Intelligent Gateways Enabling Broadband Access via Integrated Terrestrial and Satellite Systems

Luc Ottavj[1(✉)], Emmanuel Duros[1], Jacques Webert[1], Fabrice Gamberini[1],
Christian Niephaus[2], Javier Perez-Trufero[3], and Simon Watts[3]

[1] OneAccess Networks, 2455 Route Des Dolines – BP355,
06906 Sophia Antipolis, France
`{luc.ottavj,emmanuel.duros,jacques.webert,`
`fabrice.gamberini}@oneaccess-net.com`
[2] Schloss Birlinghoven, Fraunhofer-Fokus, 53757 Sant Augustin, Germany
`Christian.niephaus@zv.fraunhofer.de`
[3] Avanti Communications Ltd, 20 Black Friars Lane, London, EC4V 6EB, UK
`{javier.pereztrufero,simon.watts}@avantiplc.com`

Abstract. Satellite broadband systems will play a key role in reducing the Digital Divide by complementing terrestrial networks in the delivery of next generation broadband to users in remote and rural locations. We describe an integrated broadband delivery system to fixed users that makes simultaneous use of heterogeneous access networks in order to optimize the end-user QoE. The design of the overall network architecture and the key building blocks of the routing entities at both ends of the integrated system are presented. The paper argues that MPTCP is an appropriate mechanism to offer hybrid communications over heterogeneous links and details a specific implementation of the intelligent routing gateways. Results from in-factory validation tests of the prototyped platforms are presented and discussed.

Keywords: Satellite broadband · Integrated terrestrial and satellite systems · MPTCP · Intelligent routing · Digital divide

1 Introduction

The research project BATS (Broadband Access via integrated Terrestrial & Satellite systems), co-funded by the European Commission (EC) under the FP7 programme, addresses the delivery of Broadband (BB) future services in Europe according to the EC Digital Agenda [1] objective to reliably deliver > 30 Mbps to 100 % of European households by 2020. Next generation geostationary (GEO) broadband satellite systems will play a key role in achieving such objectives as the accelerated deployment of terrestrial broadband technology will not be able to satisfy this requirement in the most difficult-to-serve locations. Market studies indicate that in a significant number of regions of Europe, more than 50 % of premises, will lack access to superfast Broadband [2, 3], either due to a lack of coverage in areas where the revenue potential for terrestrial service providers is too low

© Institute for Computer Sciences, Social Informatics and Telecommunications Engineering 2015
P. Pillai et al. (Eds.): WiSATS 2015, LNICST 154, pp. 92–102, 2015.
DOI: 10.1007/978-3-319-25479-1_7

(unserved areas) or due to technological limitations which diminish the available end-user data rate in some suburban and many rural environments (underserved areas).

The BATS project [4] aims to bridge the potentially widening Broadband divide between urban and rural areas and fulfil the Digital Agenda targets in the underserved areas via an integrated network that combines the flexibility, large coverage and high capacity of future multi-spot beam satellites, the low latency of fixed DSL lines, and the pervasiveness of mobile-wireless access. The integrated broadband service will be delivered to the end-user via an *Intelligent User Gateway (IUG)* and *Intelligent Network Gateway (ING)*, dynamically routing traffic flows according to their service needs through the most appropriate broadband access network, with the goal of optimizing the user's Quality of Experience (QoE). Due to the heterogeneity of the various technologies (e.g. satellite broadband offers high bandwidth but higher latency as opposed to narrowband xDSL low latency links) randomly distributing traffic among the different connections despite their different characteristics could in turn affect negatively the service quality. More sophisticated methods are required which allow for a simultaneous use of the different links in a seamless manner fully exploiting their particular benefits.

In this paper we present the IUG and ING design approach followed in the BATS project and argue that MPTCP [8] is an optimal mechanism to enable Quality-of Service (QoS)-aware Link Selection, key requirement for the integration of satellite and terrestrial networks.

The remainder of this article is organized as follows. Section 2 describes the overall integrated broadband architecture and focuses on the functional design of the key enabling modules in the IUG and ING. Section 3 provides results from the in-factory tests of the prototype IUG and ING to be used in the laboratory and field trials of the project. Finally, Sect. 4 concludes the article and provides some insight into our future activities.

2 Network Architecture and Key Enabling Components

2.1 The Overall Integrated Network Architecture

As illustrated in Fig. 1, the overall network architecture comprises the three broadband access segments, namely xDSL, cellular and satellite, whose connections are terminated at the IUG on the end-user side and at the ING on the operator side. The IUG is the routing entity located at the end-user premises serving as the focal point for the integration of the terrestrial and satellite connections. As the counterpart of the IUG on the network side, the ING has the functionalities of both managing a set of associated IUGs and acting as a single connection interface to the public internet. For the downstream traffic, an ING has equivalent building blocks and routing functionalities to the IUG. The main functionality of IUG and ING is to route the outgoing traffic towards the most suitable access network segment considering the QoE requirements of each particular traffic flow and the real-time status of each of the links. Based on this, the ING shall be located closer to the Point of Presence (PoP) of the terrestrial operators involved in the integrated system, in order not to increase the latency of the services routed terrestrially (which are meant to be the most delay sensitive).

Due to TCP performance degradation over satellite links, Performance Enhancement Proxies (PEPs) [6] are currently the most commonly adopted solution to achieve good transport performance (in terms of link utilization and user experience) whatever the available TCP stack at both ends (clients and servers). The location of PEP terminations in the BATS architecture has important impacts on the overall network design. Trade-off analysis carried out in the project concluded that the best compromise between performance, impact and complexity is to locate a high capacity PEP in a central point of the network near to the PoP or the ING, alleviating the internal re-routing and synchronization issues compared to a case where the PEP is located at the nominal satellite gateway and traffic needs to be re-routed to a redundant gateway to avoid service interruption due to fading or failure at the nominal one.

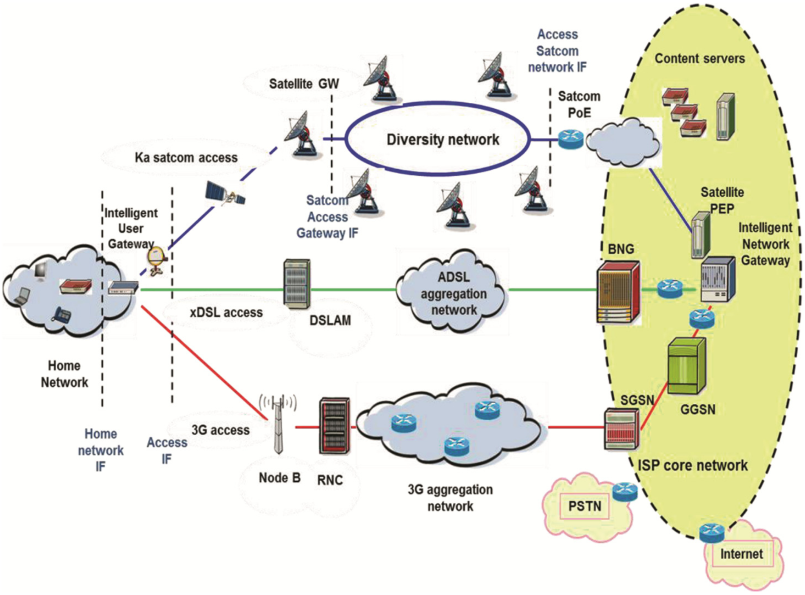

Fig. 1. Overall network architecture

2.2 The Intelligent User and Network Gateways

Due to the heterogeneity of the various access segment technologies, distributing traffic arbitrarily among them despite their different characteristics, e.g. satellite systems offer large bandwidth but introduce additional latency compared to typical terrestrial networks, has proved to be sub-optimal. A multipath routing entity such as the IUG and ING, which can seamlessly select the right access technology for the different types of network traffic, requires the ability to know the QoS requirements (e.g. bandwidth, latency and packet loss) of each particular traffic flow and to compare them with the real-time status of each of the available access links. Figure 2 illustrates the key building modules of the gateways. The heart of the system is the *Link Selection* algorithm itself, which operates on certain attributes or parameters and selects the most optimal link for a particular traffic unit. In order to do that, a module called *Link Status Estimation* evaluates the RTT, maximum bandwidth and packet losses of each of the paths between IUG and ING. In addition, the *Traffic Classification* module classifies the traffic as required to be able to take the link selection decisions and identifies the Class of Service (CoS) for QoS purposes. The IUG and ING have been conceived with a modular approach and well-defined connecting interfaces between the different components which allow for specific improvements on individual modules without modifying the entire system. In the following sub-sections details are provided on the approach BATS has followed for these different modules.

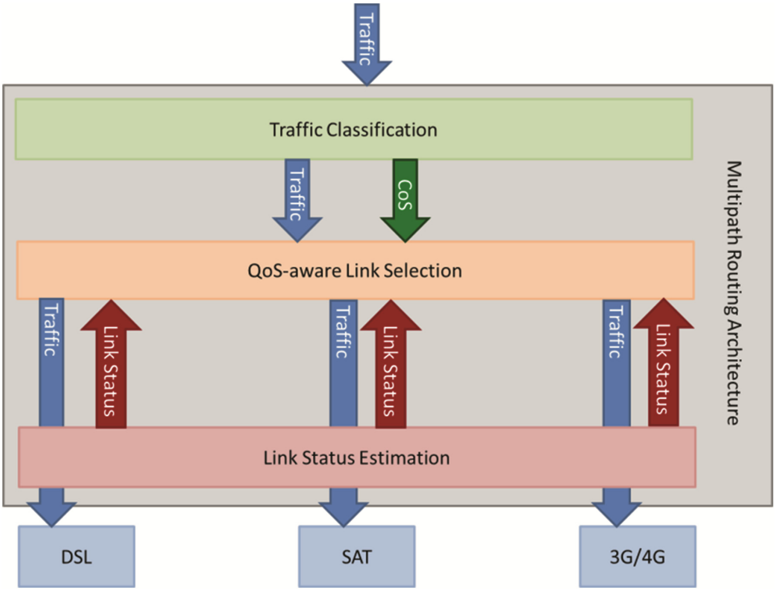

Fig. 2. Key building blocks in IUG and ING

Multipath Routing Architecture

Since TCP (and HTTP) are now the main protocols dominating the Internet traffic [7], they become a de-facto convergence layer. This trend is also driven by the Internet and content service provider infrastructures that are largely based on this protocol set. A perfect example is the evolution of the streaming services that are evolving from RTP/UDP to adaptive streaming (such as MPEG-DASH) over HTTP/TCP networks. For this reason, in BATS, MPTCP [8] has been identified as an appropriate mechanism for the core of the multipath routing architecture, as it aims at using multiple paths between a source and a destination host while providing the same interface to both the application and the network layer as in conventional TCP. Several TCP sessions, called MPTCP sub-flows, can be established on the different access links and used simultaneously whilst preserving the ordering of the packets at the far end. As with regular TCP, MPTCP is also designed to be an end-to-end transport layer protocol. Hence, in order to exploit the multipath features of MPTCP only between intermediate routers, MPTCP proxies need to be used at the IUG and ING, as shown in Fig. 3. In the BATS design, the implementation of MPTCP proxies at both IUG and ING allows the definition of a completely integrated architecture based on a common set of mechanisms independent from the heterogeneity of the access networks. Furthermore, the need for TCP acceleration in the satellite access segment perfectly fits with MPTCP proxies, as TCP PEPs can be co-located at the IUGs and INGs on both ends of the network.

In the BATS implementation, the negotiation between IUG and ING is done using TCP options to avoid losing time during payload handshakes. These TCP options need to be protected from layer 4 filtering equipment that may be in-between the IUG and ING. As an example, many firewalls and core network appliances modify the TCP options they do not understand, which would lead to a breakage of the negotiation between the BATS gateways. Therefore, the communication between IUG and ING happens within tunnels to protect encapsulated protocols including the header information. In this case, a UDP encapsulation is added in order to ease processing through NAT equipment.

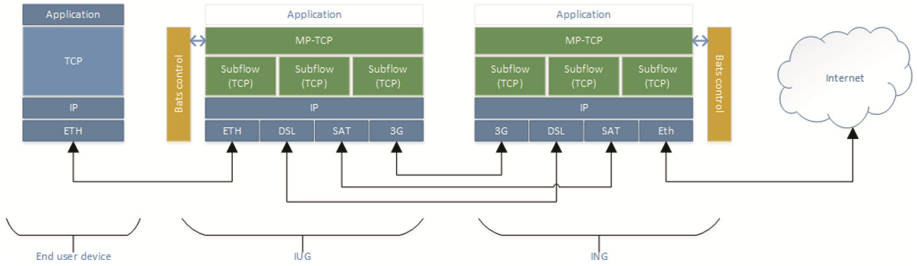

Fig. 3. Multipath routing architecture based on MPTCP proxies

Link Status Estimation Module

The link estimation module analyses each path between the IUG and the ING and provides the real-time status of different path characteristics, e.g. latency, bandwidth, loss rate, in order to assist the decision on link selection. Benefiting from the MPTCP architecture in which the IUG and ING become the end-points for the MPTCP subflows, in BATS the link estimation module is implemented by measuring information from the TCP connections, updating the statistics each time a TCP segment is sent or received.

The bandwidth estimation evaluates the traffic bandwidth towards a peer through each tunnel, estimating the maximum TCP bandwidth that a TCP connection could achieve. In order to evaluate the maximum TCP bandwidth internal variables of the TCP control block are evaluated. These include the amount of unacknowledged data in the TCP sending buffer and the number of consecutive duplicate ACKs (DUPACKs) that have been received. Upon reception of 3 DUPACKs, the TCP sender deduces that the link is congested and reduces the sending window. At this moment, the link estimation module provides the maximum bandwidth a TCP connection can achieve. In parallel, the currently used bandwidth per connection is estimated every RTT by low-pass filtering the rate of returning acknowledge packets from a peer. When an ACK is received by the source, it conveys the information that a certain amount of data corresponding to a specific transmitted packet was delivered to the destination. If the transmission process is not affected by losses, simply averaging the delivered data count over time yields a fair estimation of the bandwidth currently used by the source. When DUPACKs, indicating an out-of-sequence reception, reach the source, they should also count toward the bandwidth estimate, and a new estimate should be computed right after their reception. In addition the amount of UDP being sent needs to be accounted for. Each subflow over each connection is estimated individually and the sum of all estimations provides accurate current bandwidth estimation for each tunnel.

The latency and packet loss estimations are computed from the RTT and segment retransmission counters in the TCP kernel control blocks of each of the MPTCP subflows. This passive method allows evaluation of any path between an IUG and an ING at IP/transport layer in a technology and vendor agnostic manner without having to make assumptions on the underlying media. Another strategy could be to benefit from the implementation of the Dynamic Link Exchange Protocol (DLEP) [9] on the modems, which will then feed real time link parameters to the IUG and ING. However this protocol is currently not complemented nor widely adopted by modem manufacturers.

QoS-aware Link Selection Module

As the reference MPTCP implementation [10] always routes the traffic first towards the sub-flow with lowest RTT value, in BATS we have defined a novel QoS-aware Link Selection algorithm named "Path Selection based on Object Length (PSBOL)" that takes account of traffic requirements and the real-time status of each of the links. In the initial implementation, the link selection is performed by analyzing the so-called TCP object size, where an object is typically an Application Protocol Data Unit (APDU), e.g. an HTML request. The link selection algorithm is based on the assumption that for long objects (e.g. greater than a certain threshold, which can be selected dynamically based on the link characteristics and the class of traffic) the priority is to benefit from high bandwidth links to reduce the total transmission time, whereas for short objects is preferable to benefit from low latency and faster delivery times. Thus, for a particular TCP connection, long objects are routed via the highest bandwidth link (e.g. satellite) and short-objects are routed towards the lowest latency link (e.g. terrestrial). The discovery of objects boundaries is based on an algorithm that computes packets Inter-Arrival Times (IAT) and differentiates between the reception of packets from current or new objects. Based on this process, the size of the objects is updated on the fly with the reception of each TCP segment; hence, segments are initially considered as part of a small object and transmitted over the lowest RTT link. Once the object size exceeds the threshold, segments start to be transmitted over the highest bandwidth link.

Due to the IUG/ING architecture, optimized link selection decisions are performed independently for the inbound and outbound, which can benefit the routing of asymmetric traffic. Obviously, this approach is based on MPTCP and thus only optimized for TCP traffic. Link selection decisions for UDP traffic are based on policy based routing prioritizing first the links offering lowest latencies.

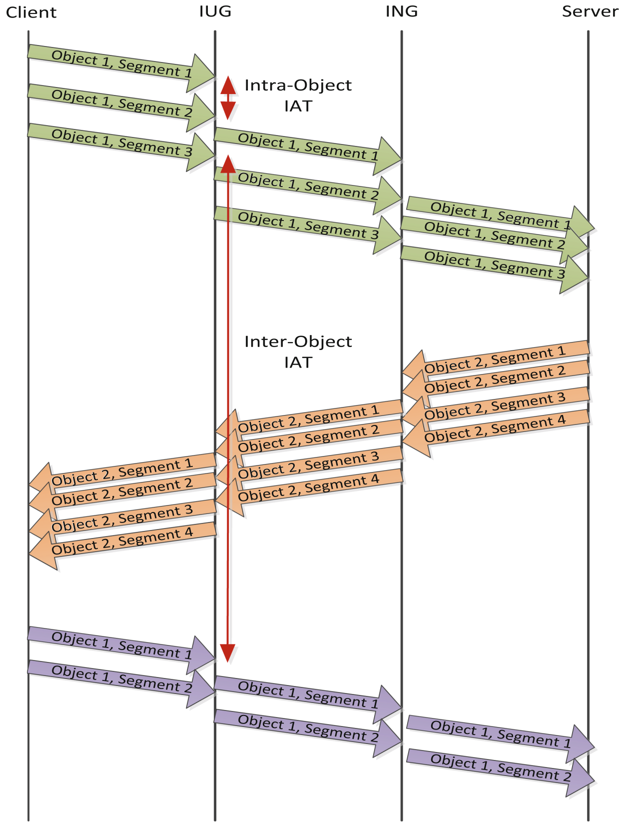

Fig. 4. TCP object boundary detection based on segment inter-arrival time

Traffic Classification Module

The proposed BATS solution identifies and sorts TCP payloads based on object characteristics. As aforementioned, the link selection decision for each packet is taken based on the size of the TCP object it belongs. It is therefore necessary to be able to compute the size of each particular object. In our implementation, objects are classified based on the IAT of the different TCP segments. The proposed process takes place in both IUG and ING and analyses the IAT of consecutive TCP segments belonging to a TCP connection. Hence, on the ING the process analyses the IAT for TCP segments received from Internet, and on the IUG the process analyses the IAT for TCP segments received from end-user devices. The algorithm compares the IAT of a segment to a threshold value to decide whether the segment belongs to the current object (Intra-object IAT being below threshold) or it is the first segment of a new object (Inter-object IAT above threshold), concept illustrated in Fig. 4. This threshold is computed from parameters that are locally available on the machine (IUG or ING) in order not to require protocols to exchange information between the two units. This method is applicable to any TCP-based application protocol and works on unencrypted as well as encrypted TCP connections.

3 Experimental Results

For the purpose of the lab and field trials of the BATS project, the IUG and ING have been prototyped on real routing platforms [11, 12]. As part of the in-factory functional, performance and stability validation tests, the software implementation of the different key building blocks has been evaluated first on a controlled environment with the IUG, ING, and network simulators for the different satellite and terrestrial paths implemented on Linux machines. Focusing on the Link Status Estimation module, Fig. 5 illustrates the performance of the bandwidth estimation when the satellite link is used and 20 parallel TCP connections are established between the two end-points. Note how during the test the bandwidth configured on the satellite simulator is reduced by half and the estimation curve is able to follow this variation. Similarly, Fig. 6 shows the output of the RTT estimation for a test during which the RTT value of the simulators is increased during a specific period of time. As we are measuring the RTT at the end-points, this accounts not only for the delay on the transmission link but also processing and buffering delays at the gateways.

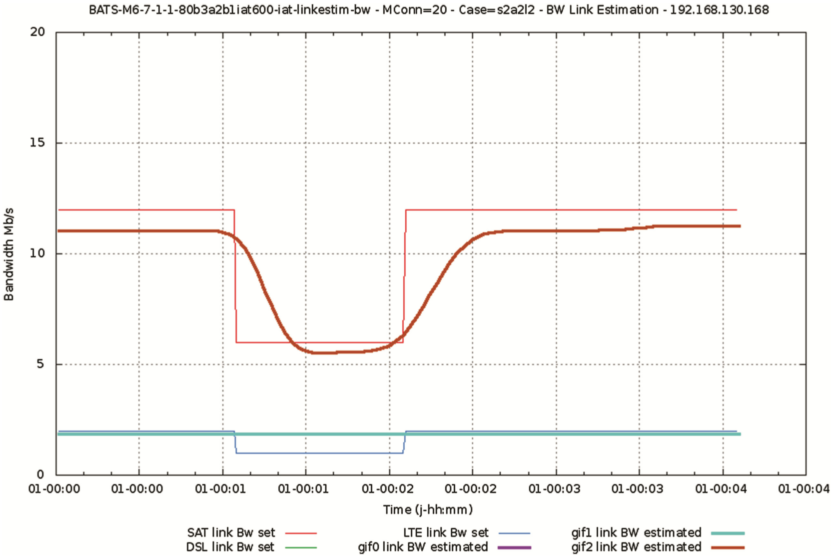

Fig. 5. Bandwidth estimation results on satellite interface with 20 parallel TCP connections *(thin red curve being real value and thick red curve being estimated value).*

In order to test the implementation in conditions as near as possible as a live experiment, the prototyped network appliances were connected to real networks (i.e., 7 Mbps satellite link, 500 kbps LTE link, and 1 Mbps DSL link) as illustrated in Fig. 7. Different configurations were tested: ADSL only, Satellite only, Cellular only, BATS implementation with Weighted Round Robin (WRR) as Link Selection algorithm with weights equivalent to each link's bandwidth, and finally BATS implementation with PsBOL as described earlier in this paper. Note that in these tests the threshold for classifying short

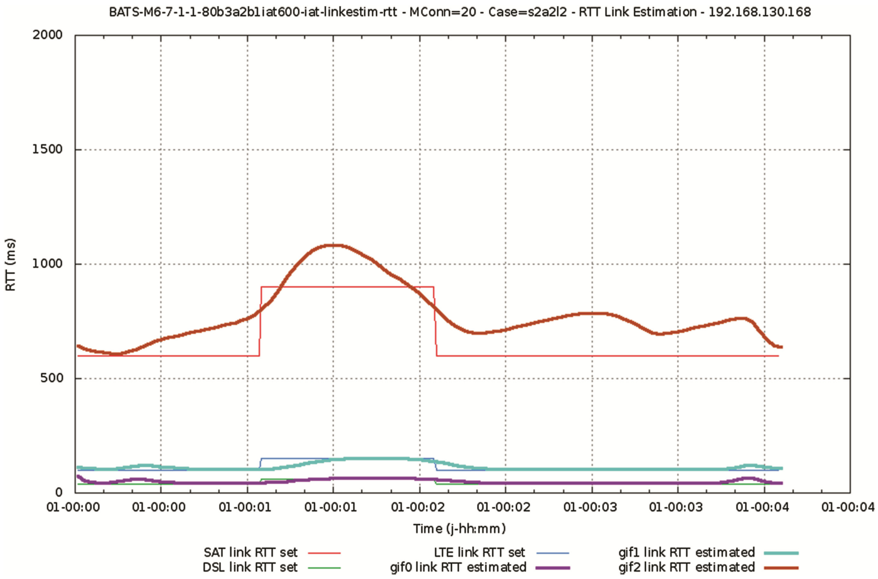

Fig. 6. RTT estimation results with 20 parallel TCP connections *(for the satellite link, thin red curve being real value and thick red curve being estimated value).*

and long objects was set to a default static value of 30 KB. In order to distinguish the BATS MPTCP implementation from the implementation in [10], it is referred in the figures as MCTCP. Figure 8 shows the result of a performance test consisting of loading 35 different websites. In each repetition of the test every website is loaded three times, results being the average of 9 repetitions. Figure 8 illustrates the average loading time and the standard deviation for each of the configurations. In this test, both configurations with the BATS architecture outperform the single link cases. However, no clear benefit is brought by PsBOL in comparison with WRR. This is due to the fact that the in WRR works well when weights are pre-defined in alignment with the available bandwidth in each connection. PsBOL shows a good performance when the link characteristics are not known a-priori or are varying over time.

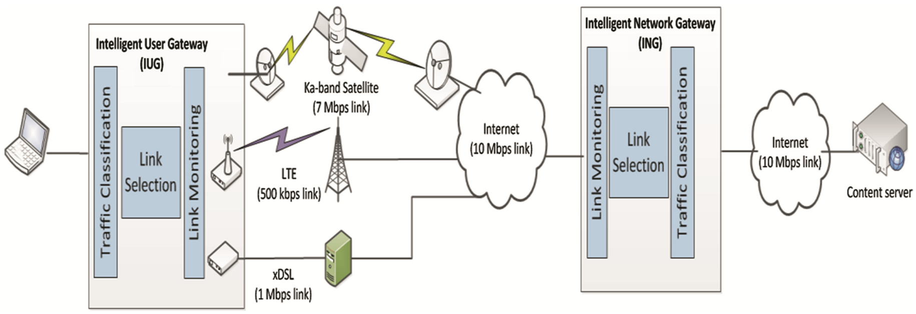

Fig. 7. Lab setup for in-factory tests with real networks

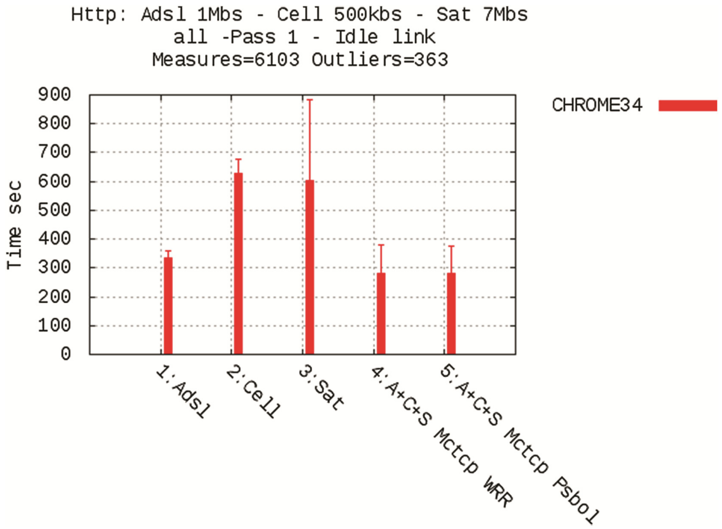

Fig. 8. HTTP acceleration test results with real networks

In the same setup, the performance of the system with SSH encrypted sessions has also been tested. Three different scenarios have been considered:

- Short Objects or Interactive: 60 commands generating small objects (e.g., ls, cd, etc.) are sent for a total amount of data equal to 48 KB.
- Long objects: 15 commands generating long objects (e.g., cat file, etc.) are sent with a total amount of data equal to 5.4 MB
- Mixed: the combination of both previous scenarios.

Table 1 details the time in seconds, averaged over 12 passes, that has been required to complete all the commands in the different scenarios and for the different configurations. As expected, for short objects BATS PsBOL matches the performance of the lowest RTT link (i.e., DSL) as it is the only one used for most of the time. For the scenarios with Long and Mixed-sized objects, PsBOL outperforms all other options benefiting from the intelligent use of the heterogeneous links.

Table 1. SSH acceleration test results with real networks (time is in seconds)

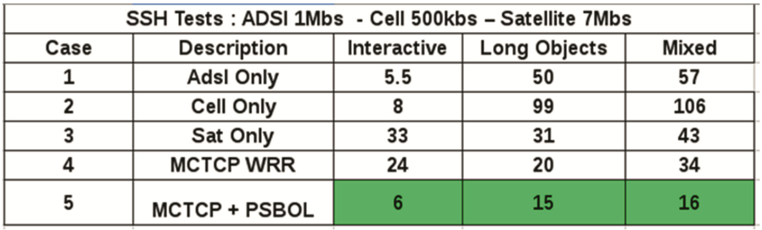

SSH Tests : ADSl 1Mbs - Cell 500kbs – Satellite 7Mbs

Case	Description	Interactive	Long Objects	Mixed
1	Adsl Only	5.5	50	57
2	Cell Only	8	99	106
3	Sat Only	33	31	43
4	MCTCP WRR	24	20	34
5	MCTCP + PSBOL	6	15	16

4 Conclusions

We have presented a novel broadband delivery system designed in the frame of the ongoing FP7 project BATS which integrates satellite with terrestrial access communication networks. We have described the design of integrated gateways in both user and operator sides of the system intelligently using all available access technologies to provide high speed broadband and improved QoE to users in under-served areas. We have discussed a specific implementation of the routing entities and provided results from the in-factory tests of the prototyped platforms.

Our future work will focus on running extensive Laboratory and Field trials with the prototype IUG and ING involving real end-users in both controlled and real-world environments to be able to assess the benefit of the designed system in terms of QoE.

Acknowledgments. This work has been supported by the BATS research project which is funded by the European Commission Seventh Framework Programme under contract number 317533. The authors of this paper would like to acknowledge the European Commission Seventh Framework Programme (FP7) and all the members of the BATS Consortium.

References

1. European Commission, A digital Agenda for Europe, COM(2010) 245, 26, August 2010
2. Point Topic, Mapping is vital to broadband investment, ITU Telecom World (2011)
3. BATS Project, Deliverable D5.2 “Cost Benefit Analysis”, April 2015. www.batsproject.eu
4. Pérez-Trufero, J., Peters, G., Watts, S., Evans, B., Dervin, M., Fesquet, T.: Broadband Access via integrated Terrestrial and Satellite systems (BATS), Ka and Broadband Communications, Navigation and Earth Observation Conference 2013, October 2013
5. Corbel, E., Peters, G., Sperber, R., Bosch, E.: TERASAT: high-throughput satellite system by 2020, Ka and Broadband Communications, Navigation and Earth Observation Conference 2013, October 2013
6. Peng, F., Cardona, A.S., Shafiee, K.: V.C.M Leung, TCP Performance Evaluation over GEO and LEO Satellite Links between Performance Enhancement Proxies, Vehicular Technology Conference (VTC) 2012, September 2012
7. Lee, D., Carpenter, B., Brownlee, N.: Observations of udp to tcp ratio and port numbers, in Internet Monitoring and Protection (ICIMP) 2010, pp. 99–104, May 2010
8. Ford, C., Raiciu, Handley, M., Bonaventure, O.: TCP Extensions for Multipath Operation with Multiple Addresses, RFC 6824 (Experimental), Internet Engineering Task Force, January 2013
9. Ratliff, S., Berry, B., Harrision, G., Jury, S., Satterwhite, D.: Dynamic Link Exchange Protocol (DLEP), Internet Draft, Internet Engineering Task Force, February 2014
10. Paasch, C., Barre, S., et al.: MultiPath TCP – Linux Kernel implementation. http://www.multipath-tcp.org/
11. OneAccess Networks, One540: Cloud Services Gateway & Multi-Service Access Router, Datasheet. http://www.oneaccess-net.com/products/multi-service-data-routers-and-eads/medium-business-regional-offices/item/69-one540
12. OneAccess Networks, WXD2450: Wan Optimization Platform, Datasheet. http://www.oneaccess-net.com/products/wan-optimization-platforms/item/131-wxd2450