage space).

For scenario S2c (dashed curve), the figure of requested qualities is similar to scenario 1 (Gw case); thus, it has a comparable QoE. In addition, from the 18th chunk, the client requests Q5 quality and obtains it directly from the ST cache. This represents a savings on the satellite link.

In conclusion, to maximise QoE, there is no need to put middle qualities in the ST cache if the best quality is already available. Moreover, such a policy makes it possible to save storage space on the ST.

The second results of S2 (d, e, and f) are presented in Figure 6. These scenarios are the same as S2 (a, b, and c) without a Gw cache.

The S2d results are the same as S2a and S1 because only the ST cache is used. It confirms our previous conclusions: Middle qualities in the ST cache and a cache in the Gw are not needed when the best quality is stored.

In S2e, Q5 is available on the ST. However, the DASH client cannot notice its presence because it begins with Q1 located at the original server. Therefore, it evaluates a poor download rate. In this case, to store only the Q5 on the ST is a pure waste of space.

S2f shows the worst misuse of the CDN in our context. In this subscenario, the ST proposes Q1 and Q4. Because the user begins with Q1, he gets the first chunk from the ST. The DASH client evaluates its download rate as excellent. So the user asks for a Q5 chunk; however, only the

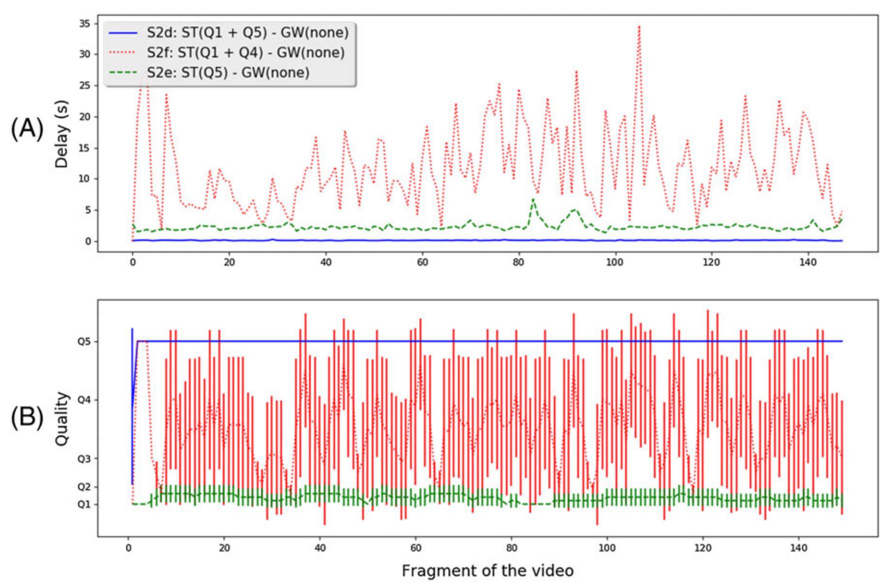

FIGURE 6 Video streaming—Scenario 2 d, e, and f. Abbreviations: Gw, gateway; ST, satellite terminal [Colour figure can be viewed at wileyonlinelibrary.com]

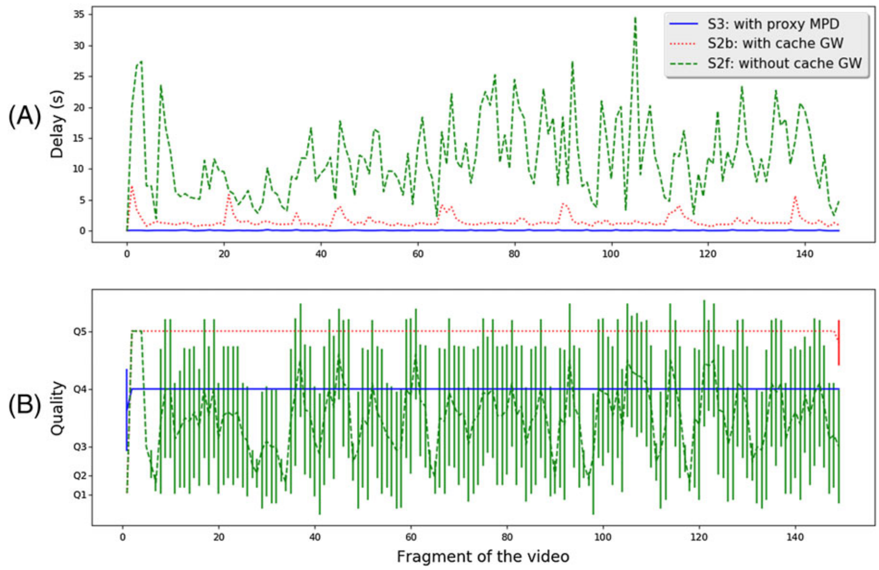

FIGURE 7 Video streaming—Scenario 3. Abbreviations: Gw, gateway; ST, satellite terminal [Colour figure can be viewed at wileyonlinelibrary.com]

original server possesses this quality. It takes approximately 20 s to get this chunk. The client measures a poor download rate. The requested quality will then oscillate between Q1 and Q5 throughout the video streaming. Furthermore and as explained in Section 4, this delay implies that the user will experience major issues with freezing images while watching the video, making him lose interest in the content. This use of the CDN puts the QoE of the user at its lowest level.

We propose through a last scenario (S3) to improve the poor behaviour of DASH in satellite context, as it has been observed in the previous scenario. We have analysed the cause of this bad behaviour. In a satellite context, the RTT is not directly linked to the available bandwidth; as a result, it may lead the DASH client to underestimate the bitrate. The previous scenario (S2) has revealed issues when only a small set of qualities is stored in the ST cache. A first problem appears when the lowest quality is not available. The user will never evaluate its bandwidth as sufficient to get a higher quality: He will only ask for the lowest one. The chunks will transit from the original server or another cache through the satellite link (as highlighted in S2e). The second problem is the most problematic, and it occurs when the ST cache does not store the highest quality. The DASH client gets the first chunks of the video from the ST cache with a short delay. It estimates a high bandwidth and asks for the highest quality chunks. As we have seen earlier (in S2f), the user gets the chunks from the server after a significant waiting period, stalling occurs, and the QoE is very low. Here, these 2 problems are due to the same cause: Clients are not aware of the video qualities cached by the ST. The solution we propose is to only inform the user of the available qualities in the ST store. Unfortunately, the DASH client used by the user is provided by the server and is optimised for terrestrial users. It is not reasonable to think that servers will change their implementations only to accommodate a minority of users like satellite users. Thus, we propose to change the manifest (media presentation description [MPD]) locally and transparently. In the DASH algorithm, this manifest is used to describe the video: the number of chunks and the available qualities. The manifest is available at the server side and can be cached. To prevent the issues highlighted in the previous scenario, we propose to modify the MPD in the ST cache. The cache in the ST will act as a transparent proxy and will change the MPD to only advertise the available qualities in the cache. Then, the DASH client will only be able to ask for these qualities, preventing the algorithm from misjudging the available bandwidth. The user gets the fragment of the video quickly and never experiences freezing images.

Figure 7 compares 3 cases: S2b (dotted), S2f (dashed), and S3 (solid), which is the same as S2f with our modified MPD.

The user gets Q4 chunks and does not experience freezing images. This proposition does not provide the best quality, Q5, and so the QoE may not be optimal. However, the QoE is far better than that for S2f and the solution without CDN. Furthermore, it saves satellite bandwidth. Finally, Q1 is not even mandatory in the modified manifest because it was only used to initiate the streaming. It makes it possible to spare more space in the ST cache and possibly to propose Q5 directly.

This solution is a workaround for satellite operators to prevent the bad behaviour of the DASH algorithm we highlighted in our scenario S2f. However, we have to mention that it breaks the adaptive behaviour of DASH and its ability to increase the quality if the link allows it. But as seen in scenario S2f, due to the high latency induced by the satellite, users are not able to get the better quality anyway.

6 | CONCLUSION

In this paper, we have proposed to use caches in satellite systems in order to improve users' satisfaction. First, we located the caches at either the Gw side or the ST side. We then compared the performances for both cases on 2 applications: web browsing and adaptive video streaming

(DASH). For these 2 applications, the ST cache shows the optimum performances in term of QoE because the requested objects are the nearest to the users. For web browsing, the page load time is almost instantaneous when there is no encrypted content. For adaptive video streaming, the best quality chunks are obtained immediately, and the playback is perfect. Nevertheless, 2 exceptions have to be highlighted. First, transparent caching (including for Gw caches) is ineffective for encrypted content. Secondly, in some cases of adaptive video streaming, experiments have revealed misuses of the ST cache. Indeed, the limited storage capacities of the ST restrict the number of video qualities available in the cache. Therefore, the cache may include an incorrect set of video qualities that will have a strong negative impact on the user's QoE. We introduce a proposition: a local modification of the manifest that successfully solves this issue. Although a cache at the Gw side does not prevent the satellite delay, it notably improves the QoE. The page load time is reduced by almost 50% for web browsing, and the best quality chunks are obtained fast enough for a perfect playback of the video (but after an initialisation phase) for adaptive video streaming. These results are mitigated by the cost of satellite resources.

All these results have been generated without simulation. All traffic has actually crossed part of the Internet to reach the original servers. The use of both an emulated and a real satellite platform strengthen our results and conclusions. As a result, this study has been used by Thales Alenia Space and CNES to advance CDN strategy into the satcom ground segment in order to better integrate terrestrial networks.

On one hand, encrypted communications are increasingly common. On the other hand, they impede effective caching functionalities. A trail that may combine both consists in the new information-centric network paradigm. Among information-centric network solutions, named data networking[22] (NDN) seems promising and meets the previous expectations. Unlike the classic Internet, caching is native in NDN. Some first analyses[23] show better performances than CDN. Moreover, the content encryption does not depend on the communicating hosts. The integrity of each content is independent from the source and the destination. It enables nodes to cache the encrypted data. Therefore, our future studies will focus on NDN integration in a satellite context in order to further enhance the QoE of end users.

ACKNOWLEDGMENT

The authors would like to thank the CNES (French spatial agency) and, particularly, the CESARS team for their help during the real satellite environment tests.

ORCID

Adrien Thibaud http://orcid.org/0000-0003-2009-5695

REFERENCES

1. Buyya R, Pathan M, Vakali A. *Content Delivery Networks*. 1st ed. Berlin: Springer Publishing Company, Incorporated; 2008.
2. ITU-T Rec. G.1031. QoE factors in web-browsing; 2014.
3. ITU-T Rec. G.1030. Estimating end-to-end performance in IP networks for data applications; 2014.
4. Fiedler M, Hossfeld T, Tran-Gia P. A generic quantitative relationship between quality of experience and quality of service. *IEEE Netw*. 2010;24(2):36-41.
5. Khorsandroo S, Noor R, Khorsandroo S. A generic quantitative relationship between quality of experience and packet loss in video streaming services. In: 2012 fourth international conference on ubiquitous and future networks (ICUFN); 2012; Phuket, Thailand:352-356.
6. Bisio I, Delucchi S, Lavagetto F, Marchese M. Transmission rate allocation over satellite networks with Quality of Experience-based performance metrics. In: 2014 7th advanced satellite multimedia systems conference and the 13th signal processing for space communications workshop (ASMS/SPSC); 2014; Livorno, Italy:419-423.
7. Dubois E, Fasson J, Donny C, Chaput E. Enhancing TCP based communications in mobile satellite scenarios: TCP PEPs issues and solutions. In: 2010 5th advanced satellite multimedia systems conference and the 11th signal processing for space communications workshop; 2010; Cagliari, Italy:476-483.
8. Border J, Kojo M, Griner J, Montenegro G, Shelby Z. RFC 3135: Performance enhancing proxies intended to mitigate link-related degradations. Internet Engineering Task Force; 2001.
9. Dayma R, Vanzara R. Improved tcp hybla: A TCP enhancement for link with high RTT and error rate. In: 2012 NIRMA university international conference on engineering (NUICONE); 2012; Ahmedabad, India:1-6.
10. Kim J-M, Lee S-Y, Kim J-H. TCP congestion window tuning for satellite communication using cross-layer approach. In: 2016 eighth international conference on ubiquitous and future networks (ICUFN); 2016; Vienna, Austria:167-169.
11. Sallantin R, Baudoin C, Chaput E, Arnal F, Dubois E, Beylot AL. Initial spreading: a fast start-up tcp mechanism. In: 38th annual IEEE conference on local computer networks; 2013; Sydney, NSW, Australia:492-499.
12. Wu H, Li J, Lu H, Hong P. A Two-layer caching model for content delivery services in satellite-terrestrial networks. In: 2016 IEEE global communications conference (GLOBECOM); 2016; Washington, DC, USA:1-6.
13. Brinton C, Aryafar E, Corda S, Russo S, Renoso R, Chiang M. An intelligent satellite multicast and caching overlay for CDNs to improve performance in video applications. In: 31st AIAA International Communications Satellite Systems Conference (ICSSC); 2013; Florence, Italy.
14. Digital video broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for broadcasting, interactive services, news gathering and other broadband satellite applications (DVB-S2). ETSI EN 302 307; 2009.
15. Backhaul technologies for small cells: use cases, requirements and solutions. Small Cell Forum - Document 049.01.01; 2013.
16. Niven-Jenkins B, Faucheur FL, Bitar NN. Content distribution network interconnection (CDNI) problem statement. RFC 6707; 2012.
17. DASH Industry Forum. http://dashif.org/
18. OpenSAND website. http://opensand.org

19. World Wide Web Consortium. Navigation timing. https://www.w3.org/TR/navigation-timing/#processing-model
20. Squid, Optimising Web Delivery. http://www.squid-cache.org/
21. Big Buck Bunny. https://peach.blender.org/
22. Zhang L, Afanasyev A, Burke J, et al. Named data networking. *SIGCOMM Comput Commun Rev*. 2014;44(3):66-73. https://doi.acm.org/10.1145/2656877.2656887
23. Ma G, Chen Z, Cao J, Guo Z, Jiang Y, Guo X. A tentative comparison on CDN and NDN. In: 2014 IEEE international conference on systems, man, and cybernetics (SMC); 2014; San Diego, CA, USA:2893-2898.

Adrien Thibaud graduated from INP-ENSEEIHT, Toulouse, France, in 2014. From 2014 to 2017, he worked as a network engineer at Viveris Technologies. Since 2017, he is a PhD student at the University of Toulouse, under a Thales/TéSA PhD grant. His work so far has been on the transport layer in the Information Centric Networks.

Julien Fasson is an assistant professor at University of Toulouse in the engineering school named ENSEEIHT. After a master in aeronautical engineering, he focused his studies on the computer networks issues. He received his PhD of Computer Science and Telecommunications in 2004. Since, he is working at IRIT laboratory, and he participates to TéSA lab too. His main interests are all focused on the integration of telecom satellite systems in terrestrial networks. This includes architecture of services as IP Multimedia Subsystem or Content Delivery Network, overlay solutions like Information Centric Network, transport protocol, link layer adaptation, signaling, on-board-processing, and so on.

Fabrice Arnal is research engineer at Thales Alenia Space (Toulouse); he joined in 2006, in the “R&D system telecom” team. He is currently work package manager of the Sat5G H2020 EU funded project, whose objective is to study and demonstrate satcom integration in 5G networks. His main fields of interests and expertise are related to satcom integration and convergence in terrestrial networks, transport layer, multicast, video streaming over HTTP, Quality of Service, access layer, and cross-layering architecture.

David Pradas received a telecommunications engineering degree from Polytechnic University of Catalonia (UPC) in 2006 and a joint PhD degree in satellite communications from ISAE-SUPAERO (Toulouse) and Autonomous University of Barcelona (UAB) in 2011. He is now a research engineer and the “R&D network/telecom” technical team lead at Viveris Technologies, which he joined in 2012. His main interests focus on satellite telecommunication networks/architectures, including satcom/terrestrial hybrid systems and multihoming, cross-layer optimization, SDN/NFV virtualization frameworks, transport layer impact, QoE/QoS performances, as well as satellite network emulation/simulation.

Emmanuel Dubois is a Network & Telecom research engineer. He joined Centre National D'Etudes Spatiales in Toulouse in 2008. He holds a PhD dealing with convergence in satellite networks. His main topics of interest include therefore network architecture convergence, cross-layering optimization and also transport layer, QoS, and resource management. He is now working on several tools for research and development in satcom topics from an open source satellite emulation platform called OpenSAND to a multipurpose metrology tool called OpenBACH.

Emmanuel Chaput received an engineering degree in computer science from ENSEEIHT, Toulouse (France) in 1992 and a PhD degree in computer networks from the Université Paul Sabatier in 1998. He is now professor at the Telecommunication and Network Department, Toulouse INP/ENSEEIHT and member of the CNRS/IRIT RMESS team (networks, mobiles, embedded, wireless, and satellites). His current research interests include IP over satellite networks architecture and vehicular communications. Within the scope of such application areas, he focuses on network architecture and resource allocation.

How to cite this article: Thibaud A, Fasson J, Arnal F, Pradas D, Dubois E, Chaput E. QoE enhancements on satellite networks through the use of caches. *Int J Satell Commun Network*. 2018;1–13. https://doi.org/10.1002/sat.1260