

This document is published in:

Trans. Emerging Tel. Tech. 23 (2012) 3, pp. 227-233

DOI:10.1002/ett.1526

© 2012. John Wiley & Sons, Ltd.

Study of the potential energy savings in Ethernet by combining Energy Efficient Ethernet and Adaptive Link Rate

P. Reviriego^{1*}, J. A. Maestro¹, J. A. Hernández² and D. Larrabeiti²

¹ Dept. Ing. Informática, Universidad Antonio de Nebrija, C. Pirineos 55, 28040 Madrid, Spain

² Dept. Ing. Telemática, Universidad Carlos III de Madrid, Avda. Universidad 30, 28914 Leganés, Madrid, Spain

Abstract: The Energy Efficient Ethernet (EEE) standard (IEEE 802.3az) has introduced some modifications to existing Ethernet standards with the aim of improving energy efficiency in the transmission of data frames. The adopted mechanism is to use different power modes of operation such that a device can be put into a low-power mode when there is no data to transmit and awake it very quickly upon frame arrival. This is expected to reduce energy consumption significantly, especially on low-loaded links. An alternative to EEE, known as Adaptive Link Rate (ALR), was also proposed to improve energy efficiency in Ethernet by dynamically changing the link speed at different traffic loads. ALR reduces the energy consumption because the lower the link speed, the less energy is consumed. However, this option is no longer considered in the EEE standard process.

In this work, the potential energy savings of combining both EEE and ALR are explored. The results show that such a combination can further improve the already efficient EEE standard given the different power overheads associated with transitions between the active and idle power modes at different link speeds. This conclusion provides a clear motivation to investigate the combination of EEE and ALR more in detail in further developments.

*

P. Reviriego, Dept. Ing. Informática, Universidad Antonio de Nebrija, C. Pirineos 55, 28040 Madrid, Spain.

E-mail: previrie@nebrija.es

1. INTRODUCTION

In the past few years, there is a growing concern with the consumption of energy worldwide for various reasons. On the one hand, making people and devices energy-aware can save a lot of money and reduce the oil-dependency of countries and governments. On the other hand, sensible habits in energy consumption can also reduce the emissions of greenhouse gases to the atmosphere. For instance, the average cost of energy in Europe per 100 kWh is of about €15 for household consumers and of about €9 for industrial end-users in 2007. The Internet core is estimated to consume about 6 TWh per year, which costs about €90–150m according to the previous estimates [1, 2].

Ethernet is a good example of technology that can be improved in terms of energy efficiency because Ethernet devices operate at full power even when no data is being transmitted. The IEEE 802.3az standard (Energy Efficient Ethernet, EEE) defines new mechanisms to reduce energy

consumption in Ethernet, with potential energy savings of 3 TWh [3]. The main idea behind EEE is to define a sleep mode to be switched to when there is no data for transmission. Such a low-power mode is expected to consume much less power than the active or transmission mode, in which the Network Interface Card (NIC) would stay only during frame transmission. This aims to make power consumption proportional to traffic load, unlike legacy Ethernet that operates at full power all the time with subsequent energy waste.

In addition to this, the Adaptive Link Rate (ALR) mechanism has also been proposed in the literature as a potential source of energy savings for NICs supporting different link speeds [3]. Essentially, ALR proposes to use the most efficient speed in terms of energy consumption for a particular traffic load observed on a link. This also implies to adjust the link speed to the most appropriate one (in terms of energy efficiency) upon traffic changes. Substantial energy reductions may be achieved because

different speeds have different energy consumption patterns. Basically, high-speed links consume much more than lower speed links. This option was however no longer considered in the EEE standard process.

This article considers the potential energy savings achievable by combining both EEE with ALR, with respect to EEE only as the IEEE 802.3az proposes. As shown in the results, both EEE and ALR can provide extra energy savings when used together. Essentially, the energy overheads caused by EEE can be alleviated if combined with ALR, as shown next.

The rest of this work is organised as follows: Section 2 overviews the fundamentals of the two approaches proposed to improve energy efficiency in Ethernet: EEE and ALR. Section 3 shows an analysis of the potential benefits of combining both approaches with simulated Poisson traffic. Section 4 shows a measurement-based practical study to show the benefits of ALR and EEE in a real case scenario. Finally, Section 5 concludes this work with a summary of its main findings.

2. ENERGY EFFICIENT ETHERNET AND ADAPTIVE LINK RATE

2.1. Energy Efficient Ethernet review

The EEE standard has defined two power modes for the operation of every EEE NIC: The idle or sleep (low-power) and the active modes. The idea of the EEE standard is to put the physical layer into the low-power mode when there is no data to transmit and wake it up only when data is ready to be transmitted. This two-state operation of NICs should result in large potential energy savings because links are normally lightly loaded, in the range of 1% to 5% [4]. Indeed, such sleep/active operation requires minor updates to the receiver elements because the channel is often quite stable. In the low-power mode, the receiver freezes the system elements, but it is capable of waking up within just a few microseconds upon data arrival. Such a wake-up time provides an upper bound for the worst-case delay that frames may suffer as a result of a mode change.

Figure 1 shows a state transition example of a given link as following the IEEE 802.3az standard [5]. In the figure, the transceiver requires T_s units of time to enter the low-power mode (*sleep time*), spends T_q units of time in the idle state (*quiet time*) and then exits the low-power mode within T_w units of time (*wake-up time*). Finally, the IEEE 802.3az standard also considers the scheduling of periodical short periods of activity T_r (*refresh time*) to ensure that the receiver elements are always aligned with the channel conditions.

Concerning energy efficiency, there is significant energy consumption during T_w , T_s , T_r and of course during frame transmission, but only a small portion of it (typically about 10% of the active power consumption [6–8]) during T_q . Clearly, large energy savings are achieved as long as the NIC spends a long time in T_q (light loads).

Table I summarises the minimum values for T_w and T_s as specified in the IEEE 802.3az standard for 100Base-TX, 1000Base-T and 10GBase-T.

2.2. Adaptive Link Rate review

During the standardisation process of EEE, another alternative was proposed to reduce energy consumption: Downgrading the link speed when there is little traffic activity [3, 9] because the lower the link speed, the less energy is consumed. Essentially, most Ethernet devices support multiple speeds; for example, current computers are provided with 10/100/1000 Mbps NICs. In the autonegotiation process during link setup [10], the highest link speed that both ends support is typically selected, even if a lower speed is sufficient for satisfying the traffic demands.

Table II provides some typical energy consumption values for physical layer devices (PHYs) at different link speeds. The values selected are in line with those reported by different manufacturers during the standardisation process [6, 11, 12]. These figures will obviously vary as new technologies are used to manufacture the integrated circuits. However, the relative consumptions of one speed versus another will tend to remain more stable as the technology scaling has similar effects on the different PHYs. Therefore, the absolute figures provided in this paper will approximately capture the current status, whereas the relative ones would also approximately apply to newer generation of devices. From the data in Table II, it becomes clear that it makes sense to use a lower speed when there is little traffic to be transmitted.

The ALR mechanism proposes an algorithm to dynamically adjust the link speed to the traffic pattern [3]. However, in ALR, speed changes occur through autonegotiation, which takes hundreds of milliseconds to complete the process. During this time, the link is down and communications are disrupted. For such unacceptable timings to be reduced, the authors in [13] have proposed the Rapid PHY Selection mechanism, whereby a frame exchange is used to re-negotiate a change of speed without the need

Table I. Proposed wake up and sleep timers for different speeds.

Protocol	Minimum T_w (μs)	Minimum T_s (μs)
100Base-TX	30.5	200
1000Base-T	16.5	182
10GBase-T	4.48	2.88

Table II. Typical physical layer device power consumption values for different link speeds.

Link speed	Average consumption
100 Mbps	200 mW
1 Gbps	600 mW
10 Gbps	4 W




Figure 1. Time required to transition between the active and low-power modes in Energy Efficient Ethernet.

for restarting the autonegotiation process. Although Rapid PHY Selection makes speed changes much faster, still these require the adjustment of a large number of elements in the receivers, including equalisers, echo cancellers, timing circuits, and so on, to the new speed, a non-negligible amount of time over which the link is down.

3. PERFORMANCE EVALUATION OF ENERGY EFFICIENT ETHERNET WITH ADAPTIVE LINK RATE

This section aims to study the benefits of combining EEE with ALR in terms of energy savings. The study comprises 100Base-TX, 1000Base-T and 10GBase-T with the power consumption values summarised in Table III.

Essentially, there are different consumption ratios per Mbps at different link speeds, thus motivating the use of ALR. Moreover, from Table III, it looks more convenient to use the highest speed rate whenever this is possible because the power cost per Mbps is lower for the higher speeds when the link is heavily loaded. However, this is not completely true as noted next.

3.1. A first estimate of power consumption: ideal behaviour

Now, let ρ give the average load of a given link. Clearly

$$\rho = \frac{T_{\text{active}}}{T_{\text{idle}} + T_{\text{active}}} \quad (1)$$

Here, ρ is estimated from the amount of time that the NIC spends in the active mode over the total time. Now, the average power consumption is defined as [9]

$$P_{\text{avg}} = P_{\text{idle}}T_{\text{idle}} + P_{\text{active}}T_{\text{active}} \quad (2)$$

Table III. Baseline power consumption values.

Link speed	Active state	Low-power state (10%; mW)	Power/Mbps at full load (mW/Mbps)
100Base-TX	200 mW	20	2
1000Base-T	600 mW	60	0.6
10GBase-T	4 W	400	0.4

where P_{idle} and P_{active} are given in Table III for the different link speeds. With the utilisation of ρ from Equation (1), Equation (2) becomes

$$P_{\text{avg}} = (1-\rho)P_{\text{idle}} + \rho P_{\text{active}} = \left(\frac{1-\rho}{10} + \rho\right) P_{\text{active}} \quad (3)$$

According to this, a traffic demand of 10 Mbps on a 1 Gbps PHY (this is 1% of load, $\rho = 0.01$) consumes 65.4 mW of power on average. The same 10 Mbps traffic demand on a 100 Mbps PHY (this is $\rho = 0.1$) consumes 38 mW. However, if the traffic demand is of 80 Mbps, this consumes 103.2 mW on a 1 Gbps PHY and 164 mW on a 100 Mbps PHY.

This simple power consumption comparison is shown in Figure 2 for all traffic loads between 1 Mbps (10^6 bps) and 10 Gbps (10^{10} bps) in a log x -axis. Obviously, the comparison between the three link speeds is only possible for traffic loads up to 100 Mbps. The same occurs between 1000Base-T and 10GBase-T whereby the power comparison is only possible for traffic loads up to 1000 Mbps. Following this ideal model, depending on the input traffic demand, the most suitable link speed varies, that is 100Base-TX for traffic loads up to 30 Mbps, 1000Base-T for traffic loads up to 1 Gbps and 10GBase-T for traffic loads beyond 1 Gbps.

So, from this ideal model, it is not always true that low-speed links consume less power than a high-speed link for a certain traffic demand. Next section repeats this figure but considering the energy overheads caused by the timers T_w , T_s and T_r of EEE.

3.2. A second estimate of power consumption: considering T_w and T_s

In the previous section, Equation (3) gives an average power consumption strictly proportional to the traffic load. However, as noted in Section 2, the amount of time spent in putting the link to sleep and waking it up consumes significant power close to P_{active} , but it was accounted as P_{idle} in Equation (2). Indeed, the amount of time (thus power) spent in waking up and putting into sleep the link may constitute a substantial portion of the total time. As an example, consider a 1500-byte data frame to be transmitted over a 100 Mbps link (100Base-TX). This frame requires $T_{\text{frame}} = 120 \mu\text{s}$ of active link operation. However, in EEE, the transmission of this frame actually

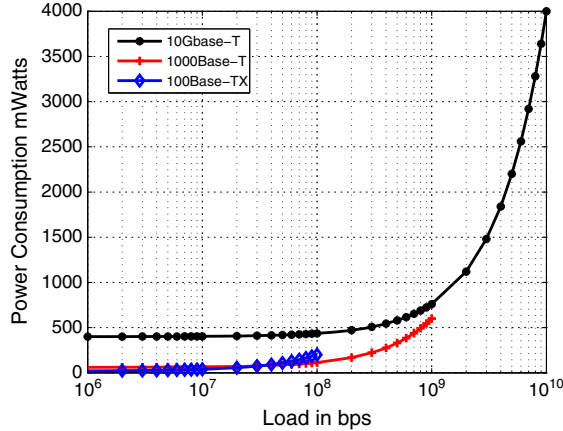


Figure 2. Energy Efficient Ethernet with Adaptive Link Rate following the ideal model of Equation (3).

requires $T_w = 30.5 \mu\text{s}$ extra to wake up the link plus further $T_s = 200 \mu\text{s}$ to send it back to sleep. In total, this accounts for $T_w + T_{\text{frame}} + T_s = 350.5 \mu\text{s}$ of link activity for a data frame transmission time of $120 \mu\text{s}$. This means that only 34.2% of the time (and power) is spent on actual data transmission. Thus, EEE introduces a power overhead of 65.8%, which is even larger for smaller data frames. Table IV summarises such power overheads for different link speeds, assuming two typical frame sizes: 1500 and 150 bytes.

Clearly, the previous power consumption model given by Equation (3) is no longer valid. For the power overhead impact on EEE to be analysed, Figure 3 shows a simulation-based power consumption study of 100Base-TX, 1000Base-T and 10GBase-T assuming data frames of 150 (Figure 3 left) and 1500 bytes (Figure 3 right). In the simulation, fixed-size data frames are assumed to arrive at the NIC following a Poisson process at different traffic loads. Under the assumption that the link goes to sleep whenever there are no data frames to be transmitted and awoken as soon as a new data frames arrives for transmission, the average values for T_{idle} and T_{active} are obtained, and the power efficiency is computed following Equation (2).

Table IV. Power overheads of Energy Efficient Ethernet for different link speeds and frame sizes.

Technology	100Base-TX	1000Base-T	10GBase-T
T_w (minimum; μs)	30.5	16.5	4.48
T_s (minimum; μs)	200	182	2.88
T_{frame} (1500 bytes; μs)	120	12	1.2
Transmitter efficiency (1500 bytes; %)	34.2	5.7	14
T_{frame} (150 bytes; μs)	12	1.2	0.12
Transmitter efficiency (150 bytes; %)	4.9	0.6	1.6

As shown, the power efficiency deviates quite significantly from the ideal load-proportional behaviour given by Equation (3) (shown in Figure 2), specially at low loads and high-speed rates. For small data frames, where the power overhead is significantly more important than for large data frames (see Table IV), this behaviour is even worse. Thus, the current timer values for T_w and T_s specified in the standard make the NICs waste most of its energy in waking up and going to sleep. Nevertheless, the power consumption is still far from the values of legacy Ethernet standards (100% power all the time).

Finally, Figure 3 goes a bit further showing, for instance, the substantial energy waste of using 1000Base-T for traffic loads below 100 Mbps with respect to 100Base-TX. This is a consequence of two issues. First, 100Base-TX has a lower power overhead than 1000Base-T according to Table IV. Secondly, in 100Base-TX, the NIC operates at higher utilisation values, and it is more likely that the overhead associated to the wake/sleep process is shared among multiple frames, thus minimising the power overheads of EEE. Clearly, at low loads, more transitions to the sleep and active states are performed per frame, thus reducing the energy efficiency. A similar behaviour is observed when comparing 1000Base-T and 10GBase-T at low loads, despite the fact that the power overheads of EEE are higher in the case of 1000Base-T.

The conclusion to this is that of course EEE saves much energy with respect to legacy Ethernet standards, but it can be further improved if used in combination with ALR. Furthermore, Figure 3 suggests to use the lower link speed possible for a given traffic demand from the point of view of energy consumption because it decreases the number of active/sleep transitions. However, it is well known that frame arrivals in Ethernet deviate from a Poisson model [14]. Therefore, the results provided in this section can only be taken as an initial approximation. In the next section, traces from real measurements will be used to validate the conclusions in a number of real scenarios.




Figure 3. Energy Efficient Ethernet with Adaptive Link Rate for Poisson traffic for (a) 150-byte data frames and (b) 1500-byte data frames.

4. A PRACTICAL CASE STUDY

This section aims to demonstrate our previous conclusions concerning EEE efficiency and ALR in the real scenario shown in Figure 4.

In this scenario, two hosts are directly attached to two 100/1000 Mbps switches, which are further connected over a 100 Mbps link. In the first scenario, the two edge links are configured to operate at 1000 Mbps, thus acting the interswitch link as a 100 Mbps bottleneck. In the second case, all three links are configured to operate at 100 Mbps. The idea is to obtain the power consumption values in the two cases and show whether or not there is any power saving improvement if the speed of the two edge links are downgraded to 100 Mbps, of course with no significant performance degradation. It is also worth remarking that this experiment may have similarities with many corporate networks, whereby their local area network operates at 100/1000 Mbps but are further limited to a smaller rate by the metropolitan area network access network.

In the two cases, a large file is transferred from one host to the other, and the frame arrival times on each communication direction are recorded. Such frame arrival times are then used to estimate the power consumption values as if EEE was employed on the links, as explained in the previous section. The objective is to derive the potential energy savings if the two ends were capable of downgrading to a lower speed values, as ALR does. Table V confirms that such a potential energy saving indeed exists.




Figure 4. Experimental setup.

As shown, there is a great energy reduction if the edge links's speeds are downgraded to 100 Mbps, especially in the downstream direction where the ACKs are transmitted, because the power overhead of the wake/sleep process is larger for smaller data frames and low utilisation values.

Furthermore, it is worth noting that these simulation-based results are close to the Poisson-based conclusions derived in the previous section. When looking at Figure 3 (right), a 71 Mbps load is expected to save around 290 mW when combining EEE and ALR, a result that is closed to the simulation-based 298 mW. In the downstream direction (ACKs), Figure 3 (left) estimates savings of around 210 mW for 150-byte packets (slightly smaller than the real ones as the ACKs are shorter than the 150-byte packets used in the simulation). This agrees with the results presented in [15] where the authors analyse the energy consumption of EEE in a number of measurement-based real scenarios and showed that the Poisson-based simulation provided a reasonable approximation to the actual energy consumption values.

In conclusion, there is a great potential energy savings when combining EEE with ALR, especially whenever there is a low-speed intermediate link that throttles the end-to-end load to below 10% on some links. In addition to this, the Poisson-based experiments accurately match the energy savings when combining EEE with ALR for the scenarios considered.

Finally, concerning the excessive power overhead of wake/sleep on small data frames like ACKs, a good

Table V. Energy consumption estimates for the file transfer experiment.

Direction	Average frame size (bytes)	Speed	T_{active} (%)	Consumption (mW)	Load (%)
File upload	1499	1 Gbps	75.96	455.76	7.11
File upload	1499	100 Mbps	78.68	157.36	71.13
ACKs direction	78	1 Gbps	58.35	350.10	0.14
ACKs direction	78	100 Mbps	44.92	89.84	1.39

Table VI. Energy consumption estimates for the data centre experiment.

Direction	Average frame size (bytes)	Speed	T_{active} (%)	Consumption (mW)	Load (%)
Input	87	1 Gbps	65.90	395.4	1.22
Input	87	100 Mbps	61.45	122.9	12.16
Output	1497	1 Gbps	72.92	437.5	52.21

strategy may be to control link speed independently on the two directions. Thus, whenever there is an asymmetrically loaded link, the two directions of communications may run at different speeds on attempts to squeeze the maximum energy savings possible. For this potential energy savings to be illustrated, a traffic trace from a data centre is used to estimate the consumption values of EEE at different link speeds. The results obtained, shown in Table VI, advice to use different link speeds on each direction of communications depending on the particular traffic load on each of them.

5. SUMMARY AND DISCUSSION

This work analyses the potential benefits of using the low-power modes defined in EEE in combination with changing link speeds as proposed in ALR. Using the two techniques in combination may bring substantial additional energy savings with respect to using EEE only. The reason for this is two-fold: On the one hand, low-link rates consume less energy than high-speed rates; on the other hand, the power overheads caused by waking up and sending to sleep the links vary significantly at different link speeds. Thus, the appropriate choice of link speed (driven by ALR) may provide further energy savings if done appropriately.

This however brings important implementation issues that need to be considered in further designs of EEE with ALR. First, energy savings vary significantly with frame sizes; thus, ALR must monitor both link load and average frame size to decide the schedule of link speed changes appropriately. Secondly, downgrading 10 times the link speed of course reduces energy consumption but increases transmission delay an order of magnitude. Finally, communications are disrupted during a speed change that may cause network performance degradation on highly loaded links.

ACKNOWLEDGEMENTS

The authors would like to thank G. Chesson from Google for the anonymous trace from their data centres used in the experiments in Section 4.

The authors would also like to acknowledge the Spanish project T2C2 (under code TIN2008-06739-C04-01) and the Medianet research grant (under code S2009/TIC-1468) for the development of this work.

REFERENCES

- Gupta M, Singh S. Greening of the Internet, In *Proceedings of the 2003 Conference on Applications Technologies Architectures and Protocols for Computer Communications SIGCOMM*, 2003; 19–26.
- Vereecken W, Deboosere L, Simoens P, Vermeulen B, Colle D, Devellder C, Pickavet M, Dhoedt B, Demeester P. Power efficiency of thin clients. *European Transactions on Telecommunications* 2010; **21**: 479–490.
- Gunaratne C, Christensen K, Nordman B, Suen S. Reducing the energy consumption in Ethernet with Adaptive Link Rate (ALR). *IEEE Transactions on Computers* 2008; **57**(4): 448–461.
- Bennett M. Server bandwidth utilization plots, 2007. Presented at IEEE 802.3 Meeting.
- IEEE Std 802.3az Energy Efficient Ethernet-2010, 2010.
- Chou J, Kuo A, Tseng D, Lin J, Huang K. Proposal of low-power idle: 100Base-TX, 2008. Presented at IEEE 802.3 meeting.
- Taich D, Tellado J, Zimmerman G, Barkan O. 10GBASE-T low-power idle proposal, 2008. Presented at IEEE 802.3az meeting.
- Grimwood M, Lin X, Wang P. Energy efficient Ethernet 1000BASE-T LPI timing parameters, 2008. Presented at IEEE 802.3az meeting.
- Nevedschi S, Popa L, Iannaccone G, Ratnasamy S, Wetherall D. Reducing network energy consumption via sleeping and rate adaptation, In *Usenix Symposium on Networked System Design and Implementation*, 2008.
- Spurgeon CE. *Ethernet: the definitive guide*. O'Really, 2000.
- Dove D. Energy Efficient Ethernet switching perspective, 2008. Presented at IEEE 802.3 meeting.

12. Kohl B. 10GBase-T power budget summary, 2007. IEEE 802.3az Task Force Presentation.
13. Blanquicet F, Christensen K. An initial performance evaluation of Rapid PHY Selection for Energy Efficient Ethernet, In *Proceedings of IEEE conference on Local Computer Networks*, 2008; 223–225.
14. Leland WE, Taqqu MS, Willinger W, Wilson DV. On the self-similar nature of Ethernet traffic. *IEEE/ACM Transactions on Networking* 1994; **2**(1): 1–15.
15. Reviriego P, Hernández JA, Larrabeiti D, Maestro JA. Performance evaluation of Energy Efficient Ethernet. *IEEE Communications Letters* 2009; **13**(9): 697–699.

AUTHORS' BIOGRAPHIES

Pedro Reviriego received the MSc and PhD degrees (Honours) from Universidad Politécnica de Madrid in 1994 and 1997, respectively, both in Telecommunications Engineering. From 1997 to 2000, he was an R&D engineer at Teldat, working on router implementation. In 2000, he joined Massana where he worked on the development of 1000BaseT transceivers until 2003 when he became a visiting professor at University Carlos III de Madrid. During 2004–2007, he is a distinguished member of technical staff at LSI Corporation, where he was involved in the development of Ethernet transceivers. He is currently at Universidad Antonio de Nebrija, Madrid. He has authored a large number of papers in international conferences and journals and has also participated in the IEEE 802.3 standardisation for 10GBASE-T. His research interests include fault-tolerant systems, performance evaluation of communication networks and the design of physical layer communication devices.

Juan Antonio Maestro received the MSc degree in Physics and the PhD degree in Computer Science from the Universidad Complutense de Madrid, Spain in 1994 and 1999, respectively. Since then, he has worked as a lecturer at several universities in Madrid, namely Universidad Complutense de Madrid, Universidad Nacional de Educación a Distancia, Saint Louis University at Madrid

and Universidad Antonio de Nebrija, where he currently leads the Computer Architecture and Technology Group. His research activities are focused on the space field, where he works on several research projects involved in reliability and radiation protection of electronic equipment. He also collaborates with the European Space Agency. He has authored a large number of technical publications in the areas of high-level synthesis and co-synthesis, signal processing and real-time systems, fault tolerance and reliability. Finally, he has also worked as Project Management Professional for several multinational companies.

José Alberto Hernández completed the five-year degree in Telecommunications Engineering at Universidad Carlos III de Madrid (Madrid, Spain) in 2002 and the PhD degree in Computer Science at Loughborough University (Leics, UK) in 2005. From 2005 to 2009, he was a post-doctoral research and teaching assistant at Universidad Autónoma de Madrid, where he participated in a number of both national and European research projects concerning the modelling and performance evaluation of communication networks and particularly the optical burst switching technology. In 2009, he moved to Universidad Carlos III de Madrid, where he became an associate professor. He has published more than 40 articles in both journals and conference in-proceedings. His research interests include the areas at which mathematical modelling and computer networks overlap.

David Larrabeiti is currently a full professor at Universidad Carlos III de Madrid (Spain) since 1998, where he teaches several modules concerning high-speed switching networks and architectures. He received the MSc and PhD degrees in Telecommunications Engineering from Universidad Politécnica de Madrid in 1991 and 1996, respectively. He has participated in a large number of both national and international research projects focused on next-generation networks for more than a decade. His research interests include the design of the future Internet infrastructure, ultra-broadband multimedia transport and traffic engineering over Internet Protocol/Generalised Multi-Protocol Label Switching backbones.