An Energy Saving Scheme for Multilane-Based High-Speed Ethernet

Kyeong-Eun Han, Choong-Reol Yang, Kwangjoon Kim, Sun-Me Kim, and Jonghyun Lee

In this paper, we propose a scheme for partially dynamic lane control for energy saving in multilane-based high-speed Ethernet. In this scheme, among the given transmission lanes, at least one lane is always operating, and the remaining lanes are dynamically activated to alleviate the network performance in terms of queuing delay and packet loss in the range of acceptance. The number of active lanes is determined by the decision algorithm based on the information regarding traffic and queue status. The reconciliation sublayer adjusts the transmission lane with the updated number of lanes received from the algorithm, which guarantees no processing delay in the media access control layer, no overhead, and minimal delay of the exchanging control frames. The proposed scheme is simulated in terms of queuing delay, packet loss rate, lane changes, and energy saving using an OPNET simulator. Our results indicate that energy savings of around 55% (or, when the offered load is less than 0.25, a significant additional savings of up to 75%) can be obtained with a queuing delay of less than 1 ms, a packet loss of less than 10^{-4}, and a control packet exchange time of less than 0.5 μs in random traffic.

Keywords: Dynamic lane control, energy efficiency, multilane, high-speed Ethernet, 40-G/100-G Ethernet.

Manuscript received May 22, 2012; accepted July 9, 2012.

This work was supported by the IT R&D program of MKE/KEIT [10041414, Terabit Optical-Circuit-Packet Converged Switching System Technology Development for Next-Generation Optical Transport Network].

Kyeong-Eun Han (phone: +82 42 860 5758, kehan@etri.re.kr), Choong-Reol Yang (cryang@etri.re.kr), Kwangjoon Kim (kjk@etri.re.kr), Sun-Me Kim (kimsunme@etri.re.kr), and Jonghyun Lee (jlee@etri.re.kr) are with the Advanced Communications Research Laboratory, ETRI, Daejeon, Rep. of Korea.

http://dx.doi.org/10.4218/etrij.12.1812.0110

I. Introduction

The recent exponential growth of traffic and the newly emerging multimedia services, such as video applications, interactive games, and social network services, require much higher data rates, increasing complexity, and a greater number of network servers. This has caused an enormous amount of electricity usage, and, together with very high capacity, energy efficiency has become a necessity for next-generation networks [1]-[6].

Many previous studies regarding power saving have mainly focused on hardware devices or network elements themselves. However, the power consumption is still high because hardware devices and network elements use the same power for data transmission regardless of their traffic load. The actual link utilization is below 5% and 30% at an edge link and backbone, respectively [7]-[10]. Therefore, data rate adaptation depending on link utilization is another method that can provide more energy efficient networks.

The IEEE 802.3az Energy Efficient Ethernet Task Force standardized an energy efficient Ethernet (EEE) approach targeting a 100-Mbps to 10-Gbps Ethernet link using 100GBASE-TX, 1000BASE-T, 10GBASE-T, 10GBASE-KR, 10GBASE-RX4, and 1000BASE-KX physical layers (PHYs). The EEE draft is optional to the existing Ethernet standards involved with the above items and provides sleep mode for energy conservation. Sleep mode provides an active state and low-power idle (LPI) state, where the data is transmitted at the full rate in the active state and the LPI state is used when there is no data to be sent. During an LPI state, the link transmits LPI packets during an inactive period and periodically transmits short refresh signals to keep the link alive and to align the receivers with the current link status. An active state is entered again when data to be sent is received [10], [11].

Meanwhile, to provide a much higher speed, the IEEE 802.3ba Energy Efficient Ethernet Task Force standardized multiple-lane-based 40-Gbps and 100-Gbps Ethernet for local server applications and the Internet backbone, respectively [12].

However, the issue of energy efficiency was not included, even though it is necessary for 40-G or 100-G Ethernet when we consider that a transmission at a higher speed consumes much more network power. In an optical 40-G or 100-G Ethernet link with multiple low-rate lanes, which is the EEE approach and provides a fast switching time and a single lane, a link is not appropriated. This may give rise to serious problems, such as a high packet loss and queuing delay, because optical devices have long on and off switching times compared to those of electrical devices. Therefore, a novel energy-efficient method fitting 40-G and 100-G Ethernet is required.

Dynamic lane control may be a solution for energy efficiency in a high-capacity optical Ethernet link based on multiple lanes [13], [14]. It dynamically activates certain lanes for transmission depending on the link utilization. For dynamic lane control, the following three requirements are needed: first, a lane decision algorithm to determine the number of required transmitting lanes in accordance with the state of the network; second, the mechanism to rearrange the number of lanes between two nodes; and, finally, a data rate control technique to transmit data at a changed rate, adaptively.

In this paper, we propose a dynamic lane control scheme for energy saving in 40-G/100-G Ethernet. This is based on the method for partially dynamic lane control for data transmission, and some parameters are defined for measuring the average link utilization and reflecting the current queue status. The lane control manager (LCM) determines the number of active lanes with this information. The reconciliation sublayer (RS) adjusts the transmitting lanes using this information by exchanging the control messages during the interframe time. Our scheme provides minimal delay, a small packet loss, and significant energy savings.

The remainder of this paper is organized as follows. In section II, an overview of 40-G/100-G Ethernet defined by IEEE 802.3ba and dynamic lane activation is given. In section III, the proposed lane control algorithm, protocol, and data rate adaptation mechanisms are described in detail. In section IV, we present and evaluate its performance in terms of delay, packet loss, energy saving rate, and average number of used and changed lanes under an OPNET simulation. Finally, in section V, we offer some concluding remarks regarding our proposal.

II. Overview

1. 40-G/100-G Ethernet Architecture

The salient feature of 40-G/100-G Ethernet is the use of a coarse wavelength division multiplexing approach. A single data stream is simultaneously distributed and transmitted through all lanes. Figure 1 shows the 40-G/100-G Ethernet architecture and data transmission defined by the IEEE 802.3ba standard [12]. The RS adapts the bit serial data of the media access control (MAC) to the parallel serial encoding data of the PHY. The physical coding sublayer (PCS) is specified as common for a 40-G/100-G PHY (known as 40GBASE-R and 100GBASE-R) and performs 64b/66b encoding, scrambling, and multilane distribution (MLD) on the transmitting side. The 40-Gb media independent interface (XLGMII) and 100-Gb MII (CGMII) are used as the service interface between the RS and PCS. These interfaces are optional and are used as a basis for the specifications in the 802.3ba standard. A physical medium attachment (PMA) primarily adapts the PCS lane formatted signal to an appropriate number of abstract or physical lanes. It may be used once or multiple times to adapt the number and rate of the PCS lanes to the number and rate of the physical medium dependent (PMD) sublayer lanes.

The 64-bit serial data transmitted from the RS is encoded by adding a sync bit that indicates the data (sync bit = 01) or control (sync bit = 10) block and is scrambled with 66-bit blocks and distributed to m virtual PCS lanes in a round-robin distribution at one time. This allows the PCS to support multiple physical lanes in the PMD sublayer. The 66-bit blocks transmitted over m PCS lanes are distributed to n electrical PMA lanes after bit multiplexing and finally transmitted to a receiver over n optical PMD sublayer lanes of the optical link. Here, the number of PMA and PMD sublayer lanes is less than or equal to the number of PCS lanes ($n \leq m$), as the 802.3ba standard specifies that the number of PMA and PMD sublayer lanes should be a divisor of the number of PCS lanes. In the 802.3ba standard, the number of PCS and PMA lanes

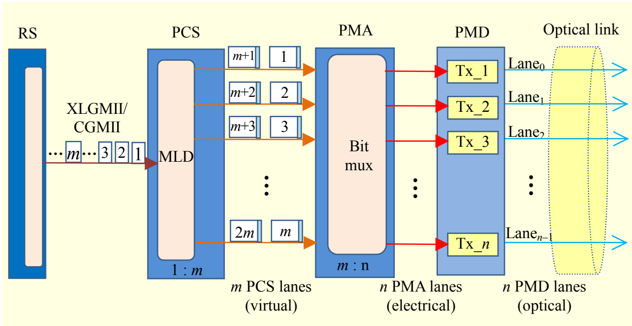

Fig. 1. Multilane-based 40-G/100-G Ethernet.

corresponds to {m=4, n=4} for 40-G Ethernet and {m=20, n=4} and {m=20, n=10} for four-lane and ten-lane 100-G Ethernet, respectively. All lanes are always used for data transmission, regardless of the link utilization.

2. Dynamic Multilane Activation

To reduce the energy consumption, we consider a dynamic lane control technique that cuts down on or expands the number of lanes according to a decrease or increase in traffic, respectively. As mentioned earlier, this is more efficient than EEE for high-speed Ethernet with multiple lanes and a relatively long switching time because it may deal independently with multiple lanes by turning on or off each lane without frequent control signaling for alignment and synchronization with the receiver. In the 802.3ba standard, all lanes of a link are always used for a transmission, and the PHY carries out such functions as synchronization and alignment after MLD based on entire lanes. Therefore, the PHY and PMD devices should support additional functions to provide dynamic lane control for energy efficiency. This issue is not discussed in this paper, however. We assume that the flexibility of the PHY and PMD devices is available at the optical PHY.

In the dynamic lane control method, the data rate is adaptively defined by how many lanes are used in a link. This method has a trade-off between the energy efficiency and network performance, such as delay and loss rate. The network has a better performance when using full lanes than when using only portions of the lanes for transmitting, but it is more energy efficient to use only portions of the lanes.

With this in mind, we first define the dynamic lane control method according to the range of lanes dealt with dynamically, which we call "fully dynamic lane control" and "partially dynamic lane control." In the former, all lanes (n lanes) are used dynamically, and they may all be inactive when no data is being sent. In the latter, however, some of the lanes (i lanes) are used statically, while the rest (n–i lanes) are used dynamically.

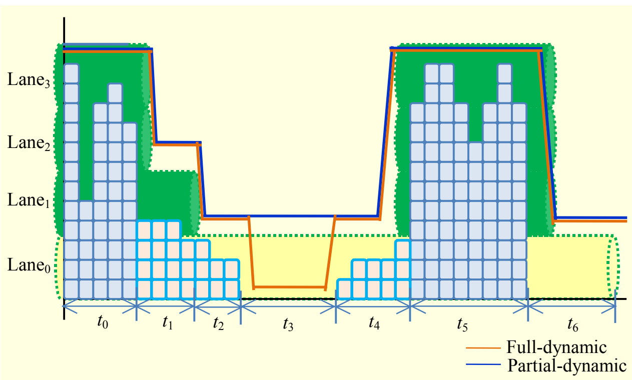

Fig. 2. Comparison of methods for fully dynamic and partially dynamic lane control (n=4).

Figure 2 shows both defined dynamic lane control methods for a four-lane-based Ethernet link. As mentioned earlier, fully dynamic lane control adjusts the number of active lanes, whereas partially dynamic lane control uses a static lane even though there is no data (in period t_3). If low traffic under the capacity of k lanes holds a dominant position for a long period (i=1 in Fig. 2), the former is more energy efficient. However, when the traffic is greater than or equal to i, these two lane control methods are equally energy efficient. Even so, there may be some limitations, such as queuing delay, packet loss rate, and control signaling, when the lane turns on at the end of time period t_3, although a static lane in the latter control may partly alleviate this. In this case, we use the method for partially dynamic lane control that is based on a single static lane. This static lane may be used to send a lane control frame and a data stream.

III. Dynamic Lane Control Scheme

1. Dynamic Lane Decision Algorithm

To determine the number of active lanes, we use the traffic load, queue threshold, current queue size, and increase and decrease rate information of the queue as key factors. The traffic load is used to calculate the average link utilization, and the remaining information in connection with the queue information is used to alleviate the packet loss and ensure fewer lane changes. The proposed algorithm mainly consists of three parts: the measurement of traffic and queue status, a determination of the number of lanes, and an update of the active lane.

Figure 3 shows a schematic of the proposed dynamic lane control scheme. The LCM determines the number of active lanes in a slightly different way, that is, by determining whether an alarm signal is received. The LCM usually calculates the number of lanes during every traffic monitoring period using input traffic and two types of queue information, and it recalculates upon receiving the alarm signal sent from the

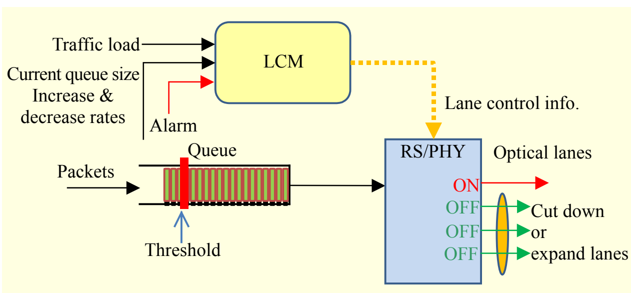

Fig. 3. Block diagram for dynamic lane control.

queue when the current queue size exceeds the determined queue threshold, regardless of the monitoring period. The determined result is sent to the RS/PHY, and the lane control mechanism carries out either lane expansion or contraction.

To calculate the average amount of traffic (R_{avg}), the input traffic is collected during every monitoring period (T) and is calculated using the following expression, where α is the weight of the traffic ($0<\alpha\leq 1$) and kt is the random time within a period ($0 < k < 1$, $t-T < kt < t$):

$$R_{avg} = \alpha R(kt) + (1-\alpha)R(t) \ .$$

The queue threshold, current queue size, and increase and decrease rates are used to reflect the current queue status. The LCM knows the queue threshold (θ_q), which is initially determined as follows, where M_q and β represent the total queue size and permitted queue occupation rate, respectively ($0<\beta<1$):

$$\theta_q = M \times \beta \ .$$

The remaining information is reported to the LCM every T. The increase and decrease rates (γ), which depend on the current queue size (M_{cur}) and average queue size (M_{avg}), are calculated as follows:

$$\gamma(t) = \frac{M_{cur}(t) - M_{avg}(t-1)}{M_{avg}(t-1)} \ .$$

This is less than zero when the amount of traffic decreases and more than zero when the amount increases. If the margin of increases or decreases expands, then the absolute value increases.

The LCM determines the number of active lanes depending on both the measured amount of traffic and the queue status. Different rules are used for deciding the number of lanes according to whether the threshold is exceeded, whether an alarm is experienced, and, in particular, whether the alarm signal is received. The blow pseudocode shows these LCM operations.

Pseudo code for determining the number of active lanes:

Definition of parameters

L: link capacity

ρ: offered load

N: number of active lane determined

N_{max}: maximum number of lanes

N_r: number of lanes required

N_c: current number of lanes

N_s: number of static lanes

M_{cur}: current queue size

R_{avg}: amount of average input traffic

θ_q: queue threshold

γ: increase and decrease rate

δ: increase and decrease threshold

Begin: After initialization, wait for the collected information and alarm signal.

If (alarm_signal) go to Case 2

Else go to Case 1

Case 1: LCM determines the number of required active lanes with traffic and queue information, such as R_{avg}, M_{cur}, and γ.

Step 1. Calculate the required number of lanes depending on the link utilization.

$$N_r = \lfloor \rho \times N_{max} \rfloor, \ \ \rho = \frac{R_{avg}}{L \times T} \ ;$$

Step 2-1. Check whether the current queue size is greater than the threshold.

If ($\theta_q \leq M_{cur}$){

$N = \max(\lfloor N_c + (N_c \times \gamma) \rfloor, N_r)$;

If ($N > N_{max}$) $N = N_{max}$;

Go to Begin;

}

Else go to Step 2-2.

Step 2-2. Check the alarm experience during the monitoring period.

If (alarm_experience){

$N = \max(N_c, N_r)$;

If ($N > N_{max}$) $N = N_{max}$;

Else If ($N < N_s$) $N = N_s$;

Go to Begin;

}

Else

If ($\gamma < \delta$) $N = N_c$;

Else {

$N = N_r$;

If ($N > N_{max}$) $N = N_{max}$;

Else If ($N < N_s$) $N = N_s$;

}

Go to Begin.

Case 2: LCM determines the number of active lanes with the maximum of the two calculated values.

$$N = \max(\lfloor N_c + (N_c \times \gamma) \rfloor, N_r);$$

If ($N > N_{max}$) $N = N_{max}$;

Go to Begin.

2. Adaptive Lane Control Protocol

To control the transmitting lanes, we use the RS-based lane control protocol, which provides no processing delay at the MAC layer and an additional buffer at the PHY for dynamic lane control. It also provides a low delay and no overhead for an exchange of control frames and a simple adaptive data rate.

Upon receiving the lane information from the LCM, the RS performs the lane control protocol by exchanging the control messages. The 64-bit control messages are used for the lane control request, ACK, and transmission through the updated lanes. These messages consist of eight 8-bit subframes (lane 0 through lane 7). The first subframe (lane 0) is set to a sequence control character representing a sequence (ordered set) frame defined for informing the link status in the standards. The fourth subframe (lane 3) contains the lane control information, and the others are set to data characters of 0×00. Lane 3 consists of three fields to represent the type of message and the number of lanes to use, such as 1-bit ID, 2-bit OP-code, and 5-bit number of lane fields. The ID is the identification bit and is set to 1 to indicate a lane control message. The OP-code indicates the type of lane control frame and is respectively set to 10, 01, and 11 in the order mentioned above. Finally, the number of lane fields indicates the number of active lanes for a data transmission.

The procedure of the control message exchange for a dynamic lane control is as follows. After receiving lane information, the RS sends a lane control request over the current active lanes. The corresponding RS receiving the request sends an ACK as a response to the agreement and obtains an alteration of the set values for synchronization and alignment depending on the number of used lanes in advance. After completing a physical lane adjustment, the RS transmits the Begin message to inform the start of a data transmission over the new active lanes. The corresponding RS informs the number of active lanes on the receive path to the PHY. After completing this, the RS inserts idle frames among data frames to be mapped into inactive electrical lanes. The inserted idle frames for data rate adaptation are dropped at the turn off transmitters, and the data frames are transmitted over active lanes.

IV. Performance Evaluation

1. Simulation Parameters

Using the OPNET simulator, we examine the performance of the proposed dynamic lane control scheme in terms of queuing delay, packet loss rate, number of used lanes and lane changes, control packet exchanges and lane change delay, and the energy savings rate.

For this simulation, we consider a 40-G Ethernet optical link with four lanes and a turn on time of 100 ms and a turn off time of 100 μs. The queue size varies from 30 ms to 50 ms in steps of 10 ms. The packets are generated with an exponential distribution in a traffic generator, and the packet size is variable, following the distribution of the IP packet length in [15]. Tables 1 and 2 show the simulation parameters and IP packet length (l) distribution considered for our simulation, respectively.

Separately, we also generate random traffic using the following defined traffic generation scenarios. For this purpose, the packets are generated in the same way mentioned above, while the ten traffic generators used have different start and holding times with a uniform distribution. Each generator

Table 1. Simulation parameters.

Parameter	Value
Link capacity	40 G
Lane capacity	10 G
Number of lanes	4
Number of static lane	1
Monitoring period (T)	500 ms
Weight of traffic (α)	0.6
Queue occupation rate (β)	0.2
Increase and decrease threshold (δ)	2.0
Queue size	30 ms to 50 ms
Number of default lane	1, 3
Turn on/off time	100 ms/100 μs
Simulation time	100 s

Table 2. Ethernet packet length distribution.

Packet length	Probability
$l = 64$	0.03
$64 < l < 322$	0.17
$322 < l \leq 580$	0.18
$580 < l \leq 1,049$	0.12
$1,049 < l \leq 1,518$	0.50
Total	1.00

Table 3. Defined traffic generation scenarios.

Scenario	Start time	Holding time
ts_1	Uniform (1.0)	Uniform (1.0)
ts_2	Uniform (0.5)	Uniform (1.0)
ts_3	Uniform (0.5)	Uniform (0.5)

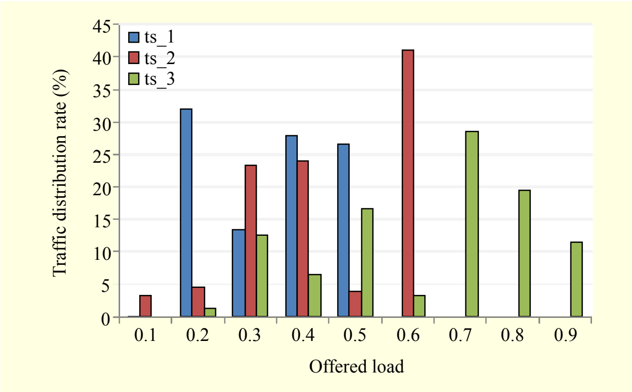

Fig. 4. Distribution of traffic in each scenario.

selects both start and holding times at random with a uniform distribution and begins creating packets based on the results. Table 3 shows three of the defined traffic generation scenarios.

Figure 4 shows the traffic distribution rate generated in the scenarios. The first scenario (ts_1) generates traffic less than or equal to 20 G (offered load=0.5) and is mixed with low (offered load < 0.3) and middle traffic (0.3 < offered load < 0.6) at a rate of almost 50:50. The second scenario (ts_2) generates traffic less than or equal to 24 G, where 24-G traffic (offered load=0.6) makes up 40% of the traffic. Finally, the heavy traffic (offered load > 0.7) holds a dominant position at 60% in the last scenario.

2. Simulation Results

We first examine the performance according to the queue size, the number of default lanes, and the offered load, shown in Figs. 5 through 7. The number of default lanes affects the performance in this case because when the offered load is higher than the capacity of the default lanes, the remaining traffic should be queuing or dropped. We assume it as 1 and 3, in which the former is the worst case in terms of performance and the latter is a case in which the average offered load is considered.

Figure 5 shows the queuing delay and packet loss rate according to the offered load. The queuing delay and packet loss mainly increase when an expansion of lanes is required because the current lane is maintained during a turn on time despite the input traffic increase. As the number of default lanes increases, it gives a low queuing delay and packet loss rate. As the queue size increases, the packet loss decreases but the queuing delay increases. The delay and packet loss increase gradually, although the offered load increases, because the number of active lanes increases as the traffic increases. When the offered load is less than 0.3, packet loss does not occur, regardless of the number of default lanes and queue size.

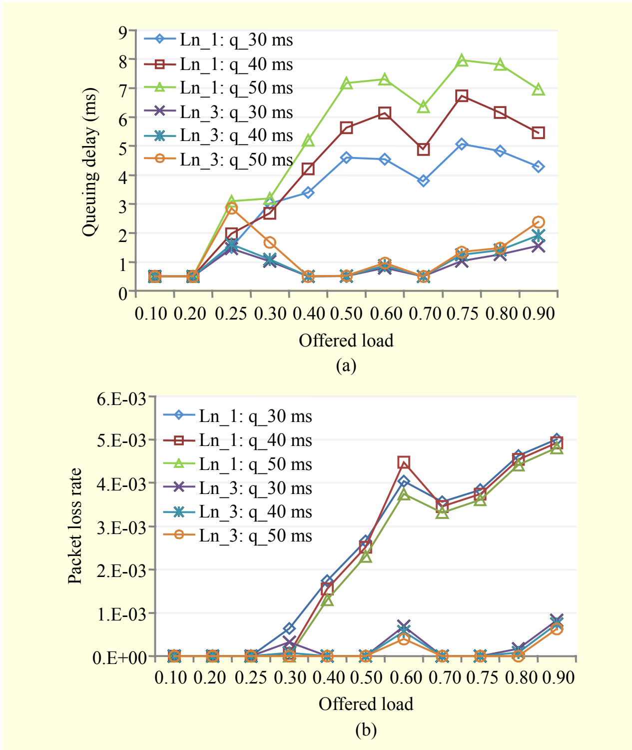

Fig. 5. (a) Queuing delay and (b) packet loss rate.

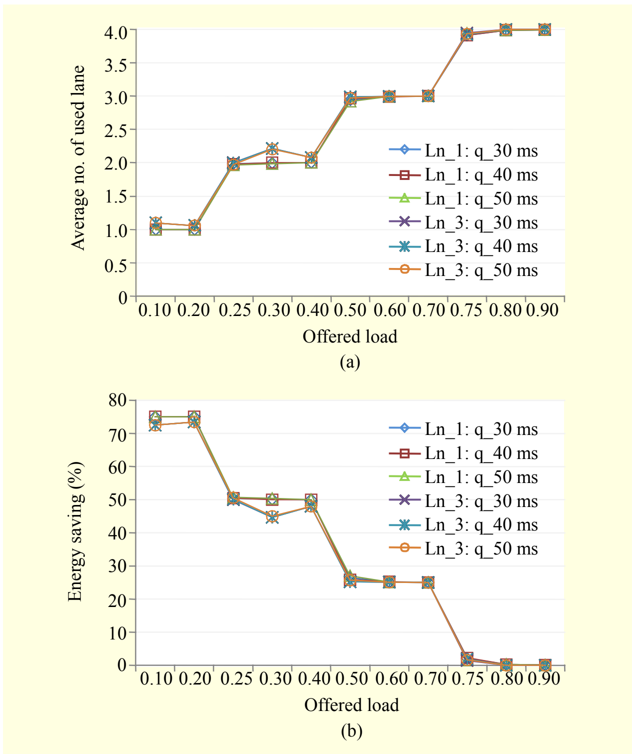

Fig. 6. (a) Number of used lanes and (b) energy savings.

Figure 6 shows the average number of used lanes and energy savings rate against using full lanes. The number of default

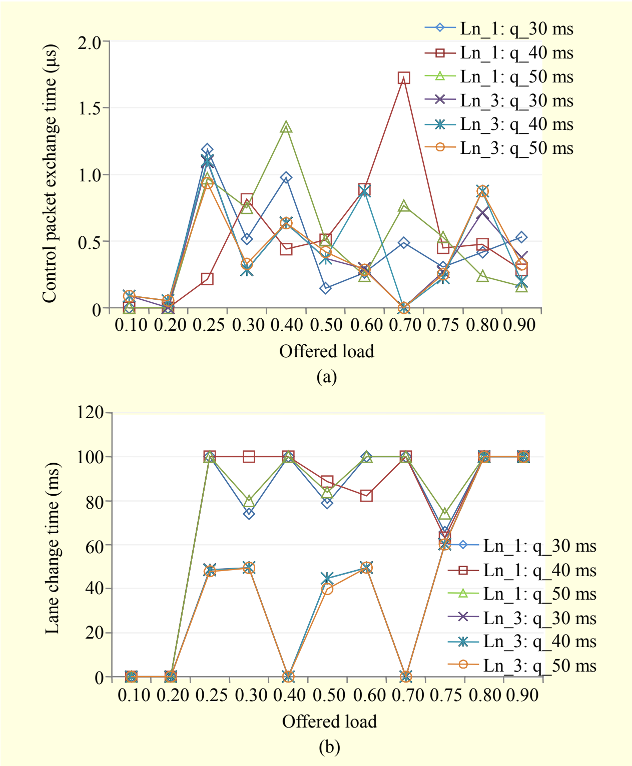

Fig. 7. (a) Control packet exchange time and (b) lane change time.

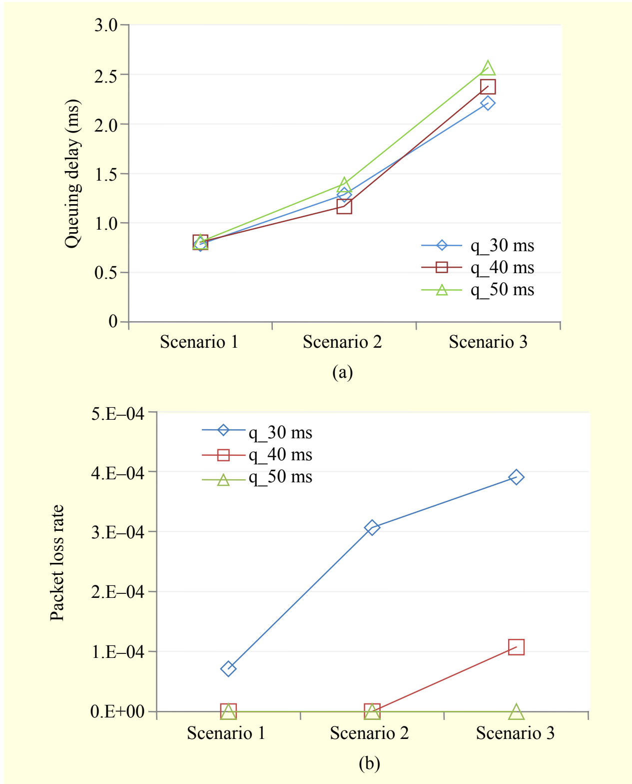

Fig. 8. (a) Queuing delay and (b) packet loss rate.

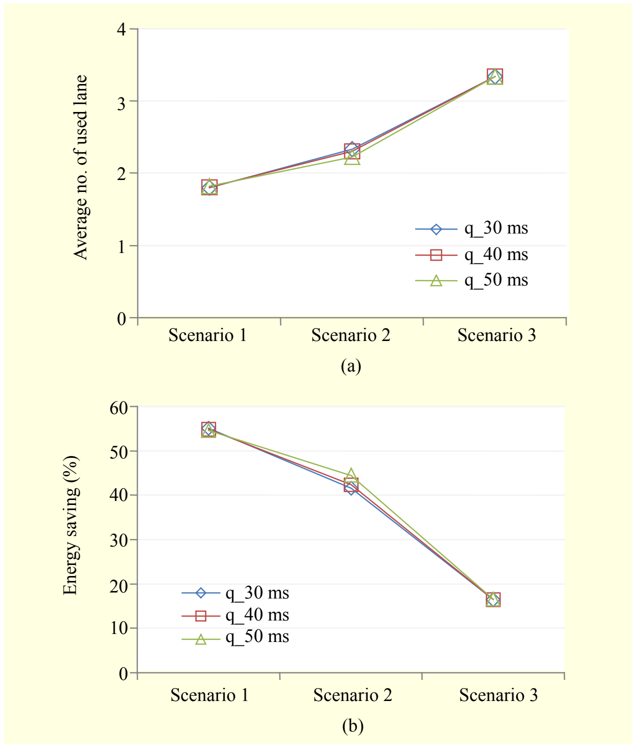

Fig. 9. (a) Number of used lanes and (b) energy savings.

lanes affects the energy savings by 5% or less and shows a similar performance. The queue size has little effect on the amount of energy being used. The energy efficiency decreases with a lane-operation unit of 10 G because the number of lanes required increases to cover the increased traffic.

Figure 7 shows the control packet exchange time and lane change time. In Fig. 7(a), the control packet exchange time does not appear regularly. It is affected not by the network traffic but by the length of the transmitting data packet because the control packets are sent during the interframe time. Lane changing occurs according to the lane-operation unit, which is based on the traffic status. If the default lanes are able to transmit the data without an increase or decrease in the number of lanes, the number is maintained and control packets are not exchanged. In Fig. 7(b), the lane change time includes the time between the completion of the lane change and the start of the transmission over the updated active lanes. The turn on/off time is the main component of the delay, although the exchange of control packets is rapid. However, the rapid delivery of new lane information is important when we consider the reconfiguration time for synchronization and alignment in accordance with the total number of active lanes at the PHY.

Next, we examine the performance according to the queue size under random traffic scenarios defined as in Table 3 and Fig. 4, where we assume one default lane.

Figure 8 shows the queuing delay and packet loss rate under

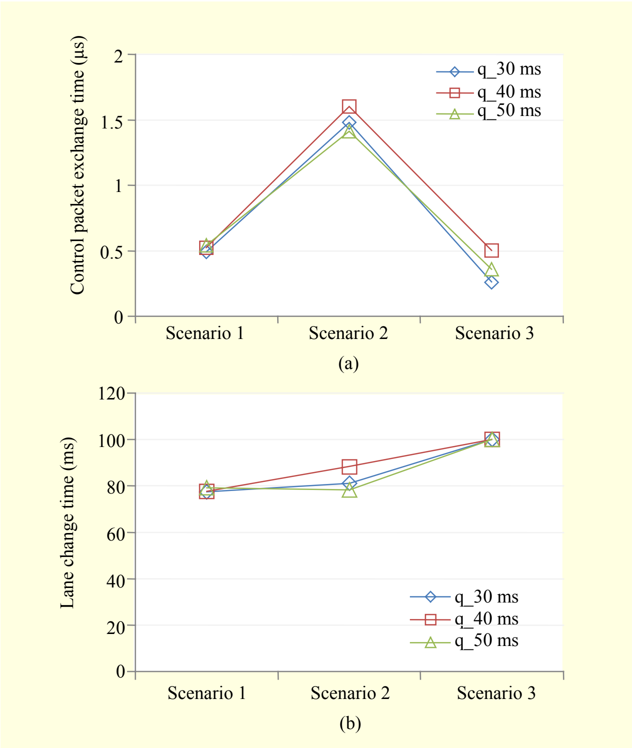

Fig. 10. (a) Control packet exchange time and (b) lane change time.

different traffic scenarios. As the queue size increases, the queuing delay increases while the packet loss rate decreases. In lower traffic, such as Scenario 1, the queuing delay and packet loss rate are less than 1 ms and 10^{-4}, respectively. In particular, when the queue size is 50 ms, the packet loss is zero, regardless of the traffic scenario, while the queuing delay only increases about 200 μs.

Figure 9 shows the number of used lanes and the energy savings. When low traffic holds a dominant position, such as in Scenario 1, a fewer number of lanes are used and the energy efficiency increases. The number of lanes is almost completely unaffected by the queue size. As shown in Fig. 9(b), the energy efficiency is better at around 55% than when a full lane is continuously used for the transmission in a low traffic scenario. This is determined by the average number of used lanes depending on the traffic and queue status. In the case of heavy traffic, as in Scenario 3, energy savings of around 20% is achieved compared with a case in which all lanes are statically used for transmission.

Figure 10 shows the control packet exchange time and lane change time. As mentioned earlier, the control packet exchange time is only affected by the generated length of the data packet, and the control packet exchanging may be completed very quickly in low traffic. The control packet exchange time is very short at around 1.5 μs (no more than 2 μs), while the lane change time is between 80 ms and 100 ms. It may take either 100 μs or 100 ms, owing to the nature of optical devices as the lanes increase or decrease. If this time can be reduced, the network performance may be significantly improved.

V. Conclusion

We proposed a scheme for partially dynamic lane control for energy saving in multilane high-speed Ethernet. In this scheme, the LCM determines the number of active lanes with the traffic load, queue threshold, current queue size, and increase and decrease rates of the queue for measuring the average link utilization and reflecting the current queue status, either periodically every monitoring period or immediately when receiving an alarm from the queue. The RS adjusts the number of lanes with the result received from the LCM by exchanging control messages during the interframe time. We observed the performance according to the queue size, number of default lanes, offered load, and random traffic scenarios. The simulation results showed that it provides energy savings of around 55% with a queuing delay of less than 1 ms and either a packet loss of less than 10^{-4} or no packet loss by increasing the queue. In addition, it has a control packet exchange time of less than 0.5 μs in a random traffic scenario. Thus, the energy savings obtained is around 75% with no packet loss and a queuing delay of less than 1 ms when the offered load is low at less than 0.25.

References

[1] C. Lange et al., “Energy Consumption of Telecommunication Networks,” *Proc. ECOC*, 2009, pp.1-2.

[2] R. Bolla et al., “Energy Efficient in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures,” *IEEE Commun. Surveys Tuts.*, vol. 13, no. 2, 2011, pp. 223-244.

[3] R.S. Tucker, “A Green Internet,” *Proc. LEOS*, 2008, pp. 4-5.

[4] B.D. Choi et al., “Symmetric Adiabatic Logic Circuits against Differential Power Analysis,” *ETRI J.*, vol. 32, no. 1, Feb. 2010, pp. 166-168.

[5] T. Yang, C. Zhao, and D. Chen, “Feedback Analysis of Transcutaneous Energy Transmission with a Variable Load Parameter,” *ETRI J.*, vol. 3, no. 4, Aug. 2010, pp. 548-554.

[6] M.C. Domingo, “Packet Size Optimization for Improving the Energy Efficiency in Body Sensor Networks,” *ETRI J.*, vol. 33, no. 3, June 2011, pp. 299-309.

[7] E. Palkopoulou, A.D. Schupke, and T. Bauschert, “Energy Efficiency and CAPEX Minimization for Backbone Network Planning: Is There a Tradeoff?” *ANTS, IEEE 3rd Int. Symp.*, 2009, pp. 1-3.

[8] C. Gunaratne et al., “Reducing the Energy Consumption of

Ethernet with Adaptive Link Rate (ALR)," *IEEE Trans. Comput.*, vol. 57, no. 4, 2008, pp. 448-461.
[9] S. Nedevschi et al., "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation," *Proc. 5th USENIX Symp. NSDI*, 2008, pp. 323–336.
[10] P. Reviriego et al., "Performance Evaluation of Energy Efficient Ethernet," *IEEE Commun. Lett.*, vol. 13, no. 9, 2009, pp.697-699.
[11] http://www.ieee802.org/3/az/index.html
[12] IEEE Std. 802.3ba-2010, "Part3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, Amendment 4: Media Access Control Parameters for 40Gb/s and 100Gb/s Operation," June 2010.
[13] H. Imaizumi et al., "Power Saving Technique Based on Simple Moving Average for Multi-channel Ethernet," *Proc. OECC*, Aug. 2009, pp. 1-2.
[14] P. Reviriego et al., "Improving Energy Efficiency in IEEE 802.3ba High-Rate Ethernet Optical Links," *IEEE J. Sel. Topics Quantum Electron*., vol.17, no. 2, 2011, pp. 419-427.
[15] S. McCreary and K.C. Claffy, "Trends in Wide Area IP Traffic Patterns," Cooperative Association for Internet Data Analysis, USA, http://www.caida.org

Kyeong-Eun Han received her BS, MS, and PhD in computer engineering from Chonbuk National University, Jeonju, Rep. of Korea, in 2001, 2003, and 2008, respectively. In 2003, she was an Up-and-Coming Young Researcher supported by the Korea Science and Engineering Foundation (KOSEF) and engaged in research on WDM-PON for access area networks. Since 2008, she has been with ETRI, Daejeon, Rep. of Korea, where she is currently a senior researcher. Her research interests are energy efficient Ethernet, multilayer network control, and OTN signaling, including optical communication networks.

Choong-Reol Yang received his BS and MS in electronics engineering from Konkuk University, Seoul, Rep. of Korea, and his PhD in electronics engineering from Chungnam National University, Daejeon, Rep. of Korea, in 1983, 1998, and 2007, respectively. In 1985, he began working with Hyundai Precision, Inc., Seoul, Rep. of Korea. Since 1992, he has been in the Optical Internet Research Department of ETRI, Daejeon, Rep. of Korea, where he is currently a principal researcher. His research interests cover the CMOS design technologies, including optical internet and optical signal processing.

Kwangjoon Kim was born in Seoul, Rep. of Korea, in 1958. He received his BS and MS in physics from Seoul National University, Seoul, Rep. of Korea, and his PhD in physics from Ohio State University, Columbus, OH, USA, in 1981, 1983, and 1993, respectively. He joined ETRI, Daejeon, Rep. of Korea, in 1984 and worked on HF communications until he enrolled in the PhD program at Ohio State University, where he worked on various linear and nonlinear optical behaviors of conducting polymers. He rejoined ETRI and began working on optical semiconductor devices with quantum wells. His current research focuses on WDM optical communication systems and high-speed optical transmission.

Sun-Me Kim received her BS in computer science from Chungnam National University, Daejeon, Rep. of Korea, and her MS in computer science from Pohang University of Science and Technology, Pohang, Rep. of Korea, in 1991 and 1993, respectively. Since 1993, she has been with ETRI, Daejeon, Rep. of Korea, where she is currently a principal researcher. Her research interests are packet-circuit-optical converged switching systems, multilayer transport network control and management in packet-optical converged networks, and IDC network control.

Jonghyun Lee received his BS and MS in electronics engineering and his PhD in communication engineering from Sungkyunkwan University, Seoul. Rep. of Korea, in 1981, 1983, and 1993, respectively. In 1983, he joined ETRI, Daejeon, Rep. of Korea. His current research activities are in optical communication systems and optical Internet technology. Dr. Lee is the head of the Optical Internet Research Department at ETRI.