

Switched Ethernet Networking over LEO Satellite

Vincenzo Mancuso^{1,2}, Giuseppe Bianchi¹, Nicola Blefari Melazzi¹, Ulla Birnbacher³

¹Department of Electronic Engineering, University of Rome “Tor Vergata”, Rome, Italy

²Department of Electric, Electronic and Telecommunication Engineering, Università di Palermo, Italy

³Institute of Communication Networks and Satellite Communications, Graz University of Technology (TUG), Graz, Austria

Abstract - Satellite onboard switching capabilities, and their possibility to enable “over-the-air” networks, is a strongly innovative achievement of modern satellite communication. The exploitation of satellite networking is especially relevant in the case of Low Earth Orbit satellites (LEO), where the reduced latency makes a LEO mesh network highly appealing and competitive with terrestrial solutions also in terms of performance. However, while the idea of a LEO mesh network is well understood, we are far from an ultimate answer on which is the best and more appropriate networking solution for moving LEO satellite network nodes. This paper promotes the adoption of existing layer-2 networking solutions over LEO networks. We believe that LEO networking may leverage from already existing and well understood layer-2 switching mechanisms, thus allowing, on one side, onboard deployment of switching fabrics and solutions already developed for terrestrial networks, and, on the other side, seamless interoperation with terrestrial networks.

Specifically, as technical contribution of this paper, we show how 802.1Q switching solutions (i.e. Ethernet Virtual LAN switching) may be adapted, with limited modification effort, to operate in a LEO scenario. Moreover, we propose an innovative approach for frame routing over a mesh of multiple spanning trees (802.1s), especially devised to dramatically enhance the performance of legacy 802.1Q solutions in the presence of moving LEO nodes.

I. INTRODUCTION

By the use of communication satellites, networks can be interconnected on a global scale, also in geographic regions with poor terrestrial backbone infrastructure. Satellites orbiting at low earth orbits, i.e. 500 to 2000 Km above the Earth surface, are called LEO satellites. Compared to communication satellites in geostationary orbit, the communication links to LEO satellites are characterized by lower propagation delay and lower link attenuation (free space loss) because of the shorter distance, resulting in the need for reduced transmission power or the usability of antennas with low directivity. Thus, LEO satellites are a good candidate to provide access connectivity and to connect remote LANs/MANs together, while providing support for high speed and real time services.

Unluckily, because of the low orbit, the communication link between a terrestrial terminal and a LEO satellite is available only for a few minutes. Hence, a fleet of many LEO satellites is needed for continuous service operation. Using inter-satellite links, a meshed network can be setup in space, enabling the establishment of communication paths between users on Earth.

However, seen from the point of a terrestrial user, the network topology is changing frequently, i.e. links to the satellites in view have to be setup and torn down regularly. Nevertheless, topology changes for the user can be calculated a priori from the known satellite motion.

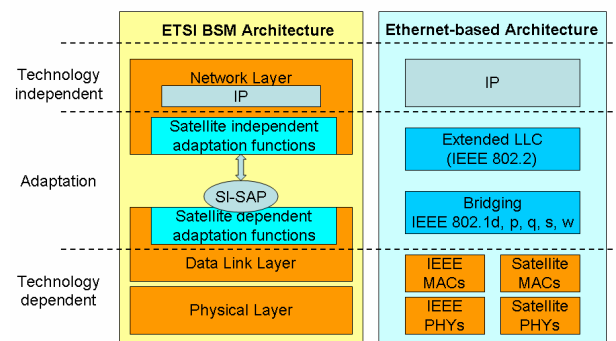


Figure 1 – ETSI BSM and Ethernet-like architectures

As for the transport of data over satellite networks, protocol stacks have been highly optimized for the given physical characteristics. However, as IP has become the dominant networking protocol, satellite dependent protocols have been reduced to Layer-1 (physical layer) and medium access part (MAC) of Layer-2. The resulting protocol stack is summarized by the ETSI BSM scheme [1][2] depicted in the leftmost side of Figure 1. That scheme has been presented by the satellite standardization community as the reference architecture, able to allow IP over satellite. In particular, it is clear from Figure 1 that a special function is needed that allows the interworking of technology-independent protocols, such as the TCP/IP protocol suite, and technology-dependent protocols, such as the satellite MAC.

It is worth noting that, in the scope of IEEE 802 project, this interworking task is performed by the Logical Link Control sub-layer, i.e. the higher part of the Layer-2 (LLC [3]). Our feeling is that the adaptation layer for “next generation” LEO satellites should be as similar as possible to LLC, with additional support for specific satellite MACs. Hence, we promote the adoption of regenerative LEO satellites endowed with switching/bridging facilities that permit a satellite interconnection with terrestrial devices operating in the higher part of Layer-2, such as Ethernet bridges and switches. In this way, satellite and terrestrial networks could be transparently merged at Layer-2. This convergence brings to an overall network that might be very complex but that should appear as simple as possible to users and administrators alike. An example of the resulting networking architecture is sketched in the rightmost side of Figure 2.

The rest of the paper is organized as follows: section II explains the rationale of the work, section III shows the switched satellite networking solution, and section IV gives insights on the impact of our solution on network services based on TCP and UDP transport protocols. Simulative results are presented in section V. Section VI concludes the paper and presents some open issues.

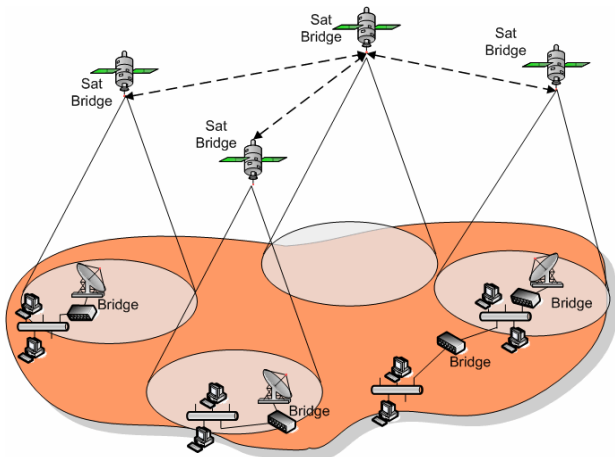


Figure 2 - Switched network with LEO bridges

II. RATIONALE OF THE WORK

The satellite network community made a significant research effort in the attempt of supporting TCP/IP networking over satellite links [4][5]. This is not an easy goal. TCP/IP protocols suffer from environmental specificities like: i) long propagation delays; ii) error prone links; iii) intermittent links.

However, satellites show some fundamental advantages, such as a distinctive broadcasting capability and potentially large service areas. Furthermore, in the case of LEO satellites, long delays and frequent errors are less of a problem, but a fleet of interworking satellites is required due to the lower orbit. As a consequence, edge-to-edge satellite paths frequently change and handover procedures are needed for inter-satellite links to provide continuous connectivity to ground stations. Actually, satellite fleet can be designed so as most of the links between satellites are stable, and, as previously stated in Section I, topology changes can be calculated a priori from the known satellite constellation.

In the following sections III to V, we show how the adoption of layer-2 switching over LEO satellite paths is a viable solution, and may be considered as an alternative to IP routing. Specifically, we propose to endow LEO satellites with IEEE 802-like bridging facilities and Logical Link Control functions. This results into: i) easier interoperability with existing devices; ii) transparency of satellites technologies as perceived by terrestrial users; iii) end-to-end service continuity; iv) protocol overhead reduction.

It is worth noting that, in case of satellites with routers onboard, they should deal with IP packets and operate a forwarding based on network-wide routing information.

Even though it could appear more elegant and appealing, the IP routing has a negative impact on network performance, especially when the topology is not static. Routing operation requires parsing of IP packets, to be extracted from MAC frames, and handling fragmentation and encapsulating IP packets in MAC data units. More important, routing information has to propagate in the network in order to be effective. Thus, upon topology changes, not only satellite routing tables have to be drastically updated, but also a

number of terrestrial routers have to learn new routes and delete old ones. With LEO satellites topology changes every few minutes and IP routing tables would be out-of-date for the majority of time. Conversely, with satellite switches, the satellite payload duty is to switch MAC frames in a limited sub-network (i.e. a single administrative domain), on the basis of the information conveyed by the MAC header.

We conclude that satellites could be endowed with routing devices (Layer-3 in the OSI protocol stack), but satellites should be managed as a separate sub-network, where switching/bridging devices (Layer-2) are more appropriate and efficient, and do not waste time and bandwidth for the extraction of IP datagrams from MAC frames.

In promoting an alternative approach based on Layer-2, our design goal has been to reuse, as much as possible, standard LAN/MAN concepts. Indeed, the IEEE 802 standardization activity has produced several fundamental results that can be usefully exploited in a satellite-based networking extension of terrestrial solutions. In particular, the adaptation layer envisioned in [1] can be provided by the IEEE 802.2 (LLC), with the support for bridging and switching standards: IEEE 802.2D, describing bridges features and the Spanning Tree Protocols (STP) [6], and STP's enhanced versions optimized for fast reconfiguration (Rapid STP [7]), and for logical network splitting (Multiple STP [8]). Other Layer-2 mechanisms that may be used in this context are the Full Duplex Switched Ethernet [9], and the virtual LAN bridging capabilities proposed by IEEE (i.e., VLANs [10] managed through Multiple Spanning Trees - MST).

However, existing standards could turn out to be not efficient enough, when ported over satellite without any modification. In particular, mobility management has to be allowed for. This leads to the adoption of special add-ons to existing spanning tree algorithms, VLAN forming and management issues, and also to QoS management tools.

In fact, implementing Ethernet-like mechanisms over satellite requires an additional effort in order to: i) maintain a high throughput; ii) minimize delay due to elaborations; iii) bound loss rate; iv) reduce reconfiguration costs.

III. MULTIPLE SPANNING TREES

Legacy (terrestrial) bridging protocols adopt reconfiguration utilities that timely check or notify network topology changes. However, bridge reconfiguration could require unfeasible long times: using STP, it may take several tens of seconds, depending on the network diameter. Rapid Spanning Tree can reduce the reconfiguration time to few seconds, or also to fractions of second for very simple connectivity changes.

Moreover, upon the reconfiguration of a spanning tree, some new bridges are likely inserted in the tree, which do not have information about users already exchanging frames. This turns in a "learning phase" during which those new bridges flood the network in order to correctly forward frames to users that they had never seen before.

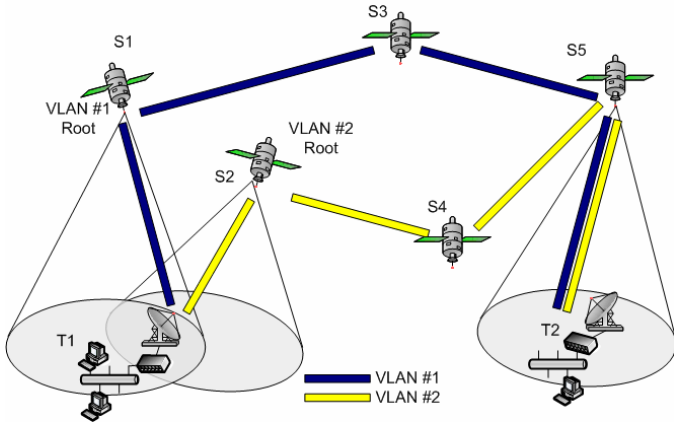


Figure 3 - Possible network topology with two VLANs available

Both effects have a relevant impact on many applications based on reactive and not reactive transport protocols (e.g., TCP and UDP respectively). For instance, discontinuities strongly affect the quality of streaming, and introduce flooding effect after handovers. Indeed, they cause TCP losses and slow-down its throughput. Thus, to improve performance, it is advisable to force bridging reconfiguration procedures in a predictive way, without waiting for a network error, and preventing long starvation periods needed to recognize the topology change and to operate the bridging reconfiguration.

To solve the issue of topology reconfiguration, we have ideated and evaluated the following mechanism. At the edge of the satellite network, frames are tagged according to a VLAN identifier. Unlike legacy (terrestrial) VLANs, tagging is not performed statically and with the goal of partitioning the network into virtual LANs. Conversely, each VLAN is associated with a specific Spanning Tree. Hence, tagging implies that frames will be routed across the Spanning Tree associated to that VLAN. Dynamic tagging is a key idea which allows to avoid the need for real-time route changes. In fact, satellite visibility changes are gradual and the LEO fleet can be designed so that any ground station has at least one LEO satellite in visibility all the time, while, during a topology transition, the ground station can communicate with at least two satellites. During a topology transition phase, at least two data paths will be available for a given end-to-end communications. We propose to label the alternate paths with different VLAN tags: end-users could exchange data over a VLAN before the topology change occurs, and switch to another VLAN during topology transition, with such transitions being predictable. In this manner, logical end-to-end paths are maintained even during physical path (namely, Spanning Tree) reconfiguration. The scenario is reported in Figure 3, where, for the sake of simplicity, a reduced number of nodes and VLANs is reproduced, and a single end-to-end communication is considered.

As a final consideration, note that our solution can be deployed over off-the-shelf bridges, which can be controlled through already existing protocols (e.g., the proprietary CISCO's VTP [11], or the standard-based GVRP [12]).

UDP		
Event	Action (No VLAN)	Action (two VLANs)
Service request from Client	Flooding	Flooding in VLAN#1
	Bridges learn Client address	VLAN#1 bridges learn Client address
Service response from Server	Switched/no flooding	Switched/no flooding
	Bridges learn Server address	VLAN#1 bridges learn Server address
Service data	Switched/no flooding	Switched/no flooding
Topology change	Re-compute ST	Switch to VLAN #2
	LAN temporarily disconnected	Re-compute CIST and VLAN#1 ST
	All frames discarded	VLAN#1 temporarily disconnected
		VLAN #2 flooded, filter. BDs are empty
Downstream (Upstream) frames after reconfiguration	Network flooded by "new" bridges until an up- (down-) stream frame is sent by Client (Server)	VLAN#2 flooded by all bridges until an up- (down-) stream frame is sent by Client (Server)

Table 1 - How topology changes affect UDP connections

TCP		
Event	Action (No VLAN)	Action (two VLANs)
Request from Client (TCP SYN)	Flooding	Flooding in VLAN#1
	Bridges learn Client address	VLAN#1 bridges learn Client address
Response from Server (TCP ACK)	Switched/ no flooding	Switched/ no flooding
	Bridges learn Server address	VLAN#1 bridges learn Server address
TCP SIN-ACK	Switched/ no flooding	Switched/ no flooding
Service data flow	Switched/ no flooding	Switched/ no flooding
Client's ack	Switched/ no flooding	Switched/ no flooding
Topology changes	Re-compute ST	Switch to VLAN #2
	LAN temporarily disconnected	Re-compute CIST and VLAN#1 ST
	All frames discarded	VLAN#1 temporarily disconnected
Downstream (Upstream) reconfig.	Flooded by "new" bridges until an up- (down-) stream frame is sent by Client (Server)	Flooded in VLAN#2 bridges until an up- (down-) stream frame is sent by the Client (Server)

Table 2 - How topology changes affect TCP connections

IV. IMPACT ON NETWORK SERVICES

Multiple VLANs allow the network provider to hide network topology changes. In fact, during a topology transition, a new path will be available *before* the old path goes down. We include these paths within different VLANs and switch to new path *during* the topology transition.

Table 1 reports what happens to frames generated in a message exchange, when UDP is used, comparing the case in

which a VLAN “hopping” is adopted when topology changes occur to a scenario without such feature. We consider the worst case scenario, where only two users are in the network and only one connection is active. Summarizing Table 1, we can say that the absence of VLANs implies service discontinuities and long flooding phases after reconfigurations (due to the unidirectional nature of UDP exchanges). Using VLANs, discontinuities of the service are avoided: packets are not filtered and end-to-end connectivity is not broken. The flooding after reconfiguration is bounded to the VLAN used after the handover, while an STP-based approach would flood the entire network.

Similar considerations can be made when TCP is used, with the exception of the occurrence of short flooding phases, due to frequent TCP ACKs in the reverse path. Table 2 collects a set of actions performed by network entities at the occurrence of specific TCP-related events (beginning and tearing down of a session, sending of data and acknowledgments, and so on). The advantage of using the management approach proposed in this paper is clear from the comparison with the “legacy” behaviour of bridges as described in the central column of the table, and the behaviour of the MST with support for off-line reconfiguration, as described in the rightmost column: VLANs reduce flooding effects and off-line VLAN reconfiguration eliminates service discontinuities.

With respect to the adoption of “legacy” VLANs, the switching between multiple VLANs does not allow saving of bandwidth after a topology change. In fact, each bridge builds up a separate filtering database for each VLAN. So, when two end-users switch from a VLAN to another, they are not known to some bridges until they send a frame. In the meantime, VLAN flooding occurs. To remedy this impairment, different solutions could be adopted. For instance, each bridge filtering database could be remotely configured by the network manager, even though this approach implies the knowledge of all active connections. A simpler solution could consist in tagging a single frame per user within the new VLAN before reconfiguration: that frame will set a new entry in any bridge, and avoid subsequent flooding of several frames.

V. PERFORMANCE EVALUATION

We simulated a scenario similar to the one depicted in Figure 3, by means of the OPNET simulator [13] with modified bridging devices. Two terrestrial bridges are adopted and a LEO fleet provides remote coverage. Users are located close to the terrestrial bridges; the longest path between two users has a delay of 100 ms, with at most one inter-satellite link. We tested both UDP and TCP-based applications with and without VLAN supports, also considering the effect of STP and RSTP when VLANs are not used. The capacity of satellite links was set to 100 Mbps, and the offered load was 1.0 Mbps both for TCP and UDP-based applications: we did not consider heavy load, since our aim was to highlight performance figures like service continuity and path discovery in the spanning tree.

As a matter of fact, simulations confirm the behaviour described in Table 1 for UDP, and the considerations about

TCP in Table 2. Figure 4 to Figure 6 depict the throughput of an unidirectional UDP connection between two remote hosts. In the simulations, a topology change occurred at $t=200s$, and one can notice that a traditional STP approach requires up to 45s to recover the path; using RSTP this time is shortened but several seconds are still required to reconfigure the large bridged network. On the contrary, preconfigured VLANs allows a seamless handover, without service discontinuities.

Figure 7 depicts the throughput of a TCP connection between two remote hosts. In the simulations, a topology change occurred at $t=200s$, and by using RSTP up to 10s are necessary to recover the path. Again, off-line configured VLANs allow ground stations to seamlessly switch between VLANs, and avoid service discontinuities. It is also remarkable the modulation effect that the path length variation and the unconnected periods, due to LEO orbits, operates on TCP throughput. As well as for service continuity, the adoption of off-line managed VLANs has a beneficial impact in terms of smoothed TCP throughput variations.

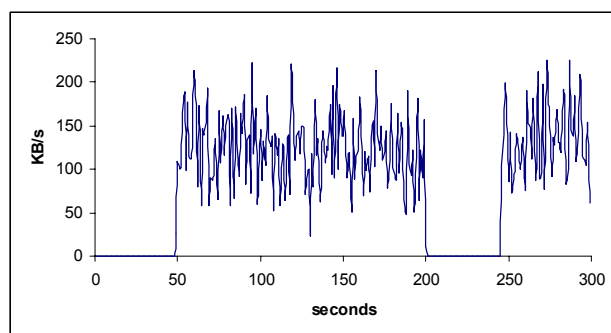


Figure 4 - UDP throughput with STP, no VLANs




Figure 5 - UDP throughput with RTSP, no VLANs




Figure 6 - UDP throughput with VLAN hopping

As for the flooding effects due to topology changes, Figure 8 and Figure 9 show results for TCP connections supporting a FTP server. Inter-arrival time between requests is exponential distributed and file size is uniformly distributed (between 1.0 and 1.5 MB), so that the load is equal to 1.0 Mbps¹. Topology changes occur each 310s. The burden of flooding data is depicted in Figure 8 when VLANs are not used: a few packets are broadcast by new bridges after each topology change, but BPDUs² are continuously exchanged in order to detect the variation of topology. Conversely, BPDUs are no more needed in case of preconfigured VLANs, so that Figure 9 shows that only a few packets are broadcast after each VLAN change, which corresponds to a topology change.

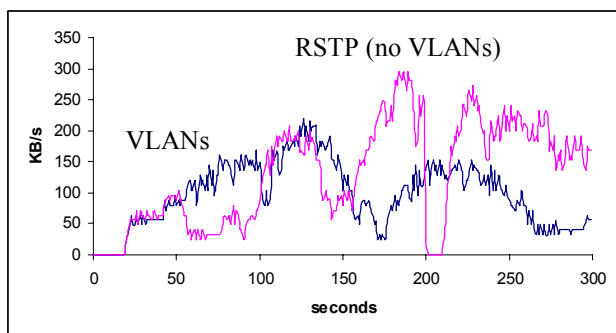


Figure 7 - TCP throughput with RTSP, and using VLANs




Figure 8 - TCP with STP (no VLANs): packets broadcast by at least one bridge due to filtering database dynamics

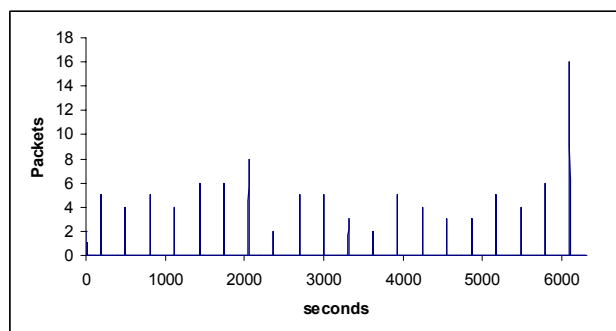


Figure 9 - TCP with VLANs hopping: packets broadcast by at least one bridge due to filtering database dynamics

¹ It is worth noting that if $RTT=2*100ms$, then the actual TCP throughput is about 2.6 Mbps, without any TCP modification to transmission window.

² Bridge Protocol Data Units (BPDUs) are control frames exchanged between bridges with a fixed (slow) rate (one BPDU each 2.0 seconds is the value suggested by the IEEE 802.1D standard and its amendments).

VI. CONCLUSIONS AND OPEN ISSUES

Our aim is designing a network that can operate in an Ethernet-like mode, being flexible and transparent to the user and to the upper layer protocols. In our vision, satellite switches are just like terrestrial legacy switches, with the only difference that they are located into a satellite payload and they move in a regular way. We assume that a satellite switch can be connected to each terrestrial switch in its coverage area. In this sense, a satellite switch can rely on multiple port operation. The rationale of pursuing an Ethernet-like solution is to go towards a full interoperability with most of existing equipments, e.g. IP terminals, able to see each other as belonging to the same LAN (VLAN), with layer 2 switches being both terrestrial nodes and satellite nodes, while avoiding IP routing intricacies. A special LLC, and the reuse of IEEE bridging facilities will permit the convergence of terrestrial and satellite-oriented MACs.

Another important issue is the possible introduction of Full Duplex Switched Ethernet (FDSE) techniques, which is by now a pervasive technology in recently deployed access network, and proposed also for future xDSL implementations.

The simulation campaign presented in this paper shows potentialities and feasibility of the proposed VLAN-based networking, and its positive impact on protocols and services.

ACKNOWLEDGMENT

This work has been co-funded by the European Commission under the FP6 IST Program, in the framework of the SatNex Network of Excellence.

REFERENCES

- [1] ETSI-TR-984: ETSI, Satellite Earth Stations and Systems (SES). Broadband Satellite Multimedia. Services and Architectures. ETSI Technical Report, TR 101 984 V1.1.1, Nov. 2002
- [2] ETSI-TR-985: ETSI, Satellite Earth Stations and Systems (SES). Broadband Satellite Multimedia. IP over Satellite. ETSI Technical Report, TR 101 985 V1.1.2, Nov. 2002
- [3] ANSI/IEEE Std 802.2, 1998 Edition, "Part2: Logical Link Control"
- [4] Byungsook Kim, Isil Sebüktekin, "An integrated IP QoS architecture - Performance," Proc. IEEE Military Commun. Conf. MILCOM 2002, Anaheim, CA, Oct. 7-10, 2002 pp. 1189-1193
- [5] M.Marchese, M.Mongelli, "On-line Bandwidth Control for Quality of Service Mapping over Satellite Independent Service Access Points", Elsevier Int. Journal of Computer and Telecommunications Networking
- [6] ANSI/IEEE Std 802.1D, 1998 Edition, "Part 3: Media Access Control (MAC) Bridges"
- [7] IEEE Std 802.1w-2001, "IEEE Standard for Local and metropolitan area networks - Common specifications Part 3: Media Access Control (MAC) Bridges - Amendment 2: Rapid Reconfiguration"
- [8] IEEE Std 802.1s™-2002 Edition, "802.1s - IEEE Standards for Local and metropolitan area networks Virtual Bridged Local Area Networks - Amendment 3: Multiple Spanning Trees"
- [9] T. Kuri, K. Kitayama, Y. Takahashi, "A single light-source configuration for full-duplex 60-GHz-band radio-on-fiber system" Microwave Theory and Techniques, IEEE Transactions on Volume 51, Issue 2, Feb. 2003 Page(s):431 - 439
- [10] IEEE Std 802.1Q™, 2003 Edition, "802.1Q - IEEE Standards for Local and metropolitan area networks - Virtual Bridged Local Area Networks"
- [11] <http://www.cisco.com>: Understanding and Configuring VTP
- [12] <http://www.alliedtelesyn.co.nz/documentation/at8700/261/pdf/garp.pdf>
- [13] The official OPNET internet site: www.opnet.com