A Virtual PEP for Web Optimization over a Satellite-Terrestrial Backhaul

Published Version:
A Virtual PEP for Web Optimization over a Satellite-Terrestrial Backhaul / A. Bujari, M. Luglio, C. E. Palazzi, M. Quadrini, C. Roseti, F. Zampognaro. - In: IEEE COMMUNICATIONS MAGAZINE. - ISSN 1558-1896. - ELETTRONICO. - 58:10(2020), pp. 42-48. [10.1109/MCOM.001.2000322]

Availability:
This version is available at: https://hdl.handle.net/11585/792056 since: 2021-03-01

Published:
DOI: http://doi.org/10.1109/MCOM.001.2000322

Bujari, M. Luglio, C. E. Palazzi, M. Quadrini, C. Roseti and F. Zampognaro, "A Virtual PEP for Web Optimization over a Satellite-Terrestrial Backhaul," in IEEE Communications Magazine, vol. 58, no. 10, pp. 42-48, October 2020

A Virtual PEP for Web Optimization over a Satellite-Terrestrial Backhaul

Armir Bujari, *Member, IEEE,* Michele Luglio, Claudio E. Palazzi, *Member, IEEE,*
Mattia Quadrini, Cesare Roseti and Francesco Zampognaro

***Abstract*—The availability of network softwarization and virtualization technology in the field of telecommunications has opened the door to a radical review of the applications, protocols and deployment models. In this evolving framework, old assumptions and constraints specific to satellite communications must be carefully re-assessed. To this aim, we revisit the role of the Performance Enhancing Proxy (PEP), replaced by a chain of custom Virtual Network Functions (VNFs) properly enabled to optimize common Web traffic performance over a backhaul dynamically enabled with a supplementary satellite link. The resulting virtual PEP (vPEP) is compliant with the breakthrough virtualization and slicing paradigms and can fruitfully exploit the advanced features of the most recent IETF technologies such as QUIC and MPTCP.**

***Index Terms*—Satellite, Virtualization, virtual PEP.**

I. INTRODUCTION

Satellite communications present physical and configuration characteristics different from those of the terrestrial wired and wireless technologies, making the adaptation of current Web technologies, often designed with the latter in mind, critical. The main factors impacting performance are: the high latency, the dynamic bandwidth assignment on the satellite interactive channel introducing possible variable delay contributions, channel fading leading to bursty losses, etc. [1]. Yet, the satellite has proven to be a viable complementary alternative to terrestrial networks, capable of providing a cost-effective coverage in crowded and unserved areas [2].

Addressing the aforementioned issues, different approaches have been proposed in the scientific literature as well as in the commercial products, enabling the satellite as a complementary technology for Internet based services [3]. In particular, the Performance Enhancing Proxy (PEP [4]) approach introduces architectural changes whereby specialized agents are deployed at the edge of a satellite link, mainly performing TCP splitting operations aimed to use an ad-hoc transport protocol for the IP data transmission over satellite.

The traditional PEP approach strives to implement a protocol isolation of the satellite segment by converting standard TCP/IP protocols with specialized protocols and/or configurations, aimed to mitigate the impairments of the satellite on standard applications.

A drawback and main obstacle of this integration approach is the need to install PEP hardware-legacy middleboxes, tailored to optimize the performance of specific data flows. Furthermore, techniques such as TCP splitting, often used to improve the performance of satellite based communications, have inherent limitations on security implementations due to the violation of the end-to-end semantics and consequent incompatibility with the main network security paradigms [5]. More in general, hardware-legacy solutions cannot follow the speed of the IETF technology evolution, making optimizations narrowed in time and specific to current protocols.

Fortunately, new technology enablers such as virtualization and network softwarization have appeared in the network domain and can be a game changer even for satellite scenarios. In this framework, the PEP can be decomposed into a set of specialized Virtual Network Functions (VNFs) fulfilling a set of different requirements. In particular, distributed VNFs allow to support paired functions as required in the traditional PEP paradigm. These functions can be hosted and deployed on off-the-shelf hardware, reducing operational costs, leveraging Network Function Virtualization (NFV) architectures [6]. The possibility to dynamically configure a number of per-service, per-flow VNFs enables the much needed flexibility to target the performance requirements of heterogeneous applications and to account for possible protocol updates.

In the envisioned approach, the resulting virtualized PEP (vPEP) can be also tasked to dynamically enable the satellite link in case of necessity (i.e. outage of the terrestrial link) or convenience (i.e. traffic offloading during congestion). Specifically, users could require some enhancements to the basic connectivity service provided by the terrestrial network operator, such as an efficient broadcast delivery, a coverage extension, extra-capacity (bonding), failover configurations, etc. Thanks to the vPEP, the satellite segment can be enabled on-demand, with users being generally unaware of its presence, enjoying an enhanced Quality of Service/Experience (QoS/QoE).

The proposed vPEP can be particularly beneficial for the Web communications, where the protocol heterogeneity, the different application requirements and the time-varying traffic patterns require a flexible and diversified per-flow management.

The reminder of the article is organized as follows: Section II provides a concise review of the PEP architecture and optimizations. Section III outlines a vPEP configuration tailored to Web traffic acceleration while Section IV presents the emulation testbed used to assess our proposal. Section V is dedicated to the experimental validation and their outcome. Finally, Section VI draws the conclusions.

A. Bujari and C. E. Palazzi are with the Department of Mathematics, University of Padua, Italy, e-mail: {abujari,cpalazzi}@math.unipd.it.

M. Luglio, M. Quadrini, C. Roseti and F. Zampognaro are with the Department of Electronics Engineering, University of Roma Tor Vergata, Italy, e-mail: {luglio,roseti,zampognaro}@ing.uniroma2.it

II. PROXY ARCHITECTURE FOR HYBRID BACKHAUL OPTIMIZATION

The proxy architecture is largely used in satellite communications to introduce domain-specific enhancements able to compensate for the performance degradation caused by the high and variable delays. The PEP architecture can be centralized, a single PEP component implemented within a single node, or distributed, two or more PEP agents typically implemented in multiple nodes. This can be referred also to a chain of proxies or chain of PEPs. In case of a centralized architecture the performance enhancement is applied to the single node of interest whereas, in a distributed architecture, the enhancement involves the cooperation of all PEP agents found in the delivery path. In the latter approach shown in Fig. 1, PEP agents could perform different operations; for instance, a PEP agent located near the client(s) i.e. edge side, should be able to generate spoofed TCP ACK packets while a PEP agent located towards the public Internet i.e. core network side, suppresses ACKs received from the destination. In this context, the distributed PEP agents support the same custom transport protocol to enable an efficient backhaul over satellite, whereas a legacy transport protocol is used to communicate with the Web end-systems. This distributed architecture is largely adopted in satellite communications and constitutes the general reference architecture for our proposed solution.

To efficiently utilize multiple heterogeneous access links, traffic engineering operations gain particular importance [7]. A study of bandwidth aggregation techniques applied at different layers of the protocol stack e.g., data link, network, transport, application and tunneling aggregation etc., has been conducted in [8] concluding that there is no silver bullet and the best solution depends on the scenario and on the considered application(s). As an example, it is shown how aggregation at the network layer can be efficient for those applications requiring high throughput (i.e. large data file transfer) but, simultaneously, the solution may negatively impact real-time applications (e.g., receiving out-of-order packets). Concerning transport layer aggregation, multipath TCP (MPTCP [9]) represents a common solution.

Regarding the tunneling approach, it can be used to overcome problems incurred using network and transport layer aggregation when the presence of firewall or middle-boxes could compromise the communication.

On the use of multiple backhaul links via proxy components, the authors of [10] identify possible wireless and wired technologies for use in an hybrid backhaul network scenario aimed at addressing an ever increasing traffic demand while improving coverage and capacity. Within the set of the proposed technologies and approaches, the satellite represents a valid alternative when, for instance, the deployment cost of a wired/wireless solution overcomes the cost of the consumed bandwidth. The authors argue how different backhaul technologies, and in particular the satellite, may by used to complement each other allowing for different QoS requirements to be met at the same time.

The adoption of the Software Defined Networking (SDN) and NFV technology in satellite ground segment systems can facilitate the integration of the satellite component within 5G [11], [12]. In this respect, the novelty of our approach lies in the adoption of VNFs based on tailored technology exclusively operating at the network layer and above, properly configured to deal with heterogeneous Web traffic and different performance requirements. In this way, there is no direct dependency with lower layers (i.e. L2), which are generally under control of the network operations. This property guarantees the flexible deployment of the proposed solutions over a large set of network configurations, without interactions with SDN underlying setups. Specifically, dynamic and differentiated forwarding path(s) can be managed at the application layer through a vPEP orchestration component, as shown in Fig. 1.

III. vPEP DESIGN FOR WEB PERFORMANCE OPTIMIZATION

The legacy PEP represented a hardware-based agent working as a circuit-level proxy, running customized network software and protocols based on a static configuration and architecture. On the contrary, the vPEP is now defined as a chain of VNFs (IP layer and above) deployed to allow an optimized integration of a satellite component. The target vPEP architecture transparently enables the satellite between two IP-enabled nodes running a set of chained VNFs. The combination of the satellite link with the default terrestrial backhaul provides a hybrid network segment, which increases the overall bandwidth and network capabilities ready to be exploited for efficient traffic offloading and service continuity.

The vPEP functional architecture and integration in 5G has already been sketched and presented in [13] along with a preliminary functional validation of a reduced vPEP solution using synthetic traffic. This work steps forward in this direction by: (i) presenting a complete protocol stack and optimizations (QUIC vs. MPTCP) tailored to real Web traffic, (ii) describing a fully fledged virtualized testbed with NFV-compliant management functionalities used for the tests, (iii) assessing the vPEP by employing real traffic demonstrating a deployment of the suitable network technology as VNFs.

A. The vPEP Network Service

On top of the sketched technological approach, the most recent standardized transport and application protocols such as MPTCP, QUIC, HTML5 etc. can be deployed, once properly customized, as VNFs. Although these protocols are designed for a terrestrial network environment, either wired or wireless, some of their features can help to achieve good performance over satellite, such as:

- multiplexing and data push (QUIC GETs over a single connection);
- fast start up (large TCP window);
- bonding by exploiting multi-link capacity (e.g., MPTCP);
- connectionless protocols (HTTP over UDP);
- 0-latency (Transport Layer Security TLS1.3);
- Web page pre-rendering (i.e. HTML5);
- modern congestion control (e.g., Bottleneck Bandwidth and Round-trip propagation time (TCP BBR)).

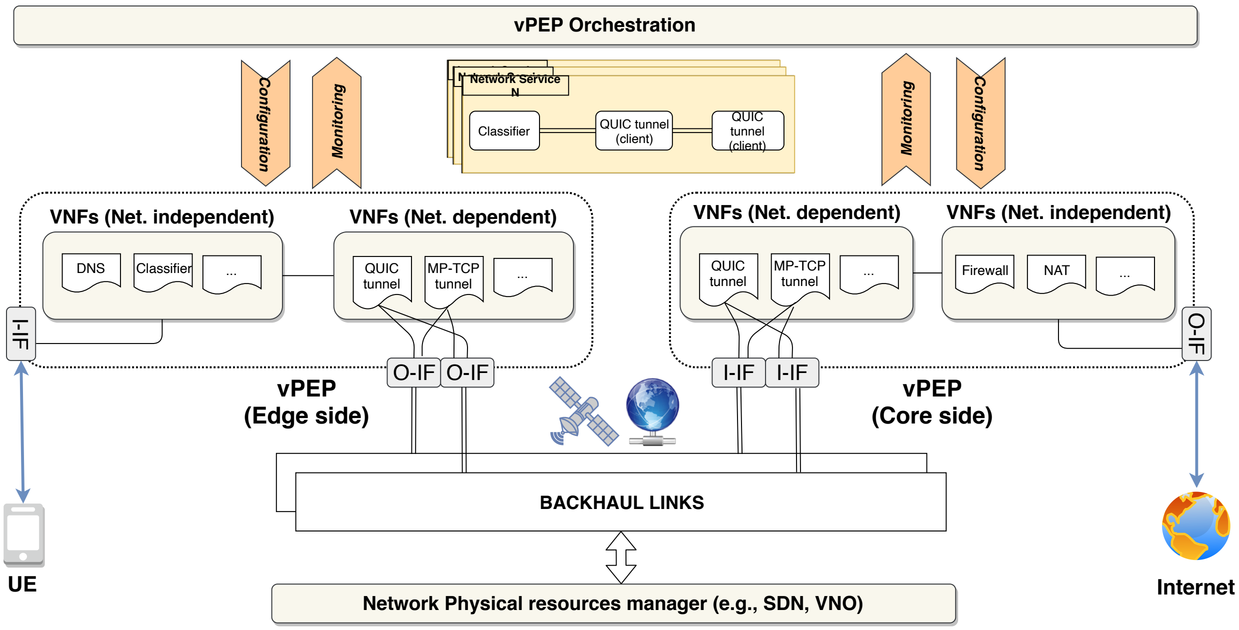

Fig. 1: The reference vPEP-based solution, enabling a multi-protocol architecture with a multifold processing capability.

In the past, the above (or similar) techniques have been pursued by researchers in satellite communications, exploited to upgrade legacy protocols and to propose satellite-specific solutions for use over PEP-to-PEP communication links. Now, these protocols and tailored optimizations can be properly combined and deployed as VNFs in the proposed vPEP solution, without the need to look for hardware-based, ad-hoc protocols for the satellite.

A VNF chain is built up to accomplish the requirements of different application flows. In the proposed vPEP deployment, short-lived HTTP traffic is tackled using a QUIC-based VNF, creating an application layer tunnel used to aggregate traffic over the backhaul, while longer data transfers exploit the bonding capability of the MPTCP protocol.

A specific VNF is in charge of traffic classification and differentiation based on some criteria e.g., rule-based approach depending on the port number and transport type, and to steer it towards the most appropriate VNF for performance optimization(s). Specifically, two alternative VNF-pairs can be selected to transport the incoming traffic over the hybrid backhaul into either QUIC-based tunnels or MPTCP ones, respectively.

Additional network independent functions might be placed by the service operator, on the basis of the applications needs, close to the user (such as local DNS resolvers) or at the server side e.g., firewall, traffic shaper, NAT, etc. Functionalities related to these entities are delivered as VNFs and their relationship is depicted in Fig. 1. The ETSI NFV specification natively supports the concept of service chains and several compliant frameworks are available. In our proof-of-concept implementation, we rely on a customized version of the Openbaton NFVO [14].

B. Aggregation, failover and offload through vPEP tunneling

The QUIC protocol is equipped with an in-band connection migration feature, providing an efficient mechanism that can seamlessly move traffic from one network to another without causing any transfer drops or significant performance degradation. On this basis, we propose the use of two QUIC-based VNFs, acting as client and server respectively, distributed among the virtualized infrastructure nodes. The QUIC connection setup leads to the introduction of the novel concept of QUIC tunnel, meant as the aggregation of web-traffic, steered on a single QUIC connection based on type (i.e. object format), source or destination network, server list, etc. Therefore, a QUIC tunnel represents the granularity accounted for traffic steering operations, i.e. *Tunnel Migration*. The number of concurrent QUIC tunnels is a flexible configuration parameter depending on the QoS requirements settled by the service provider. In the target scenario, a QUIC tunnel can be seamlessly and transparently moved across the available backhaul links without any disruption of the end-to-end service. The traffic over a QUIC tunnel could in turn be fed to prediction and traffic balancing algorithms.

Accordingly, a specialized VNF, namely *classifier*, is envisaged at the entry point of the vPEP to create and manage mappings between end-system flows and the available QUIC tunnels. For this reason, the number of QUIC tunnels must be a trade-off between a low value (to guarantee suitable values for the steady-state throughput measurements) and a high one to improve on traffic granularity and allow a better grouping of flows based on the performance requirements. The adoption of QUIC-based tunnels over the satellite reduces the amount of signaling, does not impair security and allows to actually exploit the 0-latency feature (fixed counterparts) [13]. Indeed, the resulting QUIC tunnel migration is recommended as an efficient tool to support traffic offloading and failover, in case of multiple backhaul links (supplementary or complementary).

Similar principles apply to the MPTCP tunnel approach, implemented through the client/server VNF counterparts. MPTCP tunnelling over a hybrid backhaul could be employed to augment backhaul link capacity, i.e. bonding or to add resilience in case of link failure(s). This particular optimizations are effective in case of long transfers, because in short transfers only the link labelled as "default" is most likely used. Thus, the forwarding to the MPTCP VNF pair is recommended in case of software downloads and large software upgrades, system synchronizations and bulk transfers.

IV. TESTBED ARCHITECTURE

To validate our proposal, we have built an emulation testbed comprised of both a terrestrial and a satellite link. The physical architecture of the testbed is represented in Fig. 2 and is comprised of three physical high-end server machines, where each machine represents a distinct Point of Presence (PoP) hosting testbed-related functionalities. PoP1 hosts the core building blocks of the vPEP network service and is dedicated to the deployment of functions both at the local (station) and remote (satellite gateway) side. PoP2 represents the junction between local and remote sides, hosting both the satellite and terrestrial emulation segments representing the backhaul network. Last, PoP3 hosts all services useful to validate vPEP functionalities, including the implementation of web client and servers, the end points of the communication. In addition, PoP3 can be used as an exit point able to offer Internet access to all the elements of the testbed performing gateway operations (e.g., forwarding, NAT, etc.).

A KVM/QEMU hypervisor is installed on PoP1, PoP2 and PoP3, enabling a first virtualization layer by defining a set of Virtual Machines (VMs). Within these VMs, VNFs implemented by Docker containers can be dynamically deployed as a second virtualization layer. The mix of both hypervisor and container-based technology, allows us to combine the benefits of exploiting customized kernel level functionalities (i.e. MPTCP) and a lightweight approach to virtualization. In fact the Docker container will inherit the kernel configuration of the host VM, and in particular the availability of the MPTCP protocol. All the VMs are based on Linux O.S. with default configurations for network and protocols, with the exception of the TCP receive buffer for the MPTCP VMs. The TCP receive buffer is set manually to 5 MB in order to boost the data transfer, in particular, on the satellite segments. The details of the testbed configuration, the source code for the orchestrator and the repository with the main Docker images are available online at [14].

A. Data Path

To support the communication among all testbed components, some of which reside on the same hardware, a multi-layered network configuration has been setup, resulting in the data path (dotted lines) shown in Fig. 2. The different isolated networks for the data path(s) managed in the testbed are enforced by means of VLAN tagging which allows to isolate multiple logical IP pipes among the servers, realizing separated links. In particular, a VLAN segmentation is configured on the following links:

1) between PoP3 and PoP1 in order to manage different network slices at the remote side;
2) between PoP1 and PoP2 in order to discriminate between different satellite and terrestrial backhaul links.

VLAN segmentation between PoP3 and PoP1 allows to create the *service slices* in accessing the vPEP sub-system, at both local-side (optional) and remote-side. Through this configuration, end-user applications in PoP3 are granted isolated connectivity with PoP1 (vPEP sub-system). Furthermore, within PoP1 it is necessary to arrange the backhaul connectivity through PoP2. This is as well enabled through VLAN tags, creating independent backhaul data paths. Please note that the backhaul connections are static pre-configured links, through either terrestrial or satellite emulated nodes, which can be used independently.

In addition to the L2 isolated connectivity granted by VLAN-based networking, a Docker-based overlay network is setup and mapped to the underlying data path. This results in the virtual path between end-points as shown in Fig. 2, which is representative of a realistic virtualized network configuration. The proposed service-based perspective allows to introduce a networking approach which is independent from the underlying physical infrastructure configuration. For more information on the configuration details and step-wise procedure to reproduce the testbed, we refer the reader to [14], where also a 4G cell access network is considered.

B. Backhaul Network

Terrestrial and satellite emulators are deployed and configured on the PoP2 which is connected to PoP1 with a dedicated 10 Gbps link. The testbed allows two basic backhaul types: terrestrial (default) and satellite (supplementary). For each backhaul link, different virtual links represent different physical nodes/infrastructures (i.e. satellite terminals or terrestrial network trunk) enabled in the test scenario. In fact, different connectivity services can either share the same infrastructure or can use independent backhaul links. These testbed allows to setup several independent end-to-end data paths (slices), suitable to host several service operators.

The terrestrial emulator is Linux-based and uses Linux traffic control embedded functionalities to instrument the terrestrial backhaul. In particular, it is possible to adjust for each terrestrial link the latency, the jitter, the loss, the bottleneck capacity and a buffer size. In these routers, the rules are configured at startup and different network profiles are available in the testbed, from 100 Kbps to 1 Gbps, with low RTT values.

The satellite emulator is represented by an OpenSAND installation configured with a pre-defined number of satellite stations and one gateway, in a DVB-RCS2 star setup. Different link profiles are available and each station has a return and a forward link, with a baseline physical RTT of 530 ms.

V. EXPERIMENTAL ASSESSMENT

We propose a set of scenarios to validate the approach in a realistic deployment. In each scenario, we compare differ-

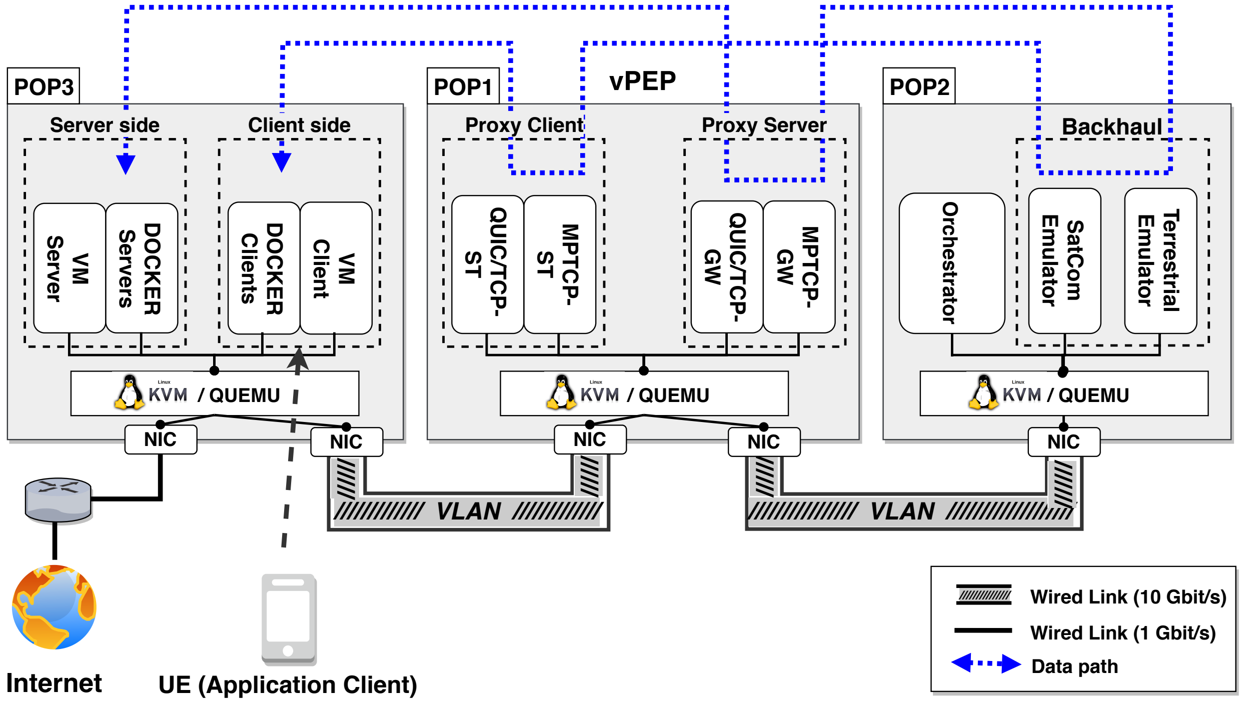

Fig. 2: Testbed components and functionalities.

ent configurations: (i) one disabling the vPEP (benchmark), adopting simple routing-only VNFs along the data path, and (ii) one where the vPEP is enabled via QUIC or MPTCP. All the proposed vPEP optimizations can be deployed on-demand and associated to parallel data paths, i.e. different service operators and applications. The reference for all the tests is a terrestrial link setup compliant with an ADSL profile (i.e. 10 Mbps downlink and 1 Mbps uplink, with an RTT of 30 ms) while the satellite link is configured as VSAT profile (i.e. 40 Mbps downlink in DVB-S2 and 2 Mbps uplink in DVB-RCS2).

A total of three scenarios are considered and discussed, all enabled by distributed chains of VNFs on a common infrastructure. The first scenario showcases the QUIC multi tunnel management feature used to aggregate and dynamically (re)distribute traffic over the available backhaul links. The second one relates to the bonding of multiple channels through the use of the MPTCP protocol capable of exploiting the aggregate capacity of the backhaul links. Last, the third scenario highlights the benefit of MPTCP also in case of possible link outages. The same benefits apply to the QUIC-based tunnel as well, herein omitted for the sake of brevity.

A. *QUIC - Multi tunnel management*

In this scenario, two docker clients running a web traffic generator are deployed on *PoP3-Docker-client*. The generated web traffic consists of real HTTP GETs issued to a web server deployed at *PoP3-Docker-server*. The traffic generators are set to produce an average download traffic of 6 Mbps and 12 Mbps, respectively. The resulting traffic pattern oscillates over time, typical of the HTTP request-response paradigm. The default terrestrial link is coupled with a supplementary satellite link that can be opportunistically used at any time thanks to the vPEP. The specific network configuration imposes a large buffer at the shared terrestrial-satellite network path, adopted in order to opportunistically exploit the additional satellite capacity. As an implication, the RTT could be subject to significant increases upon persistent congestion even when using the terrestrial link only. Such an RTT dynamic is used to design the proposed migration scheme, whereas other buffer configurations might require different approaches based on other indicators, i.e. losses, re-transmissions, data rate, etc., to be addressed in future work.

In the benchmark case, both traffic generators will establish end-to-end TCP connections through the terrestrial link only. When enabling the vPEP, the classifier function is configured to associate all web traffic coming from one generator to a high-priority tunnel called *Gold* and all web traffic coming from the other generator to a different low-priority tunnel called *Silver*.

The outcome of this test case is presented in Fig. 3. In the benchmark case, the terrestrial link is saturated by the two traffic generators while the RTT results well above its minimum value, skyrocketing from about 30 ms to nearly 500 ms (black line in the rightmost plot). This implies that all HTTP transfers, as well as other services sharing that link, experience a high latency, resulting in QoS degradation of current and forthcoming connections.

When vPEP is enabled, it is possible to (re)distribute part of the traffic to solve congestion issues. In this scenario, at time t = 33 s the terrestrial link reaches its maximum rate for more than 5 s while the RTT steadily increases, with values up to more than 5 times the reference RTT value. This RTT dynamic is deemed a persistent congestion indicator by the

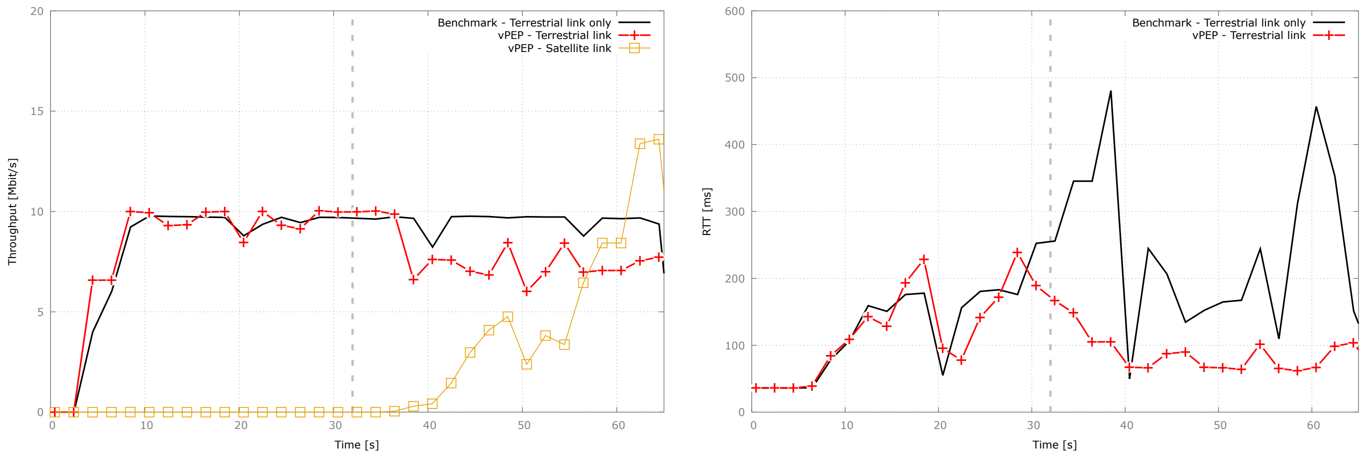

(a) Bandwidth

(b) RTT

Fig. 3: Link offloading through the QUIC multi tunnel migration.

multi-tunnel management algorithm and the traffic over the Silver tunnel is offloaded onto the satellite link. This operation does not require any start up signalling over the new link, and allows to reduce the terrestrial RTT well below 100 ms. The use of QUIC for the tunneling protocol also allows to maintain the consistency of already established connections, which are not interrupted upon migration. It is noteworthy to point out that the trigger enacting the tunnel migration can be either performed by a human operator controlling network statistics in real-time, or autonomously by the use of dedicated algorithms.

In conclusion, the vPEP effectively and transparently allows to move back and forth one or more tunnels over the available alternative paths. On a practical note, the service provider could configure the number of tunnels, the traffic classification policies and the migration rules enacted by tailored VNFs.

B. *MPTCP - Channel bonding*

Similar to the prior scenario, a web client is deployed on *PoP3-Docker-client* and a web server is deployed in conjunction on *PoP3-Docker-server*. The test starts when the client issues a download request for large file (i.e. 100 Mbytes) to the server and terminates when the download completes.

In the benchmark case, a single end-to-end TCP connection is used to download the file on either one of the available links. In case of vPEP optimization, an MPTCP tunnel is established between *PoP1-MPTCP-ST* and *PoP1-MPTCP-GW* and both terrestrial and satellite links are available.

The outcome of this test is presented in Fig. 4. In the benchmarking cases, the download is performed considering a single link. In case of use of the terrestrial link, the transfer completes at a 10 Mbps nominal rate and takes about 80 s. The download rate for the same objects, when using the satellite link instead, is 16 Mbps on average and takes a little less than 50 s. On the other hand, when vPEP acceleration is enabled, the MPTCP tunnel exploits both terrestrial and satellite links simultaneously. The bonding action of MPTCP is successful since two heterogeneous backhaul links are jointly used to download a content from the web server, and the download rate perceived by the web client is on average 30 Mbps. Note that this measured average download rate is higher than the maximum of 10 Mbps for the terrestrial link plus the average of 16 Mbps achieved on the benchmarking case over satellite. In fact, due to the specific tuning of the buffers, the satellite MPTCP sub-connection can grow faster with regard to the end-to-end connection established with default system parameters between VMs in POP3. This result shows that MPTCP bonding is effective and allows to increase the download rate of a reference large object, by exploiting the combined capacity of available links.

C. *MPTCP - Link outage*

This scenario adopts the same settings as the prior one, but during the data transfer the terrestrial link is disconnected abruptly. In this case we have no benchmark case since when a link is not available anymore, already established connections will suffer from timeouts and must be re-established by the

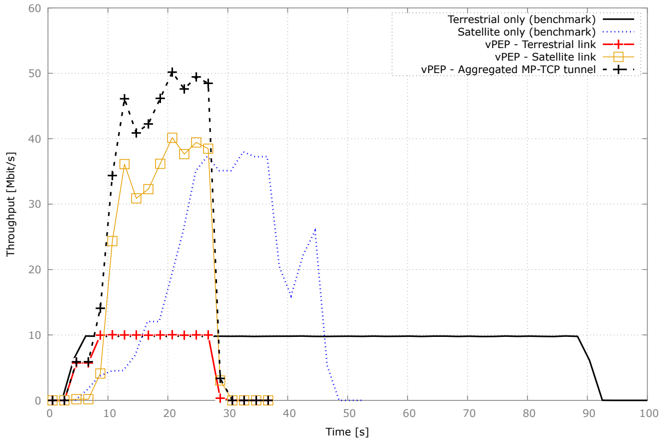

Fig. 4: Aggregated traffic benefits using MPTCP bonding via vPEP acceleration.

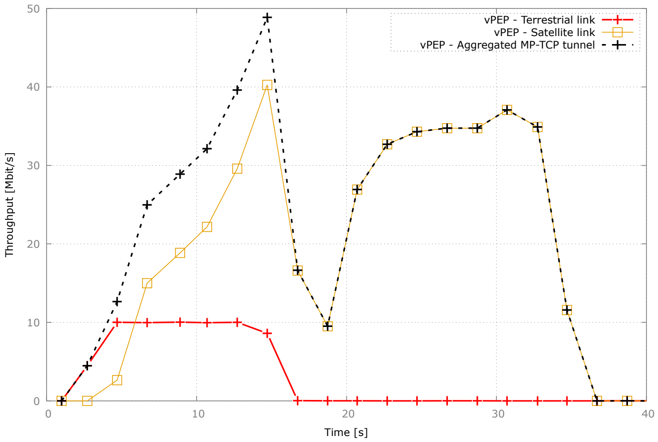

Fig. 5: Traffic continuity through MPTCP tunnel in case of link outage.

application, potentially restarting the download. When the vPEP is enabled instead, the MPTCP multi link management capability is operational, allowing the transfer to continue even when the terrestrial link is disconnected, although the download relies on the satellite link only.

The results related to this scenario are presented in Fig. 5. In the beginning, the MPTCP tunnel exploits both terrestrial and satellite links, as in the previous scenario. When the terrestrial link is not available anymore, at t = 15 s, the communication continues seamlessly on the alternate satellite path. Therefore, the bonding capability provides also resilience when facing a link outage, transparently continuing the download over the available link(s).

VI. CONCLUSION

The new communication paradigms based on NFV have a disruptive impact on the design of current and future networks, making worthwhile a review of the traditional networking and performance assumptions. Satellite communications have been generally employed as a gap-filler in scenarios not fully supported by terrestrial networks, because physical impairments required the introduction of specialized middleboxes. Unfortunately, this also created a compatibility problem with some legacy TCP/IP protocols. Now, the affirmation of the virtualization technology invalidates such a view and allows to review the satellite role. To this aim, we have here presented possible optimizations enabling a significant performance improvement on real Web traffic. This approach is based on the dynamic deployment of tailored VNFs satisfying different traffic types requirements. As a result, the satellite is involved as an enabler of advanced capabilities and not anymore as a mere physical capacity carrier.

ACKNOWLEDGMENT

This work is an outcome of ESA VIBeS project, Artes AT contract n. 4000122991/18/UK/ND [15]. Responsibility for the presented content resides with the authors.

REFERENCES

[1] C. Niephaus, M. Kretschmer, and G. Ghinea, “QoS Provisioning in Converged Satellite and Terrestrial Networks: A Survey of the State-of-the-Art,” *IEEE Communications Surveys Tutorials*, vol. 18, no. 4, pp. 2415–2441, 2016.

[2] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, “Toward 6G Networks: Use Cases and Technologies,” *IEEE Communications Magazine*, vol. 58, no. 3, pp. 55–61, 2020.

[3] P. Wang *et al.*, “Convergence of Satellite and Terrestrial Networks: A Comprehensive Survey,” *IEEE Access*, vol. 8, pp. 5550–5588, 2020.

[4] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby, “Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations,” Internet Requests for Comments, RFC Editor, RFC 3135, Accessed on: July 2020.

[5] N. Assaf, J. Luo, M. Dillinger, and L. Menendez, “Interworking between IP Security and Performance Enhancing Proxies for Mobile Networks,” *IEEE Communications Magazine*, vol. 40, no. 5, pp. 138–144, 2002.

[6] ETSI GS NFV-MAN 001, “Network Functions Virtualisation (NFV); Management and Orchestration,” Accessed on: July 2020. [Online]. Available: http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_nfv-man001v010101p.pdf

[7] ETSI TR 103 351 V1.1.1, “Satellite Earth Stations and Systems (SES); Multi-link routing scheme in hybrid access network with heterogeneous links,” Technical Report, Accessed on: July 2020.

[8] A. Abdelsalam, M. Luglio, C. Roseti, and F. Zampognaro, “Analysis of Bandwidth Aggregation Techniques for Combined use of Satellite and xDSL Broadband Links,” *International Journal of Satellite Communications and Networking*, vol. 37, no. 2, pp. 76–90, 2019.

[9] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, “TCP Extensions for Multipath Operation with Multiple Addresses,” Internet Requests for Comments, RFC Editor, RFC 6824, January 2013.

[10] O. L. A. Lopez, H. Alves, R. D. Souza, and M. Latva-aho, “Hybrid Wired-Wireless Backhaul Solutions for Heterogeneous Ultra-Dense Networks,” in *IEEE Vehicular Technology Conference (VTC Spring)*, June 2018, pp. 1–5.

[11] E. Ferrús *et al.*, “SDN/NFV-enabled Satellite Communications: Ground Segment Network Architecture for 5G Integration,” in *European Conference on Networks and Communications (EuCNC)*, 2016.

[12] T. Li, H. Zhou, H. Luo, Q. Xu, and Y. Ye, “Using SDN and NFV to Implement Satellite Communication Networks,” in *International Conference on Networking and Network Applications (NaNA)*, 2016, pp. 131–134.

[13] A. Abdelsalam *et al.*, “Implementation of Virtualised Network Functions (VNFs) for Broadband Satellite Networks,” in *European Conference on Networks and Communications (EuCNC)*, 2019, pp. 182–186.

[14] NITEL, “vPEP setup and configuration,” Accessed on: July 2020. [Online]. Available: https://bitbucket.org/fzamps/vpep

[15] VIBeS Project, “Implementation of virtualised network functions (VNFs) for broadband satellite networks,” Accessed on: May 2020. [Online]. Available: https://artes.esa.int/projects/vibes

Armir Bujari is Assistant Professor in Computer Science at the University of Padua. His research interests are focused on the design and analysis of communication protocols for wired/wireless networks and network architectures.

Michele Luglio is Associate Professor of Telecommunication at University of Rome “Tor Vergata”. He works on designing satellite systems for multimedia services both mobile and fixed.

Claudio E. Palazzi is Associate Professor in Computer Science at the University of Padua. His research interests are primarily focused on the design and analysis of communication protocols for wired/wireless networks, Internet architectures, and mobile users.

Mattia Quadrini is a PhD student at the University of Rome “Tor Vergata”. His research interests include satellite, transport and networking systems.

Cesare Roseti is Assistant Professor at the University of Rome “Tor Vergata”. His research interests include satellite, transport and networking systems.

Francesco Zampognaro is Assistant Professor at the University of Rome “Tor Vergata”. His research interests include satellite, transport and multimedia networking systems.

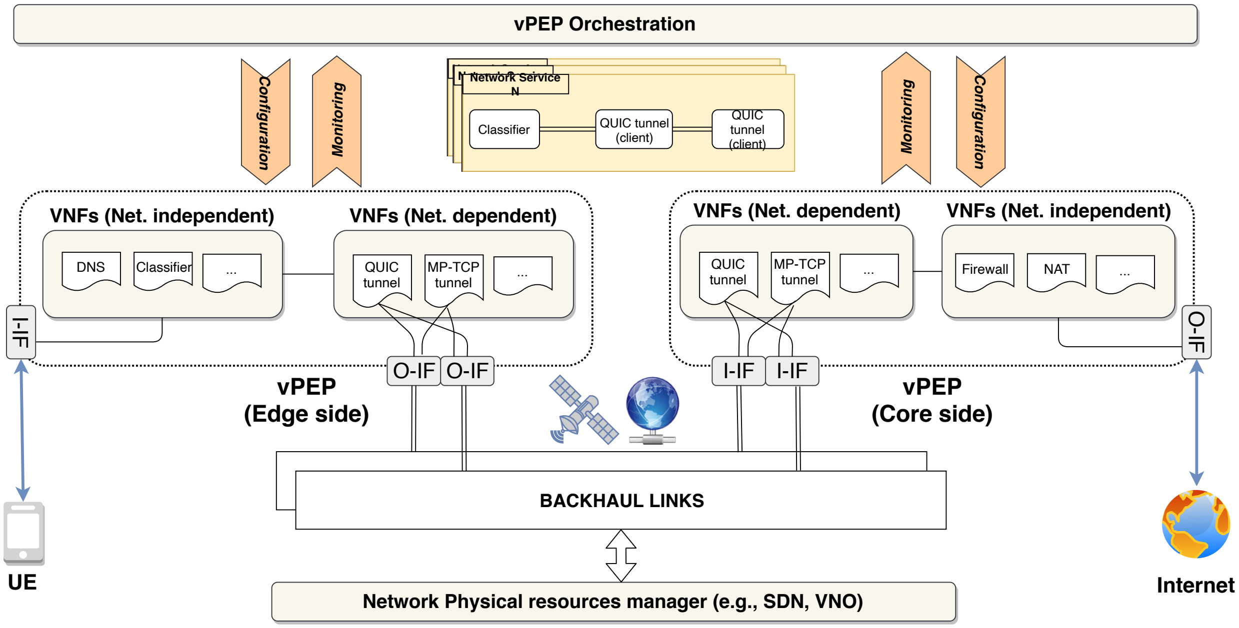

Fig. 6: The reference vPEP-based solution, enabling a multi-protocol architecture with a multifold processing capability.

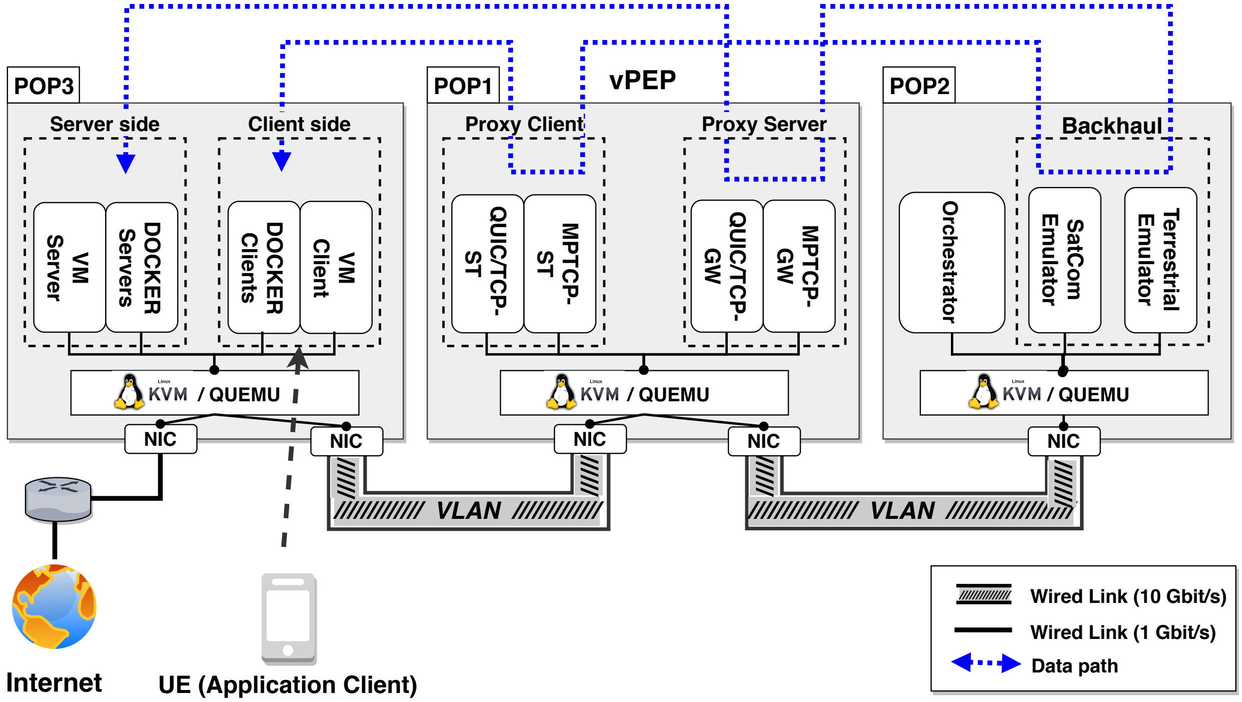

Fig. 7: Testbed components and functionalities.

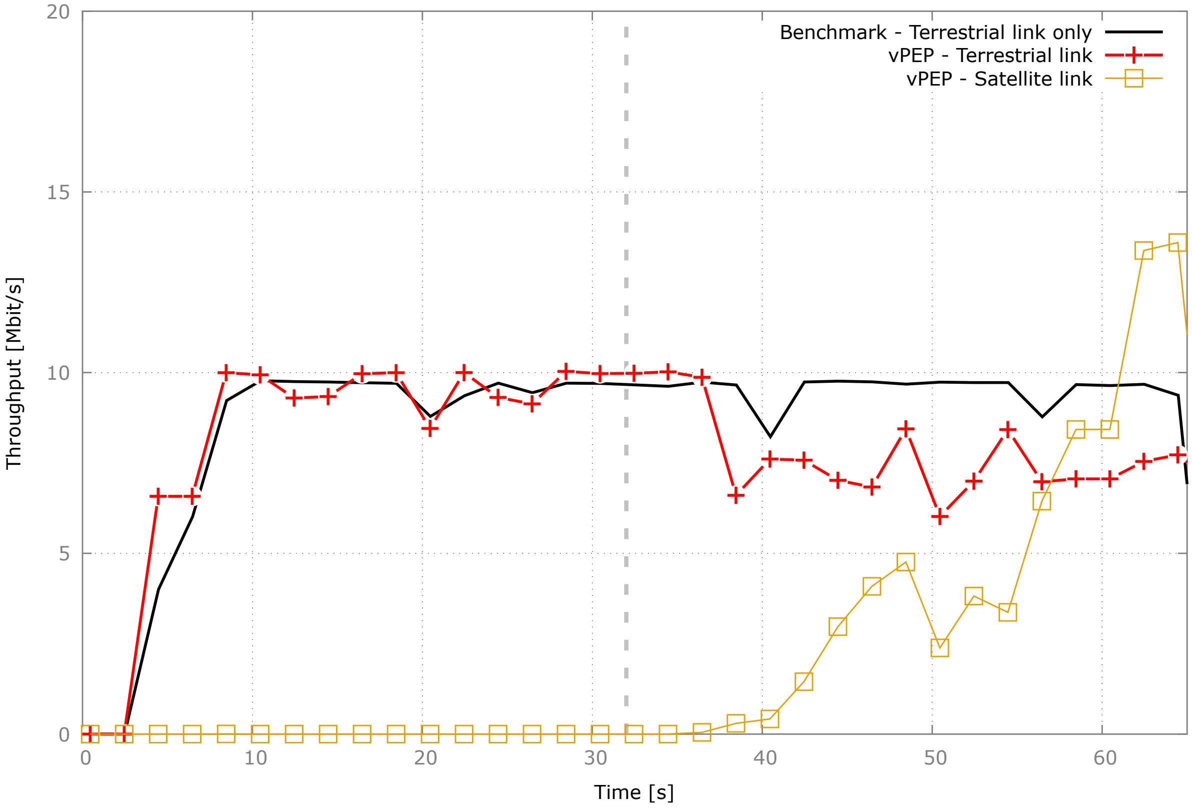

(a) Bandwidth

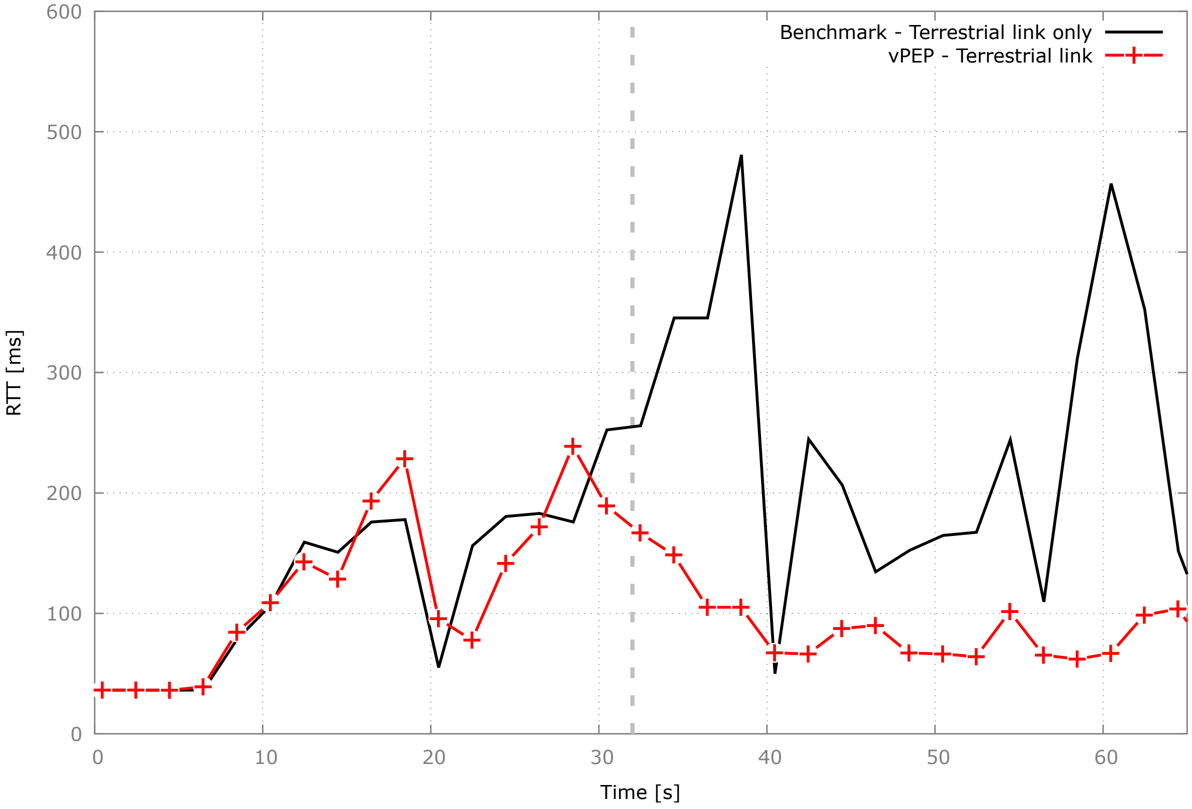

(b) RTT

Fig. 8: Link offloading through the QUIC multi tunnel migration.

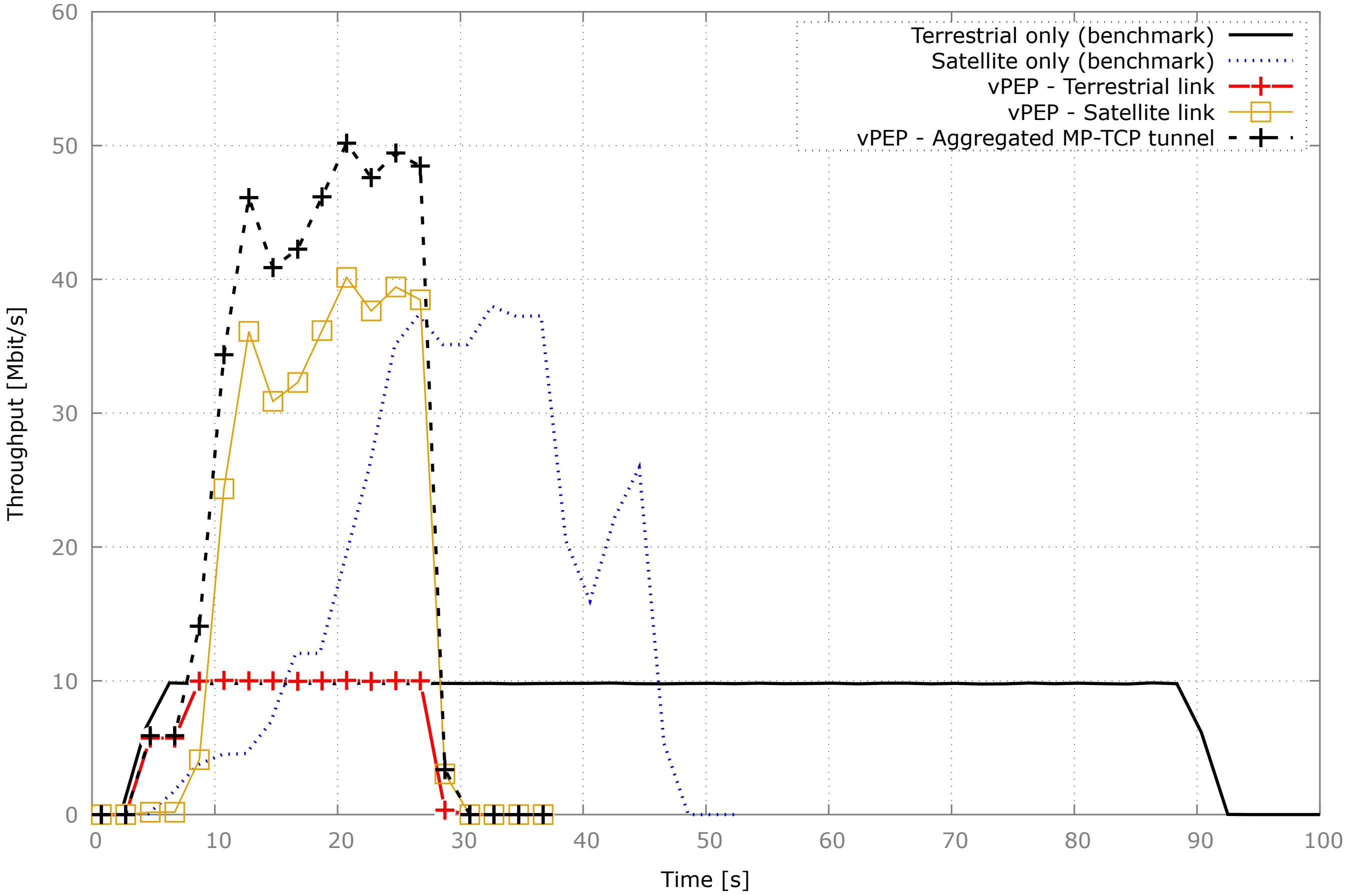

Fig. 9: Aggregated traffic benefits using MPTCP bonding via vPEP acceleration.

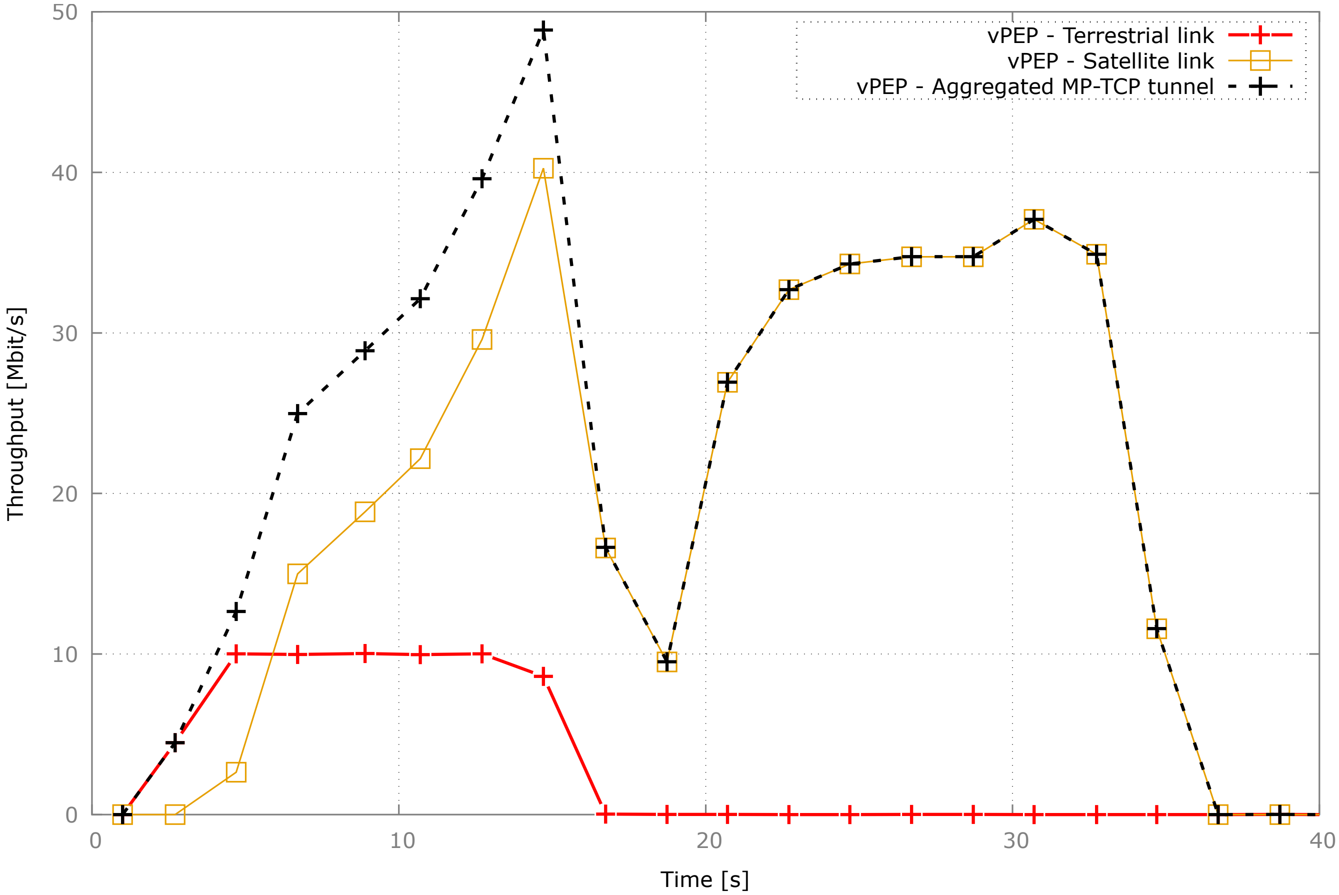

Fig. 10: Traffic continuity through MPTCP tunnel in case of link outage.