PAPER

Router Power Reduction through Dynamic Performance Control Based on Traffic Predictions

Hiroyuki ITO[†], *Nonmember*, **Hiroshi HASEGAWA**[†a)], *Senior Member*, *and* **Ken-ichi SATO**[†b)], *Fellow*

SUMMARY We investigate the possibility of reducing router power consumption through dynamic router performance control. The proposed algorithm employs a typical low pass filter and, therefore, is simple enough to implement in each related element in a router. Numerical experiments using several real Internet traffic data sets show the degree of reduction in power consumption that can be achieved by using the proposed dynamic performance control algorithm. Detailed analysis clarifies the relationships among various parameter values that include packet loss ratios and the degree of power savings. We also propose a simple method based on the leaky bucket model, which can instantaneously estimate the packet loss ratio. It is shown that this simple method yields a good approximation of the results obtained by exact packet-by-packet simulation. The simple method easily enables us to derive appropriate parameter values for the control algorithm for given traffic that may differ in different segments of the Internet.
key words: *router, performance control, energy efficiency*

1. Introduction

The amount of Internet traffic has rapidly and continuously increased all over the world [1], since broadband access such as ADSL and FTTH have become widely deployed. This trend will continue because broadband services such as 4-k cinema, ultra-high definition video [2] or e-science [3] are on the horizon. The rapid increase in traffic has resulted in huge power consumption levels at routers, which necessitates the development of energy efficient routers. Indeed, it is estimated that the power consumption of the telecommunication network will double from 2009 to 2017 due to the growth in the power consumption of Layer 3 IP/MPLS backbones routers [4]. In order to reduce effectively the power consumption of future networks, minimizing the increase in power consumption in IP/MPLS core routers is essential. Core routers consume an almost constant high level of power regardless of their loads. Indeed, Chabarek et al. [5] reported that router power consumption in the idle state is only 10% less than that for the fully loaded state. Almost the same data were reported in [6]. This is because Internet traffic rapidly fluctuates and so it is very difficult to predict future traffic volume, which is necessary to control routers. The processing performance of present routers is hence kept high so that they can process the maximum traffic volume at any time. A typical large core router consumes approximately 30% of the total power in a packet processing engine and approximately 35% for its power supply and blowers [7]. The power consumption of the power supply and blowers increases in conjunction with the power consumption at the rest part of router. Thus, the router power consumption can be effectively reduced by reducing the power consumption of the processing engines, which consume almost the same level of power regardless of their processing load.

In order to reduce the power consumption when the traffic volume is low, several studies considered dynamic performance control of transport systems to achieve energy savings [8], [9]. These studies investigated the effect of load concentration that is attained by traffic engineering, where traffic is concentrated at some routers, and the rest of the routers are turned off. A conventional study [8] proposed a method to achieve power saving in IP-WDM networks by turning off line cards or IP routers that are not necessary for each time interval. The information on traffic variation in a network must be given *a priori* and this minimizes the number of network disruptions. In [9], Idzikowski et al. formulated an optimization for route assignment and line card activation to minimize energy consumption. The formulation uses all the traffic matrix information including future traffic volume for perfect planning and exact evaluation of power consumption. Although significant improvements were verified, it is hard to apply this method directly to real networks due to the difficulty in obtaining exact predictions of future traffic. In [10], dynamic activation of links and route relocation using only the observed traffic volume were considered. The candidate links to be turned off are selected preliminarily considering the connectivity between nodes.

Other studies focused on dynamic performance control of the transport systems themselves whose processing ability can be changed in a stepwise manner. The advantage of this approach is that the technique can be introduced locally without affecting other parts of a network and without applying network wide control. Studies [11] and [12] focused on dynamically changing the link speed according to the current traffic information. In [11], the link speed of Ethernet ports of a PC and LAN are adaptively altered between two data rate states based on the current link or buffer memory utilization. This method is shown to be effective when traffic is bursty and the average utilization ratio is low. This occurs typically only when a few users are connected to the link. In [12], power consumption reduction in switches attained by employing link aggregation was proposed. Switches measure the arriving traffic volume and manage the total capacity for each link, that is, increase or decrease in the num-

Manuscript received January 17, 2012.
Manuscript revised June 5, 2012.
[†]The authors are with Nagoya University, Nagoya-shi, 464-8603 Japan.
a) E-mail: hasegawa@echo.nuee.nagoya-u.ac.jp
b) E-mail: sato@nuee.nagoya-u.ac.jp
DOI: 10.1587/transcom.E95.B.3130

ber of active physical links between switches. However, the power consumption of the links is not dominant (routing processing is dominant) in a network and hence power management must be done using other techniques to achieve substantial power reduction. Recent high-speed routers consist of multiple packet processing engines and are equipped with several interfaces. Some of the routers are expected to be equipped with a function that can control the number of activated packet processing engines so that the power consumption can be reduced. Currently, the number of activated engines cannot be dynamically changed; however, dynamic performance management techniques have been developed and will be implemented (This architecture was suggested in [13]). This will reduce useless power consumption when the traffic volume is low. In order to make the best use of these advanced router capabilities, it is very important to develop an efficient algorithm that can determine the minimum number of necessary activated routing engines that will not cause packet loss due to limited processing capabilities.

The optimal processing performance level that can assure routing performance can be determined based on accurate traffic prediction. For example, traffic prediction methods based on wavelet multi-resolution decomposition were presented in [14] and [15]. Although they try to achieve accurate prediction, the prediction error is not sufficiently low to allow efficient dynamic router performance control. The point of dynamic performance control is to minimize the number of activated packet processing engines while avoiding buffer overflows, even when detailed traffic prediction is not available. In order to achieve effectively dynamic power management in each part of the routers, a simpler control algorithm is preferable. To attain this, we need to develop a simple algorithm that estimates the upper bound of slowly changing components in traffic variations since a rapidly changing component in the traffic variations can be effectively handled by setting an appropriate number of packet buffers. Another important point is that the algorithm should rely on only a few limited parameters that can be easily determined so that the algorithm can adapt to various kinds of traffic flows that exhibit different characteristics.

The objective of this paper is to propose a method that allows dynamic router performance control and to analyze its effectiveness. We first propose an algorithm that can periodically estimate the necessary processing speed at the next time interval. The algorithm applies a typical low pass filter to derive the exact fluctuation information of low frequency components in the traffic variations. The gap between the level determined by the low frequency component and the necessary processing power at each time interval is determined using a multiple of the standard deviation of the high frequency component. The sum of these values determines the necessary routing performance at each operational timing. We proposed an initial version of this algorithm in [16], the purpose of which is dynamic path control but it did not consider the management of high-frequency component traffic, which is possible with the use of buffers. The low pass filter we use is an IIR (Infinite Impulse Response) filter of low degree and therefore a very low computational load is necessary. In this paper, the effectiveness of the performance of the proposed algorithm is investigated using two schemes: several packet-by-packet exact simulations that re based on real traffic data and a scheme using a simple router processing model based on the leaky-bucket model. The former needs substantial computer resources and requires a long computation time but the results are exact, while the latter can be done within a few seconds but the results are approximations. The numerical experiments verify the effectiveness of the proposed dynamic performance control method and determine parameter values that yield target packet loss ratios. We show that the results obtained by the two evaluation methods have an almost linear relationship and, therefore, the approximation evaluation based on the leaky-bucket model can be a computationally efficient alternative to the exact but computational power demanding evaluation. The proposed method utilizes only two parameters, the cutoff frequency of the low pass filter and the magnification factor that determines the router performance of each time period, which reflects the high-frequency traffic fluctuation. With the instantaneous measurement of buffer overflow, these two parameter values can be easily optimized, and therefore, the method is applicable to a diverse range of traffic types.

2. Preliminaries

We assume a router that is equipped with multiple routing engines, the routing performance for which can be switched to several states where each state corresponds to the number of activated routing engines. That is, its power consumption and routing performance are proportional to the number of activated engines and they are represented by an N level step function, $\frac{k}{N}$ $(k = 1, 2, \cdots N)$. Here the power consumption in the idle state is set to zero for simplicity (this can be easily modified). The maximum routing speed, S [bps], is determined so that the router can process the peak traffic volume. We also assume that the router periodically measures the incoming traffic volume, trf, at every short time interval, Δt, and control the switches to the adequate processing performance at much longer time intervals. The value of Δt is determined considering the buffer size of commercially available routers and to lessen the processing burden of traffic measurement. For example, typical core routers place buffer memories with capacity of around 100 ms at each port given the maximum line speed. Considering this, in this paper, Δt is set to 100 ms. Figure 1 illustrates the router model used in this paper. The router input/output buffers are placed before/after the routing engines, and these two buffers are separated for simplicity. The input buffer memory size is fixed and independent from the routing engine performance.

In the following, we focus on parameters for the processing engines and input buffers. In this context, packet loss occurs only when the input buffer memory overflows due to an insufficient number of activated routing engines. In the next section, we propose how to estimate the required

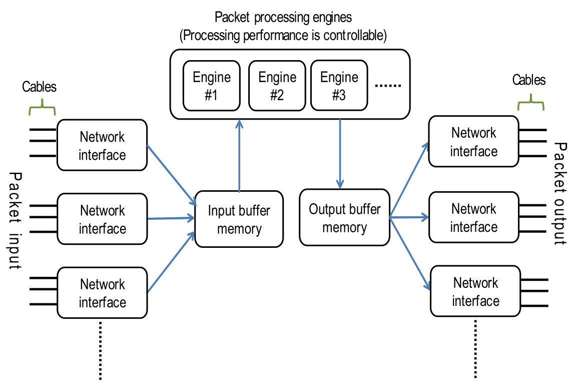

Fig. 1 Router model equipped with dynamic performance control.

number of engines to activate. Once a decision is made to change the routing performance, this decision is reflected immediately. However, fast changing of the routing performance is difficult since frequent turning off/on of the routing engine is difficult. We introduce a notion of *sleep time,* which is the shortest acceptable interval between subsequent performance changes. The effect of changing the *sleep time* is evaluated by numerical experiment.

3. Dynamic Routing Performance Control Algorithm

In this section, we propose a method that dynamically estimates the necessary routing performance level based on the history of the received traffic volume information. The objective is to determine the required routing performance at the next operation timing, $\frac{k[m+1]}{N}S$, using the past traffic volume, $\{\ldots, trf[m-1], trf[m]\}$. If we do not take the input buffer into consideration, the processing performance should satisfy $\frac{k[m]}{N}S \geq trf[m]$ while keeping the extra amount of processing performance, $\frac{k[m]}{N}S - trf[m]$, sufficiently small. This inequality can be violated to some extent when the input buffers are used and as long as the buffer overflow is within the specified value. As we stated in Sect. 2, the processing performance is fixed for the specified *sleep time* after each performance change. Thus, if the routing performance is changed at operation timing m, value $k[m]$ should satisfy the following inequality.

$$trf[m+i] \leq \frac{k[m+1]}{N}S \text{ for all } i \in \{1, 2, \cdots, sleep\ time/\Delta t\}.$$

It has already known that network traffic exhibits long range dependency [17], [18] and it is hard to predict exactly the traffic volume at the next time instant. Thus, we divide the traffic volume into a slowly changing component, $l_{trf}[m]$, and a fluctuating noisy component, $trf[m] - l_{trf}[m]$. We determine component $l_{trf}[m]$ by the output of a typical IIR low-pass filter, the Butterworth filter [19]. The transfer function of a discrete time Butterworth filter of second degree is given by

$$H(z) = \frac{j\sqrt{2}\omega_c/2}{1 - \exp[-\sqrt{2}(1+j)\omega_c T/2]z^{-1}} - \frac{j\sqrt{2}\omega_c/2}{1 - \exp[-\sqrt{2}(1-j)\omega_c T/2]z^{-1}}$$

where ω_c is the cut-off frequency and T is the sampling interval. In this paper, we set $T = \Delta t$. With normalization at $z = 1$ and some substitution, we have

$$H_n(z) = \frac{H(z)}{H(1)} = \frac{A_3 z^{-1}}{1 - A_1 z^{-1} + A_2 z^{-2}}$$

where A_1, A_2, and A_3 are real values. Through the inverse z transformation, we have

$$y[n] = A_1 y[n-1] - A_2 y[n-2] + A_3 x[n]$$

where $(x[n])_{n \in Z}$ and $(y[n])_{n \in Z}$ are the input and output sequences, respectively. Next, we estimate the essential part of the variation in the traffic volume $\widehat{trf}$ at the next interval, between m to m+*sleep time*, using the above filter. Since the variation in $l_{trf}[m]$ changes smoothly and slowly, we employ the following assumption at $t = m$.

$$\widehat{trf}[m+i] \approx l_{trf}[m] = A_1 l_{trf}[m-1] - A_2 l_{trf}[m-2] + A_3 trf[m] \text{ for all } i \in \{1, 2, \cdots, sleep\ time/\Delta t\}.$$

Next, we discuss how to evaluate the necessary number of activated routing engines. Let σ be the standard deviation from $\widehat{trf}[m+1]$ at $t = m$ to $trf[m+1]$, i.e., $\sigma = (E[|\widehat{trf}[m+1] - trf[m+1]|^2])^{1/2}$. This value, σ, characterizes the average behavior of the high-frequency component. Therefore, we assume that σ can be estimated based on the history of the change in traffic volume or the behavior of recent past traffic. With fixed positive value t, the necessary routing performance at $m+1$ is predicted by

$$\frac{k[m+1]-1}{N}S < l_{trf}[m] + t \times \sigma \leq \frac{k[m+1]}{N}S$$

If $k[m] \neq k[m+1]$ and m are not in the *sleep time*, the processing performance is changed. If the prediction error, $trf[m] - \widehat{trf}[m]$, follows a normal distribution and the *sleep time* is set to 0, the probability that $trf[m]$ exceeds $\widehat{trf}[m] + t \times \sigma$ (hereafter exceeding probability) for $t = 3$ yields is almost 0.13%. The numerical experiments on real traffic data, 5 minute sequences of arriving traffic volume, show that $t = 3$ yields an exceeding probability that is approximately 1.1% higher than that for the normal distribution [16]. The setting of t ratios will be evaluated by numerical experiment in Sect. 4.1. In the next section, we will evaluate the packet loss probability using various real traffic types employing exact packet-by-packet simulation.

4. Numerical Experiments

In this section, we evaluate the dynamic performance control effect of a router. We assessed the effects using five sets of traffic volume data observed in different networks or in different years. The data sets “Leipzig1” and “Leipzig2”

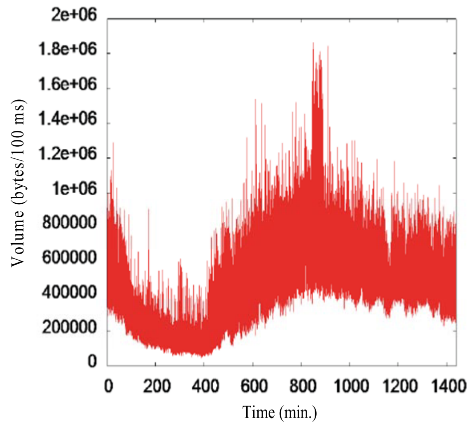

(a) Leipzig1

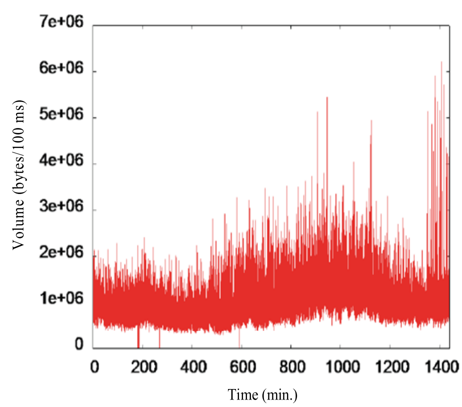

(b) WIDE2008

Fig. 2 Fluctuation in traffic volume.

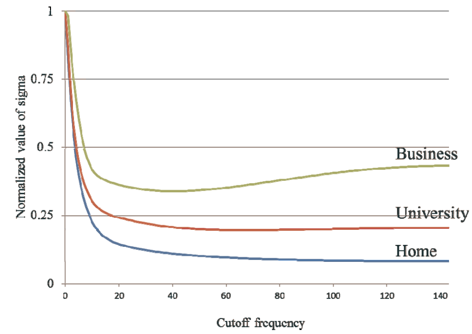

Fig. 3 Relationship between cutoff frequency and normalized standard deviation for three different traffic sequences in [23].

were recorded at the University of Leipzig Internet access link [20]. The data were collected by the NLANR PMA Project. The other data sets "WIDE2008," "WIDE2009," and "WIDE2010" were traces of a trans-Pacific link in the WIDE backbone network measured in 2008, 2009, and 2010, respectively [21], [22]. The maximum traffic volume of each sequence is less than 100 Mbps, however, the following discussion will also be valid for traffic at higher line speeds since the statistical multiplexing effect will be enhanced and traffic prediction error be reduced. Each data set consists of pairs of arrival times and the length of the received packets. We converted each data set to construct the arriving traffic volume every 100 ms. These sequence data sets are supplied to a router and the performance is analyzed when the proposed control scheme is applied. Figure 2 shows some of these converted traffic data sequences. We consider several parameters and metrics in this work. Some parameters are mostly hardware oriented ones and their values are given in advance or limited to a certain range due to the hardware limitation. Moreover, there exist tradeoffs between parameter values that must be considered when finding an appropriate parameter setting. For example, the cutoff frequency of the low-pass filter will affect the standard deviation σ. Figure 3 shows the relationship of these two parameters for three real traffic data sequences in [23] for home, business, and university clients. Almost the same trend is observed regardless of traffic data type. In this case, when the cut off frequency is around 30 freq./day (i.e. $=60\pi$ radians), the standard deviation will be significantly smaller than the value at the origin. Although the ideal way is to show the results for all possible parameter settings, it is impossible to evaluate and illustrate the results since so many parameter values need to be varied. Instead, we focus our investigations on parameter value ranges of practical importance. One of the most important metrics is the packet loss ratio. For realizing real-time services including VoIP and teleconference or real-time video streaming, the acceptable packet loss ratio standardized by ITU-T is 10^{-3} to 10^{-5} [24]. In the application of the proposed method, the acceptable packet loss ratio is therefore given first, then parameters are set to reasonable values, and then the ranges of the other parameter values are set to appropriate ones; finally, we evaluate the power consumption reduction. In this section we analyze the relationships between the packet loss ratio and the other metrics. We assume the following fixed conditions throughout this section. The standard deviation, s, is evaluated on the same converted traffic volume data in advance since the available data length is limited (a few days). However, in practical situations, s will be estimated using observed traffic variations during desired periods in the past that can maximally consider the periodicity of traffic, i.e., day and week periodicity. The number of steps of the controllable packet processing performance, N, is set to 20. The maximum of routing performance S is set at the maximum traffic volume whose resolution is 100 ms. In each experiment, the header of the data sequence with the length of 30 min is used to stabilize the output of the low-pass filter. The results corresponding to the first 30 min are not utilized for analysis and the performance is evaluated for the rest of the duration. We assume that the relationship between the routing performance and power consumption is basically proportional as shown Fig. 4.

In this paper, we developed two evaluation methods to measure the efficiency of dynamic performance control. The first method, referred to as exact evaluation, analyzes the exact behavior of packets by simulating the router processing.

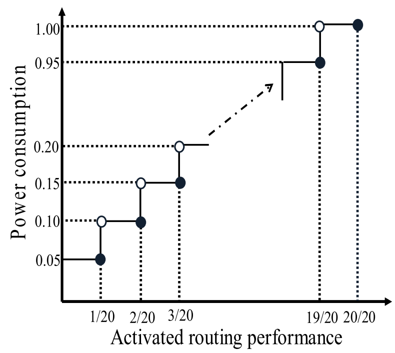

Fig. 4 Relationship between active routing performance and power consumption.

The exact evaluation was done by using commercially available software [25]. By generating packets with the same timing and size as found in the observed data, the exact evaluation analyzes the packet loss ratio, which is the ratio of the number of dropped packets to that of all received packets. The second method, referred to as the approximation evaluation, estimates the packet loss by using the leaky bucket model for a router and the same traffic volume data with the time resolution of 100 ms. That is, at every 100 ms, the differences in the received traffic volume and current processing speed are summed up as far as the sum is positive. When the value exceeds the size of the buffer, it is counted as overflow. We evaluate the performance of the proposed dynamic performance control using the ratio of the number of time samples of overflow to that for all time samples. In the following, we call this ratio the violation probability. The exact evaluation produces very accurate results by emulating a real router using real traffic while the evaluation requires much longer computational time than that necessary for the approximation evaluation. The correlation of these two evaluation results is discussed in Sect. 4.2.

4.1 Effect of Dynamic Performance Control

In this section, we discuss the results of the exact evaluation. The *sleep time* is fixed to 10 s. The cut-off frequency of the low pass filter is set to 60π radians per day to evaluate the basic performance of the proposed router dynamic performance control. The effects of the cut-off frequency are discussed in Sect. 4.3. First, we investigate how to set value t so that the specific packet loss ratio is satisfied. Figure 5 shows the variation in the packet loss ratio against t. The buffer memory size is set from $\frac{S}{20}$ to $\frac{S}{2}$, that is, the memory size corresponds to the maximum total traffic amount for 50 [ms] to 500 [ms]. For the WIDE2008, WIDE2009, and WIDE2010 data sets, some large t values incur no loss. At $t = 3$ and the buffer memory size of $\frac{S}{10}$, the packet loss ratio is less than 0.1% for five traffic data sets. Regardless of the buffer memory size, the packet loss ratios are almost the same for small t values while the variation among traffic data sets becomes large as t increases. This means that we can set $t = 3$ or 4 when the packet loss ratio of 0.01% is acceptable. The above discussion has been validated in our early study [23] that examined real traffic sequences captured in different services (clients were home, business and university users). Although their traffic characteristics are different, they exhibit almost the same trend. Figure 6 shows

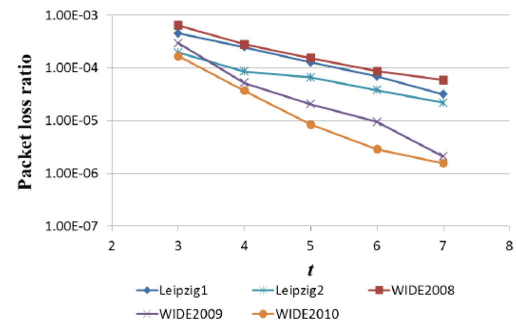

(a) Buffer memory size is S / 20

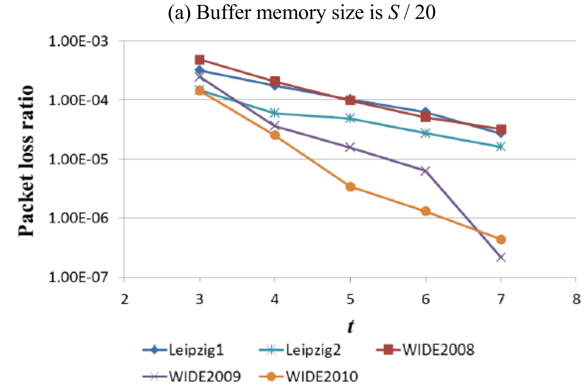

(b) Buffer memory size is S / 10

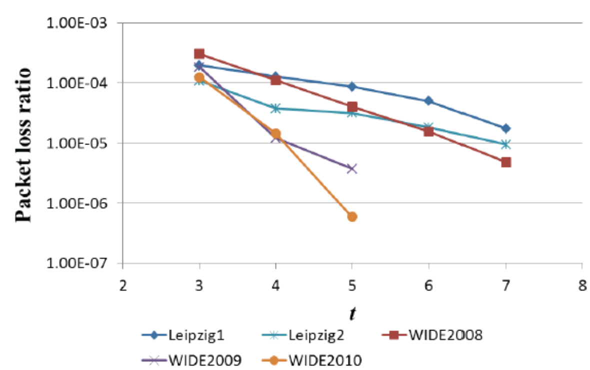

(c) Buffer memory size is S /5

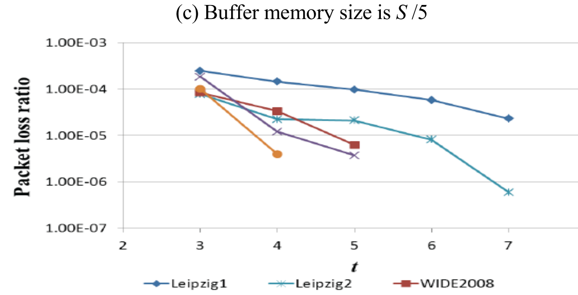

(c) Buffer memory size is S /2

Fig. 5 Variation in packet loss ratio against t.

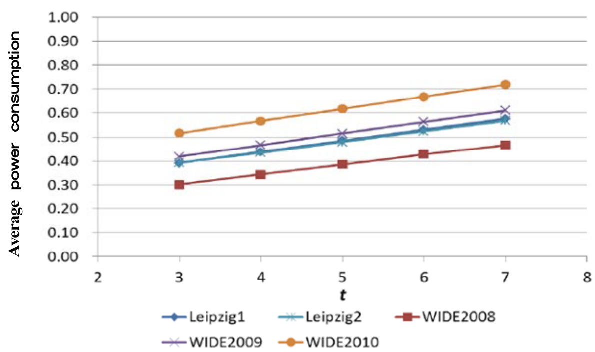

Fig. 6 Variation in average power consumption against t.

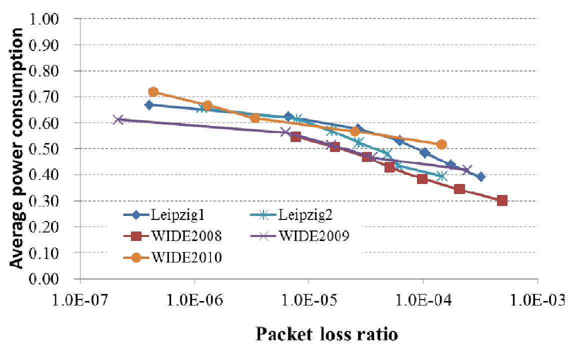

Fig. 7 Power consumption evaluation.

the variations in the average power consumption of a router for one day against t when dynamic performance control is applied. This variation is independent from the buffer memory size. The power consumption increases lineally when t increases for all five data sets. Thus, if we calculate the average power consumption for two different t values, we can es timate the average power consumption for other t values. Thus, the achieved power savings can be evaluated for the desired packet loss ratio using the relationships between the packet loss ratio and t shown in Fig. 5 and using Fig. 6. Figure 7 shows the variations in the power consumption against the packet loss ratio when buffer memory size is set at $\frac{S}{10}$, i.e., the memory can store all the received traffic for at least 100 [ms] when it is initially empty.

Each point is derived by changing the t values (each t value determines the packet loss ratio) where t is changed from 3 to 7 for WIDE2009 and WIDE2010 and from 3 to 9 for the other data sets. The curves confirm the trade-offs between the packet-loss ratio and the reduction in the power consumption. The reduction in power consumption attained with dynamic performance control depends on the traffic characteristics and given tolerable packet loss ratio. When the tolerable packet loss ratio is approximately 0.01%, the power consumption can be reduced from approximately 45 to 60% compared to that without dynamic performance control.

Figures 8 and 9 show the variations in the power consumption against the packet loss ratio when the buffer memory size is changed from $\frac{S}{20}$ to $\frac{S}{2}$ for the Leipzig1 and

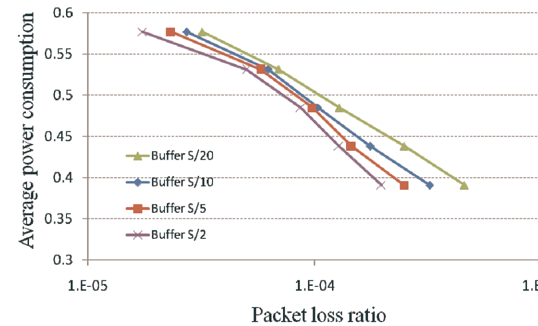

Fig. 8 Effect of buffer memory size (Leipzig1).

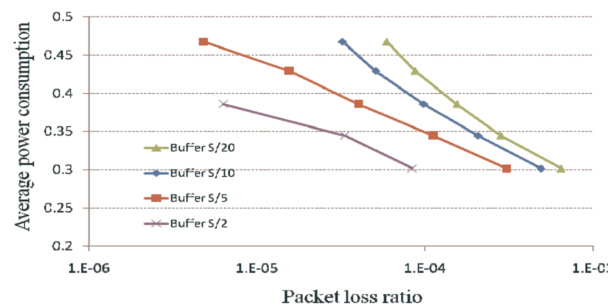

Fig. 9 Effect of buffer memory size (WIDE2008).

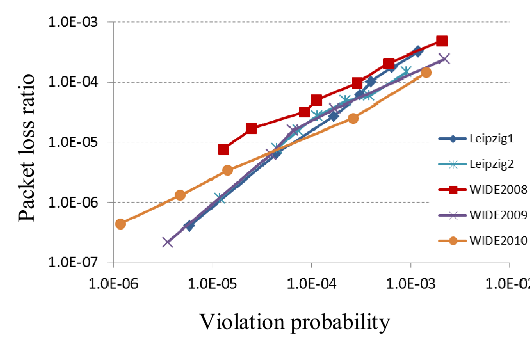

Fig. 10 Relationship between violation probability and packet loss ratio.

WIDE2008 data sets, respectively. The effect of increasing the buffer memory size depends on the traffic data sets. A larger impact is observed for the WIDE2008 data set. For the WIDE2008 data set, the average power consumption difference between buffer memory size $\frac{S}{20}$ and $\frac{S}{2}$ reaches about 15% when the packet loss ratio is 10^{-4}.

4.2 Comparison of Two Evaluation Methods

In this section, we compare the two proposed evaluation methods. Throughout this comparison, we use the following values: the *sleep time* is set to 10 s, the buffer memory size is $\frac{S}{100}$, and the cut-off frequency of low pass filter is set to 60π radians per day. Figure 10 shows the packet loss ratio against the violation probability for five traffic data sets. The packet loss ratio is obtained by the exact evaluation and the violation probability is derived by the approximation evaluation. Comparing these results, the packet loss ratio is always smaller than the violation probability, and we observe a strong correlation between the two ratios. Since the computational time for the violation probability is much shorter than that for the packet loss ratio, calculating the violation probability is a useful alternative and provides good approximation for packet loss ratios. In order to search for appropriate parameter ranges for the router control for given traffic, the approximation evaluation can be utilized as a computationally efficient technique.

4.3 Selection of Cut-Off Frequency and Sleep Time

In this section, we investigate the effect of *sleep time* and the appropriate cut-off frequency using the Leipzig1 data sets. We employ the leaky-bucket based approximation evaluation in this section. Throughout this comparison, we use the following values: the buffer memory size is $\frac{S}{100}$ and the value of t is set to 5. Figure 11 shows (a) the router average power consumption, (b) the violation probability, and (c) the frequency of the switching operations for routing performance adaptation during 23.5 hours, i.e., 1 day minus 0.5 hours for filter output stabilization, against the cut-off frequency of a low-pass filter. The sleep time is changed from 1 [s] to 1000 [s]. Since almost the same trends are observed for cases t= 3 to 7, we show the results when t = 5. The power consumption depends mainly not on the *sleep time* but on the cut-off frequency, and the effect of increasing the cut-off frequency on the reduction in the power consumption is less than 10% for the tested cut-off frequency range as shown in Fig. 11(a). If we increase the cut-off frequency, σ decreases and the frequency of the performance switch increases, which is thought to be the major reason for the reduction in the average power consumption against the increase in the cut-off frequency. However, as shown in Fig. 11(b), when the *sleep time* is long, 100 s or 1000 s, the violation probability gradually increases when the cut-off frequency is increased. This is because when the *sleep time* becomes long, adaptability of the router performance to the traffic variations will become limited as is explained below. Figure 11(c) shows that the frequency of the performance switching increases as the cut-off frequency increases to the level that is determined by the *sleep time.* The saturation levels increase almost in an inversely proportional manner to the *sleep time* (When the *sleep time* increases one order of magnitude, the saturation level decreases almost one order of magnitude). This indicates that when the cut-off frequency is high and the *sleep time* is long, the frequency of the performance switch is limited to a relatively small value (when the *sleep time* is 1,000 [s], it is limited to 10^2), and this limits the adaptability of the router performance to traffic variations.

As discussed above, the three performance measures, router average power consumption, violation probability, and frequency of switching operations are related to each other. As observed from Fig. 11(a), the router power reduction becomes large when the cut-off frequency becomes large. If the acceptable packet loss ratio is 10^{-3}, the allowable violation probability is larger than 10^{-3}, which is derived in Fig. 10. In this case, the cut-off frequency can be 10^5, which is derived from Fig. 11(b). The cut-off frequency determines the frequency of the router performance switches and the frequency strongly depends on the *sleep time* as shown in Fig. 11(c). From an operational point of view, the frequency of the router performance switch should be minimized. When the cut-off frequency is 10^5, we should adopt the *sleep time* of 1000 [s] so that the performance switch frequency can be minimized and very efficient router control can be attained with a relatively large reduction in the router power. Using these results, we can determine appropriate parameters for router performance control depending on the required quality of service (QoS) such as the packet loss ratio or we can evaluate the available power reduction depending on the possible router performance switch interval (*sleep time*).

Here we note the following. In Sects. 4.1 and 4.2,

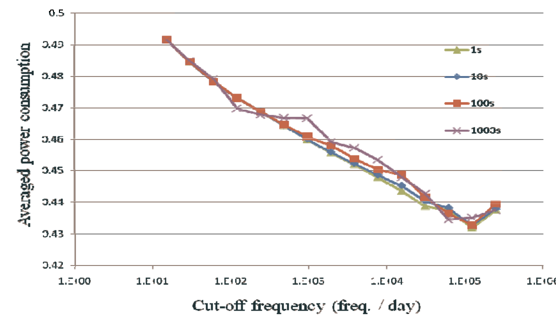

(a) Variation in average power consumption

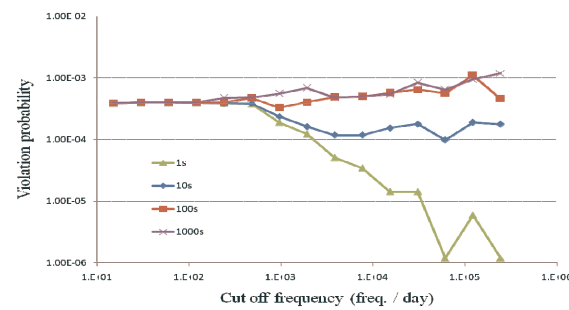

(b) Variation in the probability of violation

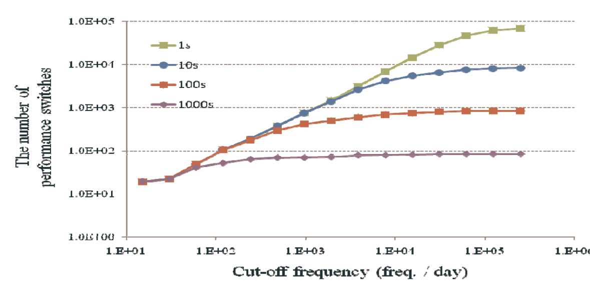

(c) Variation in the frequency of performance switches

Fig. 11 Effect of changing cut off frequency and *sleep time*.

the cut-off frequency is set to 60 πgadians per day, i.e. 30 freq./day. With this setting, output of the low-pass filter is well smoothed since the period of the sinusoidal wave of the cut-off frequency is 48 minutes, which is sufficiently longer than the sleep time discussed in this section. Indeed, as shown in Figs. 11(a)–(c), for the cut-off frequency value, the results do not depend on the *sleep time*. Hence, the results presented in Sects. 4.1 and 4.2 do not change when the *sleep time* is changed from 1 [s] to 1000 [s]. Moreover, this trend is expected to be demonstrated by the other traffic types since similar low cut-off frequencies are used.

5. Conclusion

In this paper, we investigated the reduction in router power consumption that is achieved by applying the proposed dynamic routing performance control scheme. The proposed dynamic routing performance control algorithm employs a typical low pass filter and is therefore simple enough to introduce to each related element in a router. In the numerical experiments using several real Internet traffic data sets, it was shown that the power consumption is reduced up to 50–70%. It was also verified that the leaky bucket model can yield packet loss ratios that well approximate those indicated by packet-by-packet exact simulations. We verified that the estimation of the packet loss ratio can be performed in a short time, within a few seconds. Based on the evaluations, we can determine the appropriate parameter values that are needed for the performance control algorithm in a short time. Please note that this process is necessary only when commencing operation and an update is not always necessary until apparent traffic characteristics change, or it will be done periodically, say, once a day. It was shown that by using the parameter values, efficient router control is achieved and substantial router power reduction is attained.

Acknowledgments

A part of this work was supported by the NEDO Green IT project.

References

[1] http://www.jpix.ad.jp/jp/technical/traffic.html
[2] ITU-R Recommendation BT.1769, "Parameter values for an expanded hierarchy of LSDI image formats for production and international program exchange," 2006.
[3] M.D. Brown, ed., "Blueprint for the future of high-performance networking," Commun. ACM, vol.46, no.11, pp.30–77, Nov. 2003.
[4] C. Lange, D. Kosiankowski, C. Gerlach, F. Westphal, and A. Gladish, "Energy consumption of telecommunication networks," Proc. ECOC 2009.
[5] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," IEEE INFOCOM, 2008.
[6] D. Kharitonov, "Time-domain approach to energy efficiency: High-performance network element design," GLOBECOM Workshops, 2009.
[7] J. Baliga, R. Ayre, K. Hinton, and R.S. Tucker, "Photonic switching and the energy bottleneck," Proc. Photonics in Switching, 2007.
[8] Y. Zhang, M. Tornatore, P. Chowdhury, and B. Mukherjee, "Time-aware energy conservation in IP-over-WDM networks," Proc. Photonics in Switching, 2010.
[9] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, "Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios," ONDM, 2010.
[10] V. Soteriou and L-S. Peh, "Dynamic power management for power optimization of interconnection networks using ON/OFF links," Proc. High Performance Interconnects, 2003.
[11] C. Gunaratne K. Christensen, and B. Nordman, "Managing energy consumption costs in desktop PCs and LAN switches with proxying, split TCP connections, and scaling of link speed," Int. J. Network Mgmt, vol.15, pp.297–310, 2005.
[12] Y. Fukuda, T. Ikenaga, H. Tamura, M. Uchida, K. Kawahara, and Y. Oie, "Performance evaluation of power saving scheme with dynamic transmission capacity control," GLOBECOM Workshops, 2009.
[13] M. Yamada, T. Yazaki, N. Matsuyama, and T. Hayashi, "Power efficient approach and performance control for routers," ICC Workshops, 2009.
[14] Y. Luan, "Multiresolution traffic prediction: Combine RLS algorithm with wavelet transform," Proc. ICOIN, 2005.
[15] V. Alarcon-Aquino and J.A. Barria, "Multiresolution FIR neural-network-based learning algorithm applied to network traffic prediction," IEEE Trans. Syst., Man, Cybern. C, Appl. Rev., vol.36, no.2, pp.208–220, 2006.
[16] H. Ito, H. Hasegawa, and K. Sato, "Impact of path granularity and operation interval on dynamic path network control," Proc. ACP, 2009.
[17] S. Mallat, "A wavelet tour of signal processing," Academic Press, ISBN 0-12-466606-X, 1999.
[18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," IEEE/ACM Trans. Netw., vol.2, no.2, pp.1–15, Feb. 1994.
[19] A.V. Oppenheim and R.W. Schafer, Discrete-Time Signal Processing, Prentice Hall, ISBN 0-13-083443-2, 1999.
[20] The NLNR web page is not present now. The same traffic trace is available at the Waikato Internet Traffic Storage project at http://www.wand.net.nz/wits/
[21] K. Cho, K. Mitsuya, and A. Kato, "Traffic data repository at the WIDE project," Proc. USENIX Annual Technical Conference, 2000.
[22] http://mawi.wide.ad.jp/mawi/
[23] H. Ito, H. Hasegawa, and K. Sato, "Effect of path granularity on dynamic path network control," IEICE Gen. Conf. 2009, B-7-80, March 2009. (in Japanese)
[24] "Network performance objectives for IP-based services," ITU-T Recommendation Y.1541, Feb. 2006.
[25] OPNET Modeller, OPNET Technologies Inc. Available from http://www.opnet.com

Hiroyuki Ito received the B.E. and M.E. degrees in electrical and electronic engineering and information engineering from Nagoya University, Nagoya, Japan in 2009 and 2011, respectively. In 2011, he joined Fujitsu Co., Inc. His major interests include dynamic network control for energy saving.

Hiroshi Hasegawa received the B.E., M.E., and D.E. degrees all in electrical and electronic engineering from Tokyo Institute of Technology, Tokyo, Japan, in 1995, 1997, and 2000, respectively. From 2000 to 2005, he was an assistant professor of the department of communications and integrated systems, Tokyo Institute of Technology. Currently he is an associate professor of Nagoya University. His current research interests include photonic networks, multidimensional digital signal processing and time-frequency analysis. Dr. Hasegawa received the Young Researcher Awards from SITA (Society of Information Theory and its Applications) and the IEICE (Institute of Electronics, Information and Communication Engineers) in 2003 and 2005, respectively. Dr. Hasegawa is a member of IEEE.

Ken-ichi Sato is currently a Professor at the Graduate School of Engineering, Nagoya University, and he is an NTT Research and Development Fellow. Before joining the university in April 2004, he was an executive manager of the Photonic Transport Network Laboratory at NTT. His R&D activities cover future transport network architectures, network design, OA&M (operation administration and maintenance) systems, photonic network systems including optical cross-connect/ADM and photonic IP routers, and optical transmission technologies. He received the B.S., M.S., and Ph.D. degrees in electronics engineering from the University of Tokyo, in 1976, 1978, and 1986, respectively. He has authored/co-authored more than 250 research publications in international journals and conferences. He holds 35 granted patents and more than 100 pending patents. He received the Young Engineer Award in 1984, the Excellent Paper Award in 1991, and the Achievement Award in 2000 from the Institute of Electronics, Information and Communication Engineers (IEICE) of Japan, and the Best Paper Awards in 2007 and 2008 from the IEICE Communications Society. He was also the recipient of the distinguished achievement Award of the Ministry of Education, Science and Culture in 2002. His contributions to asynchronous transfer mode (ATM) and optical network technology development extend to co-editing the IEEE Journal on Selected Areas in Communications (three special issues) and the IEEE Journal on Lightwave Technology (special issue), organizing several workshops and conference technical sessions, serving on numerous committees of international conferences including the OFC and ECOC, authoring a book, Advances in Transport Network Technologies (Artech House, 1996), and co-authoring thirteen other books.