

NSIS-Based Quality of Service and Resource Allocation in Ethernet Networks

M. Carmo, B. Carvalho, J. Sá Silva, E. Monteiro, P. Simões, M. Curado and F.
Boavida

University of Coimbra,
Laboratory of Communications and Telematics
DEI / CISUC

Polo II, 3030 Coimbra, PORTUGAL

Tel.: +351 239 790000 Fax: +351 239 701266

{maxweel, sasilva, edmund, psimoes, marilia, boavida}@dei.uc.pt,
becarv@student.dei.uc.pt

Abstract. The definition of the IEEE 802.1Q and IEEE 802.1p standards provided Class of Service (CoS) capabilities to Ethernet networks and, consequently, allowed new QoS services to be deployed. This is used by the Subnet Bandwidth Management (SBM) protocol, an RSVP-based protocol that provides IntServ-like services at Ethernet level. This paper proposes an alternative way to quality of service provision and resource allocation in Ethernet networks, based on the emerging IETF NSIS framework. The proposed approach was validated as a proof of concept by simulation, showing the ability of NSIS to provide QoS differentiation in Ethernet scenarios.

1 Introduction

Ethernet is a network technology supported by a set of international standards that offer pragmatic solutions for communication. Its success is due not only to its low-cost but also to its simplicity.

Ethernet solutions are widely used in local area domains. However, with the development of Gigabit and 10 Gigabit variants, Ethernet seems to have gained a new momentum as a technology for use in all telecommunication networks. This, along with the fact that it is a flexible and switched technology, can be regarded as an omen of the success of Ethernet in future communication systems, especially when large networks are concerned. In fact, already nowadays Ethernet is not only used in Campus and LANs networks, but also in Small Office Home Office Networks, Metropolitan Area Networks (MAN) [1], Wide Area Networks (WAN) and MAN residential scenarios (Ethernet Passive Optical Networks - EPONs).

This paper presents a proposal for the support of QoS and resource allocation in Ethernet networks, developed in the scope of the EuQoS project (<http://www.euqos.org>).

EuQoS is an FP6 IST Integrated Project with the aim of proposing, developing and studying end-to-end QoS support for Internet applications. This will be achieved through the research, integration, testing, validation and demonstration of end-to-end QoS technologies that can support advanced QoS-demanding applications such as voice, video-conferencing, video-streaming, educational, tele-engineering and medical applications.

EuQoS targets a wide range of network technologies, from access networks – including Ethernet, UMTS, and WiFi – to core networks. As such, one of the research lines of the EuQoS project consists of the development of solutions that provide QoS in Ethernet networks.

The proposals presented in this paper were studied by simulation and are currently being refined in order to serve as a basis for prototyping. In order to adequately fundamnet and present the developed as well as the on-going work, the paper is organized as follows: the next section presents a global view of QoS in Ethernet, identifying the base standards and providing the motivation for extending QoS to such networks; Section 3 provides background and related work, by presenting the current approach to resource allocation in Ethernet, namely the RSVP-based Subnet Bandwidth Management protocol (SBM); Section 4 presents the main proposal of this paper, addressing the use of NSIS for resource allocation in Ethernet and presenting a simulation-based validation of this proposal; Section 5 summarizes the contributions and identifies the guidelines for further work.

2 Quality of Service in Ethernet

This section presents the base standards for QoS support in Ethernet and discusses the motivations for and support of such developments.

2.1 Standards Support

CSMA/CD is the media access control mechanism that was initially developed to give the possibility for two or more devices to share a common media. This mechanism worked well for 10 Mbps but revealed problems at higher data rates. Nowadays, CSMA/CD networks are hardly used. With the advent of full-duplex and switched solutions, it is possible to build collision-free tree and star topologies, connecting terminal equipment to switches.

Original Ethernet standards did not support QoS functionality. IEEE 802.1Q [2] can be considered a first step in this direction, by defining the operation of Virtual LAN (VLAN) Bridges.

The IEEE 802.1p standard (which is part of IEEE 802.1D [3]) supports the provisioning of expedited traffic in a LAN network, based on the use of VLAN tags. VLAN tags have two parts: the VLAN ID (12-bit) and the Prioritization field (3-bit). IEEE 802.1p defines this latter field, allowing the prioritization of traffic into 8 levels, thus providing the basic support for QoS differentiation in Ethernet.

The user priority values defined in the IEEE 802.1p standard are presented in Table 1. Each level is associated with a specific traffic type.

Table 1. User priority values recommended by IEEE 802.1p standard

IEEE 802.1p User Priority	Traffic Type
7 (highest)	Network Management
6	Voice
5	Video
4	Controlled Load
3	Excellent Effort
0	Best Effort
2	Undefined
1 (lowest)	Background

2.2 Rationale for QoS in Ethernet

It is important to distinguish the use and need of QoS mechanisms in the different types of Ethernet networks. Although the technology principles are the same, the purpose of implementing QoS policies can be different.

A SOHO, also called a *virtual office*, refers to the small office or home environment and the business culture that surrounds it. In this type of environment, the implementation of QoS mechanisms in order to differentiate traffic inside the network is not critical, as the number of active pieces of equipment is generally low.

Ethernet has found widespread acceptability in campus and small scale environments, and nowadays is the supporting technology of most LANs. With the use of different applications, with a wide range of requirements, it is crucial to differentiate the traffic and to manage it properly.

Current MAN infrastructures are built not only for a voice-centric world but also for data-centric world. However, in order to use Ethernet to implement MANs, it is necessary to offer and maintain the same level of QoS of traditional voice-based applications. The use of over-provisioned solutions, typical in some of the recent systems, leads to low resource utilization, high costs and poor scalability.

The main reason Ethernet is being considered in WANs is because of its low cost. Initially, Ethernet will be used as a switch-to-switch interconnection for multiplexing different traffic, offering inter operation between different vendors' implementations. The definition of new standards and policies is crucial to the mapping of traffic classes between different operators.

The high modularity and scalability of Ethernet solutions enables operators to supply broadband services to subscribers, including data, voice and video, through a cost effective architecture in residential networks, namely through the use of EPONs. However, in these networks, each link supports a set of subscribers, which produce

very bursty traffic in contrast with MANs and WANs, where the bandwidth requirements are relatively smoother due to the aggregation of many traffic sources.

Given the variety of environments to which Ethernet is applicable, which nowadays covers the whole spectrum of networks, it is clear that if quality of service is to be provided to applications, it must be supported in this type of technology.

2.3 Support for QoS in Ethernet

Nowadays, there is a huge number of pieces of equipment, such as switches and network interface cards, that provide some level of IEEE 802.1p implementation. In addition, there are pieces of software, like device drivers and APIs that support the provisioning of QoS through the 802.1p standard.

Several network cards available in the market also support the 802.1p specification, allowing hosts to set the user priority of the frames, using appropriate operating system device drivers and routines.

Several operating systems provide traffic prioritization support at the application and host levels. Applications can use appropriate APIs to handle the 802.1p parameters, and network administrators can use traffic management tools to apply QoS requirements in applications that are not QoS-aware.

In [4] a MIB module related to the management of MAC bridges (IEEE 802.1D) is defined, specifying objects for the management of Traffic Classes (IEEE 802.1p) and Enhanced Multicast Filtering. [4] also defines a MIB module for the management of Virtual LANs, specified by the IEEE 802.1Q standard.

This diversity of implementations and support confirms the acceptance of the 802.1p specification by the market.

3 RSVP-Based Resource Allocation in Ethernet

The traffic class differentiation provided by the 802.1p, together with some signaling and admission control mechanisms, can provide a close approximation to the Integrated Services (IntServ) paradigm at the level of Ethernet.

The IETF Integrated Services over Specific Link Layers (ISSLL) Working Group [5] was responsible for the mapping of QoS mechanisms of upper layers into layer 2 (link layer). A direct result of this work was the definition of the Subnet Bandwidth Management protocol (SBM) [6], a signaling protocol for shared and switched IEEE 802-style networks, like Ethernet.

Specifically, the SBM protocol describes a method of performing the reservation of LAN resources for RSVP-enabled data flows along the path where switches and routers offer support to do it.

An important module of the SBM definition is the Designated SBM (DSBM), a protocol entity that resides in a switch/router and is responsible for the admission control in a network segment.

All the RSVP request messages crossing the Ethernet network are redirected to the DSBM entities that will check the availability of the required resources. Figure 1 presents a LAN where the router needs to forward an RSVP PATH message to the host. In case (a) the signaling message reaches the host without take into account details of the Ethernet network.

In case of a SBM scenario (b), the signaling message is sent to the DSBM entities, instead of the destination node. After processing the request and adding themselves to the path, the DSBMs forward the message until it reaches the host. In the same way, RSVP RESV messages, crossing the managed segment, also need to pass through the DSBMs entities.




Fig. 1. RSVP and SBM signaling on Ethernet LAN scenarios

Each DSBM entity maintains information about the available resources on a given segment, being responsible for sending a RESV_ERR message to the requester if the request cannot be granted.

To ensure the existing of resources for the flows inside a LAN, [7] describes an approach to map the characterization parameters (e.g. bandwidth and delay bounds) [8] and services defined in the IntServ model (e.g. controlled load and guaranteed services) to the Ethernet traffic class parameters (Section 2.1). This mapping is performed in such a manner that different traffic flows requiring similar grade of service are aggregated into the same traffic class (with the same user priority value).

4 NSIS-Based Resource Allocation in Ethernet

This section presents the proposed resource allocation approach. As this approach is based on the use of Next Steps in Signaling (NSIS) [9], the section starts with a brief description of this IETF framework. Then, the basic proposal is presented and justified. The section ends with some results of a simulation-based evaluation.

4.1 Next Steps in Signaling

NSIS is a signalling framework being developed by the IETF, based on various signalling protocols, of which the RSVP is the corner stone. This framework is used for application signalling, in order to install and maintain flow state in the network, similar to other protocols such as the aforementioned RSVP.

NSIS can work on a per-flow basis, although it allows for flow aggregation based on the DSCP field of IP headers or on tunnels. In addition, NSIS works on a hop-by-hop basis (using NSIS-aware nodes in the network), as opposed to an end-to-end approach. The states related with the data flows are handled by NSIS-aware nodes, but this does not mean that every node in the network must support NSIS. NSIS has been specified in order to be usable in different parts of the Internet, for different needs, without requiring a complete end-to-end deployment.

NSIS considers two paradigms for resource reservation signalling, referred to as path-coupled and path-decoupled.

According to the path-coupled paradigm, the signalling messages are routed through NEs (NSIS Entities) that are on the data path only. These messages can use various addressing styles, with messages either explicitly addressed to adjacent NEs (normally referred to as NSIS hop) on-path, or addressed to an NSIS receiver with the RAO option (Router Alert Option) in the data packets, allowing the messages to be intercepted by other NEs along the data path. This allows NEs to identify neighbours for future information exchange.

In the case of the path-decoupled paradigm, the messages are routed to nodes (NEs) which are not assumed to be on the data path, but which are aware of it. The initial effort on NSIS development has been focused on the path-coupled paradigm, and some issues are still open in respect to the path-decoupled paradigm.

Figure 2 provides a simple scenario of NSIS path-coupled configuration. The NSIS messages are transparent to nodes that are NSIS-unaware, such as node R2 in the figure. A single message from Sender to Receiver establishes a session. On the way, nodes R1 and R3 will establish an association between them as adjacent NEs.




Fig. 2. Simple path-coupled signaling

NSIS decomposes the overall signalling protocol suite into a generic (lower) layer and specific upper layers for each specific signalling application. These are: NTLP (NSIS Transport Layer Protocol) and NSLP (NSIS Signalling Layer Protocol).

In the lower layer, NTLP (also known as GIST [10], General Internet Signaling Transport) offers transport services to higher layer signalling applications for two purposes: sending and receiving signalling messages between neighbouring hops (NSIS entities), and exchanging control and feedback information, through the use of

GIST-query, *GIST-response* e *GIST-confirm* messages. NTLP messages can be delivered using existent transport protocols such as TCP, UDP, SCTP, HIP/IPsec, etc. By choosing the transport protocol to use, it is possible to guarantee security and reliability as needed.

Above the NTLP layer, there is the NSLP (NSIS Signalling Layer Protocol) layer, which generically stands for any protocol within the signalling application layer. Several NSLPs can be running independently of each other. This separation between transport and application allows NSIS to implement different QoS mechanisms, and extensibility even to non-QoS purpose. Although NSIS is used to QoS signalling in the first hand like in the QoS-NSLP specification [11], it should be possible to develop NSLPs for other signalling purposes that use different types of network control state, such as firewalls, NATs, and so on. NSLP uses GIST for message exchange.

4.2 NSIS in Ethernet

SBM, presented in Section 3, is RSVP-based. Although the NSIS protocol is quite similar to the RSVP protocol, namely in what respects support of unicast path-coupled signaling and soft state, NSIS has left out some complexities associated with RSVP like the multicast support. QoS NSLP also extends the set of reservation mechanisms like the support of sender or receiver-initiated reservations, bi-directional reservation and the support of reservations between arbitrary nodes. Thus, it makes sense to explore the use of NSIS for Ethernet resource allocation.

Exploring the usage of the NSIS over Ethernet makes even more sense if one considers that Ethernet is a widely spread technology, with more than 90% of today's data traffic originating from and terminating in it [12].

In the context of the EuQoS project, the authors propose the replacement of the SBM/RSVP solution by the NSIS approach for resource allocation in Ethernet networks. The basic approach is depicted in Figure 3, and is similar to the general NSIS approach.

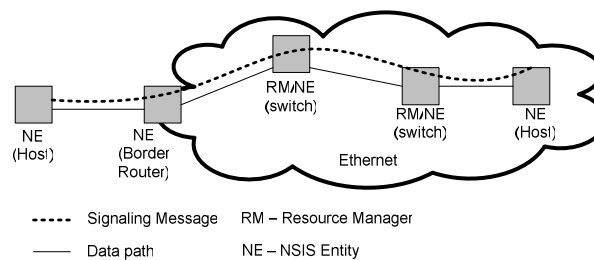


Fig. 3. Ethernet resource reservation through NSIS

The Resource Manager (RM) is the entity responsible for the admission control in the Ethernet domain, having similar functionality to the DSBM. This entity is NSIS-aware and includes a NSLP layer through which it is possible to receive signaling messages. These are used to install/maintain states along the Ethernet network.

Using this approach, it is possible to offer QoS guarantees even if the path between the sender and the receiver crosses an Ethernet network.

4.3 Validating NSIS in Ethernet

In order to validate the use of NSIS and its ability to support traffic differentiation in Ethernet, a preliminary simulation study was performed. Note that, at this stage, the objective of the study was not to compare NSIS with SBM/RSVP, but only to verify that NSIS is capable of supporting the required QoS functionality.

For the purpose of the study, the authors specifically developed several new modules for the NS-2 simulator, which allow the simulation of NSIS-based QoS provision over Ethernet networks. The modules, presented in [13], comply with the IEEE 802.1Q and IEEE 802.1p standards, implementing its traffic prioritization specifications.

Specifically, a queue object was developed that simulates an Ethernet switch/bridge and has the following properties:

- It can be configured to use up to eight virtual queues to group incoming LAN packets into separate traffic classes, according to IEEE 802.1p recommendations;
- The forwarding mechanism is implemented in such a way that packets from a given virtual queue are selected to transmit only if a higher order virtual queue is empty at the time of selection (strict priority queues).

To study the advantages of using priorities in Ethernet networks, the authors evaluated the transmission of heterogeneous signals in different topologies and with different traffic conditions. Next, one of these studies is presented, where a VoIP signal is transmitted over an overloaded Ethernet network, not using and using traffic priorities.

The topology of the referred scenario is depicted in figure 4. The link between the switch and Host 4 constitutes the bottleneck. The traffic from Host 1 to Host 4 is modeled as G.711 VoIP traffic without Voice Activity Detection (VAD) support. The other two traffic sources are modeled as exponential ON/OFF traffic.

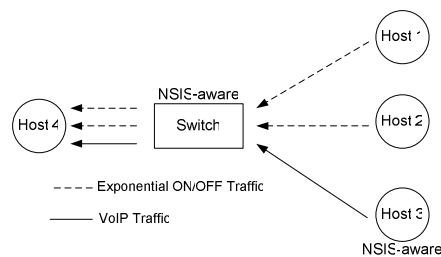


Fig. 4. Used LAN topology

A simplified NSLP layer implemented on the switch allows it to be signalled to process a flow with a certain priority (according to the priority levels listed on Table 1). Host 3 participates on the signalling process by sending to the switch a request

asking for a highest priority for the VoIP traffic flow. The exponential on/off traffic packets are kept with a constant low priority along all the simulation time.

Figure 5 shows the one-way-delay (OWD) experienced by each of the traffic flows when no prioritization is performed in the switch. As expected, all the traffic presents similar behaviour, experiencing high delay in the face of network congestion.

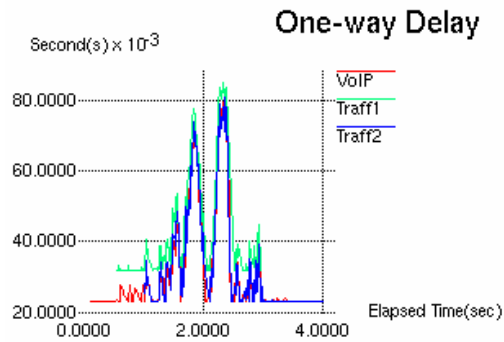


Fig. 5. Packet delay without traffic prioritization

Figure 6 presents the one-way-delay for the traffic flows, after the signaling message sent by Host 3 set the switch to prioritize the VoIP traffic in relation to the other two flows. The VoIP delay remains low during all of the simulation, as a result of the NSIS-based resource allocation.

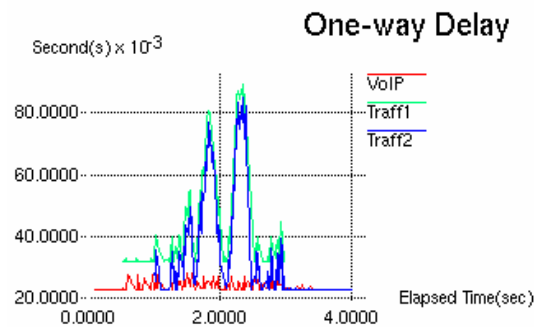


Fig. 6. Packet delay with traffic prioritization

4.4 NSIS-based Resource Allocation in the EuQoS Project

Considering the current lack of Ethernet devices supporting NSIS protocol, two solutions have been envisioned to allow its deployment in the EuQoS project. However, the proposed solutions require some modification on the LAN's layout.

4.4.1 Centralized Resource Allocation

This solution considers the use of just one RM element, a standard computer, which will perform the resource allocation for flows trying to get access to the LAN (Figure 7). The RM element needs to be fed with details of the network topology including link information and assumptions about the switch capabilities. All the standard traffic (which did not explicitly ask for resource reservation) is aggregated on a lower priority queue by the edge switches to avoid disturbance of the flows which resources were reserved by the NSIS signaling. In this way, the RM does not need to be aware of each traffic flow present on the LAN.




Fig. 7. Centralized resource allocation

4.4.2 Distributed Resource Allocation

In this case several RM elements are responsible for performing the resource reservation related tasks. As each RM manages a sub-set of resources (i.e. links), it is not necessary for them to know the entire network topology. As occurs on a standard IP NSIS signaling, it is necessary to establish a path inside the LAN by contacting the successive RM elements between the router and the destination node (Figure 8). Hosts inside the LAN can contact the closer RM to start the resource reservation process.

This solution is more flexible and allows a fine control over the traffic flows. By exploring the capabilities of some switches available on the market, it is possible to discover the RM elements between the sender and transmitter without a previous knowledge of the network topology. Naturally, this option demands more implementation effort.

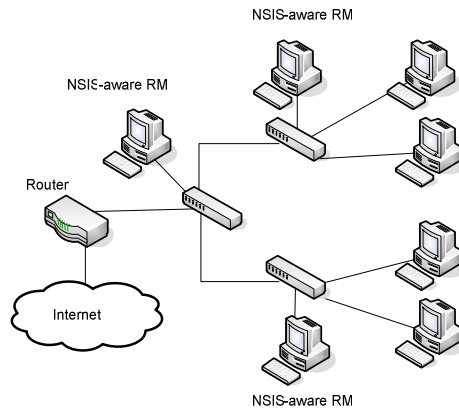


Fig. 8. Distributed resource allocation

5 Conclusion

As Ethernet becomes widely used in LAN, MAN and WAN scenarios, it is important to propose and study solutions for quality of service provision over this type of underlying network technology.

This paper proposed an approach to QoS provision in Ethernet, based on the emerging NSIS framework. In order to validate the proposal, simulation studies were performed which showed the potential of the proposed approach to differentiate traffic.

Following the promising simulation results, at the moment the authors are further analysing the overall issues related to the deployment of NSIS in Ethernet networks, and implementing the NSIS framework. Specifically, the GIST implementation is well under way and the NSLP protocol specification is being ultimated. Comparison of the SBM/RSVP-based solution with the NSIS solution will also be made in the near future, both by simulation and prototyping. Also, the analyses of the NSIS for QoS signaling purposes in heterogeneous scenarios involving Ethernet and WiFi is to be performed in future works.

References

1. Metro Ethernet Forum, <http://www.metroethernetforum.org>
2. 802.1Q, Standards for Local and metropolitan area networks; "Virtual Bridged Local Area Networks"; IEEE Computer Society 2003
3. 802.1D, Standards for Local and metropolitan area networks; "Media Access Control (MAC) Bridges"; IEEE Computer Society 2004

4. A. Simt, P. Langille, A. Rijhsinghani, K. McCloghrie; "Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions"; RFC 2674; 1999
5. Integrated Services over Specific Link Layers (issll), <http://www.ietf.org/html.charters/issll-charter.html>; 2001
6. R. Yavatkar, D. Hoffman, Y. Bernet, and F. Baker; "SBM Subnet Bandwidth Manager): Protocol for RSVP-based Admission Control Over IEEE 802-style Networks"; RFC 2814; May 2000
7. M. Seaman, A. Smith, and E. Crawley; "Integrated Service Mappings on IEEE 802 Networks"; RFC 2815; May 2000
8. S. Shenker J. Wroclawski; "General Characterization Parameters for Integrated Service Network Elements"; RFC 2215; September 1997
9. R. Hancock, G. Karagiannis, J. Loughney and S. van den Bosch; "Next Steps in Signaling (NSIS): Framework"; RFC 4080; June 2005
10. H. Schulzrinne, R. Hancock, "GIST: General Internet Signaling Transport draft-ietf-nsis-ntlp-09"; Internet Draft; February 2006
11. J. Manner, G. Karagiannis, A. McDonald, "NSLP for Quality-of-Service signalling draft-ietf-nsis-qos-nslp-09.txt"; Internet Draft; January 2006
12. M. McGarry, M. Maier and M. Reisslein, "Ethernet PONs: A Survey of Dynamic Bandwidth Allocation (DBA) Algorithms"; IEEE Communications Magazine, vol. 42, no. 8, pp. S8-S15, Aug. 2004
13. M. Carmo, et al; "Ethernet QoS Modelling in Emerging Scenarios"; Proceedings of 3rd International Workshop on Internet Performance, Simulation, Monitoring and Measurement, IPS-MoMe 2005, pp.90-96, IST MoMe Cluster