

PAPER

Background TCP data transfer with inline network measurement

Tomoaki TSUGAWA^{†a)}, Go HASEGAWA^{††b)}, and Masayuki MURATA^{†c)}, *Members*

SUMMARY In the present paper, ImTCP-bg, a new background TCP data transfer mechanism that uses an inline network measurement technique, is proposed. ImTCP-bg sets the upper limit of the congestion window size of the sender TCP based on the results of the inline network measurement, which measures the available bandwidth of the network path between the sender and receiver hosts. ImTCP-bg can provide background data transfer without affecting the foreground traffic, whereas previous methods cannot avoid network congestion. ImTCP-bg also employs an enhanced RTT-based mechanism so that ImTCP-bg can detect and resolve network congestion, even when reliable measurement results cannot be obtained. The performance of ImTCP-bg is investigated through simulations, and the effectiveness of ImTCP-bg in terms of the degree of interference with foreground traffic and the link bandwidth utilization is also investigated.

key words: background data transfer, inline network measurement, congestion control, available bandwidth, TCP (Transmission Control Protocol)

1. Introduction

Due to the rapid development of networking technologies in both access and core computer networks, as well as the sudden increase of the Internet population, various IP-based network services are emerging and currently co-exist on the Internet. Although these services compete for network link bandwidth, Transmission Control Protocol (TCP) [1] currently plays a major and important role for avoiding and solving network congestion collapse by using the congestion control algorithm [2] between the sender/receiver endhosts. Thus, TCP provides effective usage and fair sharing of network resources among competing data transmission flows.

However, some of the Internet services do not necessarily require fair resource allocation with respect to other flows and should be operated in the background through prioritization mechanisms. For example, in Content Delivery/Distribution Networks (CDNs) [3–5] such as Akamai [6], Web servers transfer various types of data (e.g., backup, caching [7], and prefetching [8,9]), in addition to the data in response to the document transfer request from Web clients. In this

case, the user-requested data should be transferred with the higher priority than the other traffic. Data backup and synchronization in Storage Area Networks (SAN), online updating of operating systems (e.g., Microsoft's Background Intelligent Transfer Service [10]), and data caching in peer-to-peer network [11, 12] are other examples of tasks that should be performed without affecting the foreground traffic.

In previous studies, such prioritized behaviors were realized by either IP-based mechanisms or application-based mechanisms. In IP-based mechanisms, such as DiffServ [13], the Internet routers are equipped with prioritization mechanisms and process the incoming packets according to pre-defined prioritization policies. For instance, Assured Forwarding [14] in DiffServ has four classes and three dropping levels at the router buffer to differentiate the incoming flows. However, such mechanisms have well-known shortcomings in scalability, because the prioritization mechanisms should be implemented on all routers between the sender and receiver endhosts.

In application-based mechanisms found in [15, 16], the prioritization mechanisms are provided by upper-level programs, or sometimes by service administrators. For example, cache synchronization and prefetching in CDNs is performed when there is little user-requested foreground traffic. Data backup is usually scheduled to be performed at midnight in order to avoid degrading the throughput of other higher-prioritized flows during the daytime. In such cases, the programs/administrators must monitor the network traffic to determine the time during which the network is underutilized. However, successfully realizing such mechanisms is difficult due to large fluctuations in Internet traffic. Moreover, since each application deploys a prioritization mechanism according to its service requirements, we cannot estimate its performance especially when multiple applications with different prioritization mechanisms co-exist in the network.

Therefore, TCP-based approaches such as TCP Nice [17] and TCP-LP [18], which are herein referred to as *background TCP*, have been introduced in order to handle background (lower-prioritized) data transfer on the Internet. These approaches observe the Round Trip Times (RTTs) of the data packets of a TCP connection and decrease the congestion window size when the RTTs increase, whereas the original TCP Reno contin-

[†]Graduate School of Information Science and Technology, Osaka University

^{††}Cybermedia Center, Osaka University

a) E-mail: t-tugawa@ist.osaka-u.ac.jp

b) E-mail: hasegawa@cmc.osaka-u.ac.jp

c) E-mail: murata@ist.osaka-u.ac.jp

ues to increase its congestion window size until packet loss occurs, regardless of increases in RTTs. Generally, on the network congestion in the actual networks, the RTT increases before the packet loss occurs since the incoming packets to the bottleneck router are first stored at the buffer, and when the buffer is fully utilized, the incoming packets are discarded. TCP Nice and TCP-LP utilize this characteristic and detect the network congestion earlier than the TCP Reno by observing the change of RTT.

Although both TCP Nice and TCP-LP can realize data transfer without affecting other higher-prioritized traffic, these protocols are unable to utilize the available bandwidth of the network efficiently. This is because the degree to which the congestion window size can decrease when the RTTs increase is fixed, and is too large, regardless of the network condition, similarly to TCP Reno which halves the window size when packet loss occurs.

In the present paper, a novel background TCP mechanism based on bandwidth measurement is proposed, with the goal of achieving both background transfer and network bandwidth utilization. The proposed background TCP variant uses the inline network measurement mechanism proposed in [19, 20], which can measure the available bandwidth of the network path between sender and receiver endhosts. This inline network measurement mechanism uses the data and ACK packets of a TCP connection for the measurement task, without injecting additional traffic, which is ideal for background data transfer. The proposed mechanism sets the maximum value of the congestion window size of the sender TCP by using the measurement results of the available bandwidth. In addition, an RTT-based mechanism that dynamically determines the degree to which the congestion window size can decrease according to the observed RTT value is employed, whereas TCP Nice and TCP-LP use a constant degree for the possible decrease.

The remainder of this paper is organized as follows. Section 2 describes the purpose of background TCP data transfer and presents a discussion of the problems of existing mechanisms. In Section 3, ImTCP-bg, a new background TCP data transfer mechanism that uses an inline network measurement technique, is proposed and its characteristics are discussed. In Section 4 presents simulation results that are used to evaluate the performance of ImTCP-bg. Finally, in Section 5, conclusions are presented and areas for future study are outlined.

2. Background data transfer with TCP

TCP adjusts the data transmission speed by changing the congestion window size in response to network congestion. The TCP algorithm allows a TCP sender to continue to increase its congestion window size additively until network congestion is detected. TCP de-

creases the window size multiplicatively when network congestion occurs. As an indicator of the network congestion, TCP Reno uses packet losses in the network (referred to herein as a *loss-based mechanism*). On the other hand, TCP Nice and TCP-LP introduce another congestion indicator, namely the increase of RTTs for data packets (*RTT-based mechanism*). These protocols provide background data transfer without affecting the foreground traffic based on the following assumption. Consider an output link of an Internet router equipped with an output buffer. When the packet incoming rate of the traffic destined for the output link is larger than the output link bandwidth, the excess traffic is stored in the output buffer, which causes some queuing delay, and eventually results in the packet losses when the buffer becomes full. That is, for a TCP connection, the RTTs usually increase before packet losses occur when the network is congested. Therefore, TCP Nice and TCP-LP connections can detect network congestion earlier than TCP Reno connections.

The present paper considers the following two objectives for background data transfer:

1. no adverse effect on the foreground traffic
2. full utilization of the network link bandwidth

That is, a perfect background data transfer mechanism can fully utilize the bandwidth that is unused by the foreground traffic, while not degrading the performance of the foreground traffic. However, realizing such a complete mechanism is quite difficult because a trade-off relationship exists between these two objectives. The difficulty in realizing a good background data transfer mechanism lies in balancing this trade-off relationship. For example, TCP Nice and TCP-LP are unable to efficiently utilize the available bandwidth of the network path, especially when the number of background TCP connections is small [17, 18]. This is mainly because these protocols use fixed parameters in detecting network congestion and decreasing the congestion window size. That is, these two background TCP variants set the parameters by which to satisfy objective (1), while sacrificing objective (2). Opposite to this, when these TCP variants change the controlled parameters in order to improve the utilization of link bandwidth, i.e., satisfying the objective (2), it is obvious that the change causes the increase in the degree of interference with other traffic. In addition, the optimal parameter setting to the trade-off relationship depends on the network condition strongly. Therefore, RTT-based mechanisms such as those of TCP Nice and TCP-LP cannot completely solve this trade-off problem due to their trial-and-error nature. These mechanisms continue to increase the window size until the value of RTT reaches its threshold, that is, until the indication of network congestion appears, and then decrease the window size to some degree in order to avoid the congestion.

In order to satisfy the above two objectives, another congestion indicator is proposed, i.e., the available bandwidth of the network path between the sender and receiver hosts (*bandwidth-based mechanism*). An available bandwidth means the residual bandwidth, that is, the bandwidth which no other connections are currently using. The available bandwidth is the most straightforward information by which to describe background data transfer. If the TCP sender obtains the available bandwidth information exactly and quickly, then an ideal background data transfer mechanism, in terms of both the background nature and link utilization can be created.

Many algorithms and tools by which to measure the available bandwidth of network paths have been proposed in the literature [21–25]. However, the existing methods cannot be directly employed for the newly proposed background TCP because these methods utilize numerous test probe packets and require too much time to obtain a single measurement result. For instance, PathLoad [23] sends several 100-packet measurement streams for a measurement. PathChirp [25] is a modification of PathLoad for the purpose of decreasing the number of probe packets. However, the required number of packets to be sent at one time in PathChirp is still large. In order to address this problem, a method referred to as Inline measurement TCP (ImTCP) has been proposed in [19, 20]. ImTCP does not inject extra traffic into the network, but rather estimates the available bandwidth of the network path from data and ACK packets transmitted by an active TCP connection in an inline fashion. Since the ImTCP sender obtains bandwidth information every 1–4 RTTs, ImTCP can follow the traffic fluctuation of the underlying IP network well. In addition, because the ImTCP mechanism is implemented at the bottom of the TCP layer, various types of TCP congestion control mechanisms can include this measurement mechanism. Therefore, the ImTCP mechanism is integrated into the proposed background TCP in order to obtain the available bandwidth information of the network path.

However, the RTT-based mechanism cannot be discarded even when the bandwidth-based mechanism is employed, because ImTCP does not always provide accurate measurement results for the available bandwidth. For example, when the congestion window size of the ImTCP sender is small, ImTCP cannot measure the available bandwidth. Furthermore, the measurement accuracy of ImTCP depends on the network environment, e.g., the RTT, the physical link bandwidth, and the number of active connections. When the measured available bandwidth value is inaccurate, the background data transfer based on the measured value may affect the foreground traffic. Therefore, the RTT-based mechanism should be used in conjunction with the bandwidth-based mechanism.

In the next section, the mechanism of ImTCP-bg,

a new background TCP data transfer mechanism based on the available bandwidth measurement technique of ImTCP is described. The proposed mechanism also employs an enhanced RTT-based mechanism, which dynamically determines the control parameters.

3. ImTCP-bg: ImTCP background mode

Basically, in this paper we do not care which traffic should be transferred by background TCP mechanism. That is, we do not introduce a mechanism that automatically selects the transfer mode (normal or background) for traffic from upper-layer applications. Instead, we prepare the interface to select the transfer mode by socket option. Therefore, each application should select the transfer mode when it establishes a new TCP connection.

ImTCP-bg consists of two major mechanisms: a bandwidth-based mechanism with inline network measurement and an enhanced RTT-based mechanism for adjusting the window size when the first mechanism does not work well. In this section, we first introduce the outline of the inline network measurement technique, ImTCP, and then describe the detail of each proposed mechanism. The ImTCP algorithm is described in detail in [19, 20].

3.1 Outline of ImTCP algorithm

ImTCP measures the available bandwidth of the network path between sender and receiver hosts. In TCP data transfer the sender host transfers a data packet and the receiver host replies an ACK packet against the data packet. ImTCP measures the available bandwidth by using this nature. Concretely speaking, ImTCP adjusts the interval of data packets according to the measurement algorithm, and then calculates the available bandwidth by observing the change of ACK intervals.

During every measurement, ImTCP searches for the available bandwidth only within a given search range. The search range is a range of bandwidth that is expected to include the current available bandwidth. By introducing the search range, ImTCP can avoid sending probe packets at an extremely high rate. ImTCP can also keep the number of probe packets for the measurement quite small.

ImTCP has some cases when the measurement result is unavailable or unreliable. According to [20], when the current window size is smaller than the number of packets required for a measurement, ImTCP does not measure the available bandwidth. In addition, when the other traffic send data packets burstly, the intervals of ImTCP data packets are disturbed by the bursty traffic and it makes the measurement result inaccurate. Although ImTCP give the measurement result in this case, the result is not reliable.

3.2 ImTCP-bg mechanisms

3.2.1 Bandwidth-based mechanism

In ImTCP-bg, the congestion window size is controlled using the available bandwidth information of the network path between the sender and receiver hosts, as measured by the ImTCP mechanism. ImTCP-bg smoothes the measurement results of ImTCP using a simple exponential weighted moving average, as follows:

$$\bar{A} \leftarrow (1 - \gamma) \cdot \bar{A} + \gamma \cdot A_{cur} \quad (1)$$

where A_{cur} denotes the current result of the available bandwidth measured by ImTCP mechanism, γ is a smoothing parameter, and \bar{A} is the smoothed available bandwidth. Note that the appropriate setting of γ will be described in Section 4. The ImTCP-bg sender then sets the upper limit of the congestion window size ($maxcwnd$) using the following equation:

$$maxcwnd \leftarrow \bar{A} \cdot RTT_{min} \quad (2)$$

where RTT_{min} is the minimum RTT experienced throughout the lifetime of the connection.

As shown above, the proposed bandwidth-based mechanism is quite simple: the upper-limit of the congestion window size is simply set as the product of the measured available bandwidth and the minimum RTT. In Section 4, this simple mechanism is demonstrated to be effective for background data transfer.

3.2.2 Enhanced RTT-based mechanism

The effectiveness of the above-described bandwidth-based mechanism depends largely on the accuracy of the measurement by ImTCP of the available bandwidth. In [20] the authors demonstrated that ImTCP can give the reasonably accurate measurement results every 1–4 RTTs. However, ImTCP does not always provide reliable measurement results, as explained in Section 3.1, and may result in the congestion of the bottleneck link. Therefore, the RTT-based mechanism is employed to quickly detect and resolve the undesirable network congestion. The RTT-based mechanism is enhanced in order to be used with the bandwidth-based mechanism.

ImTCP-bg detects network congestion using only the current and minimum values of RTT, whereas TCP Nice and TCP-LP also use the maximum RTT, which is difficult to observe in the actual network. When an increase in RTT is detected, the ImTCP-bg sender decreases its congestion window size immediately in order to resolve the congestion. Next, denote \overline{RTT} as the smoothed RTT value that is calculated by the traditional TCP mechanism and RTT_{min} as the minimum RTT. Here, δ (> 1.0) is the threshold parameter to

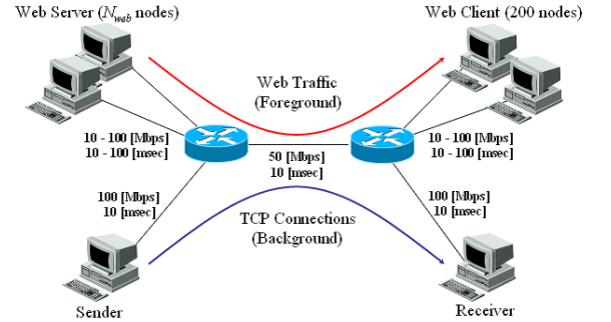


Fig. 1 Network model

judge whether network congestion occurs. The ImTCP-bg sender detects the network congestion when the following condition is satisfied:

$$\frac{\overline{RTT}}{RTT_{min}} > \delta \quad (3)$$

When Equation (3) is satisfied, it means that the queuing delay occurs at the bottleneck router. Since we treat the increase of queuing delay as the indication of network congestion, ImTCP-bg decreases its congestion window size according to the following equation in order not to affect the foreground traffic.

$$cwnd \leftarrow cwnd \cdot \frac{RTT_{min}}{\overline{RTT}} \quad (4)$$

Here, $cwnd$ is the current congestion window size, \overline{RTT} is the smoothed RTT value and RTT_{min} is the minimum RTT. Equation (4) implies that ImTCP-bg determines the degree of decrease of the congestion window size based on the ratio of the current value of RTT and its minimum value. Thereby, ImTCP-bg avoids unnecessary the underutilization of the link bandwidth while maintaining the background-based data transfer. Note that this modification of the RTT-based mechanism is effective because the bandwidth-based mechanism is used concurrently.

4. Performance evaluations

In this section, simulation results are used to evaluate the performance of ImTCP-bg, as proposed in Section 3, and ns-2 [26] is used for the simulations. Traditional TCP Reno, TCP Nice and TCP-LP were chosen for performance comparisons. Issues in parameter settings of ImTCP-bg are first described in Subsection 4.1, and then the performance of ImTCP-bg is compared to that of TCP Nice and TCP-LP in Subsection 4.2.

The network model used in the simulation is depicted in Figure 1. This model consists of sender/receiver hosts, two routers, and links between the hosts and routers. The bandwidth of the bottleneck link is set to 50 Mbps, and the propagation delay is 10 msec. A DropTail discipline is deployed at the

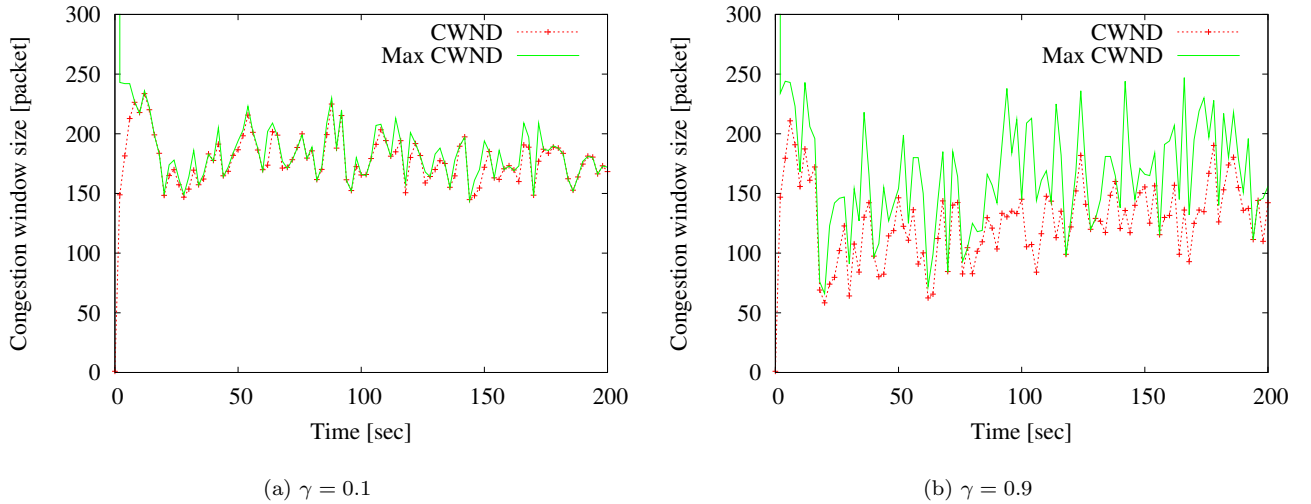


Fig. 2 Change of congestion window size and its upper limit

router buffer, and the buffer size is set to 1000 packets. The packet size is 1500 Bytes. Web traffic is assumed to be foreground traffic. N_{web} Web servers transfer Web documents to 200 Web clients. The bandwidth of the access link of each Web node is set randomly between 10 and 100 Mbps, and the propagation delay is also a random value between 10 and 100 msec. The amount of foreground Web traffic is adjusted by changing N_{web} . In addition, one or more TCP connections are established in order to perform background data transfer. The performance of the background TCP variants are compared with respect to the following: the transfer time of the foreground Web traffic, the queue length of the bottleneck link buffer, the throughput of the background data transfer, and utilization of the available bandwidth. The control parameters for TCP Nice and TCP-LP are configured according to [17, 18].

4.1 Parameter settings

4.1.1 γ setting

First, the effect of the smoothing parameter of the measured available bandwidth of Equation (2), which is denoted as γ ($0 < \gamma < 1$), is considered. If γ is set to a larger value, then measurement errors have a greater influence on ImTCP-bg, and changes in the network environment can be detected more rapidly. On the other hand, if γ is set to a smaller value, the ability to detect changes in the available bandwidth may be degraded, but the instantaneous measurement error can be filtered out. Therefore, γ should be set in a way such that these two performance aspects are balanced.

Figure 2 shows the changes in the congestion window size and its upper limit, tuned by ImTCP-bg, as a function of time, when γ is set to 0.1 (Figure 2(a)) and




Fig. 3 Effect of parameter γ

0.9 (Figure 2(b)). Figure 3 shows the average utilization of the available bandwidth as a function of γ . Here, N_{web} is set to 20 and δ is set to 1.2. These figures indicate that when the value of γ is large, the congestion window size tends to be small. ImTCP-bg increases its congestion window size additively until the congestion size reaches its upper limit. Opposite to this, when the congestion window size exceeds its upper limit, ImTCP-bg decreases the congestion window size to the upper limit immediately. Therefore ImTCP-bg cannot utilize the available bandwidth effectively when the value of γ is large. Based on the above considerations, γ is set to $\frac{1}{8}$ in the following simulations. This value is small enough to utilize the available bandwidth effectively and is typically used for calculating the smoothing RTT for TCP. However, we cannot conclude that this value of γ is always appropriate. In order to evaluate the effect of γ in detail, we need experiments in actual network envi-

ronments. We will consider the issue in the parameter setting as a future research.

4.1.2 RTT threshold δ

Next, the RTT threshold δ ($\delta \geq 1.0$) of Equation (3) is considered. In order to determine the appropriate value, δ was changed to various values and a number of simulations were conducted. Figure 4 shows the changes in the available bandwidth utilization and average queue length at the bottleneck link buffer as functions of δ , where γ is set to $\frac{1}{8}$ according to the results in Subsection 4.1.1, N_{web} is set to 20 and the number of ImTCP-bg connections is set to 1, 2, 5, and 10.

Figure 4 shows that if δ is set to a very small value, then ImTCP-bg cannot utilize the available bandwidth effectively, even though the queue length is small. This is because when ImTCP-bg uses small δ , the instantaneous measurement error is sensed and the congestion window size is frequently decreased. Note that the larger the average queue length, the larger the affect on the foreground traffic. Therefore, δ should be set so as to balance the degree of interference with the foreground traffic and the utilization of available bandwidth. From the extensive simulation results, including those shown in Figure 4, $\delta = 1.2$ is determined to be a good selection in order to balance these two requirements.

4.2 Performance comparisons

4.2.1 Case of one connection

First, the simulation results are presented for the case in which one background data transfer connection is established, and the degree of interference with the foreground traffic and the utilization of network bandwidth are evaluated. The number of Web servers, N_{web} , is changed from 10 to 50. Figure 5 shows the change in the average download time of foreground Web documents and the average throughput of the background TCP connection. The results labeled as “available bandwidth” in Figure 5(a) and “no background traffic” in Figure 5(b) show the results for the case in which no background data transfer exists.

Figure 5(a) shows that although the TCP Reno connection achieves the highest throughput, the throughput exceeds the available bandwidth. Furthermore, Figure 5(b) shows that the average download time of the foreground Web documents is larger than for the case in which no background traffic exists. That is, TCP Reno cannot be used for background data transfer. On the other hand, for TCP Nice, TCP-LP and ImTCP-bg, the average download time in Figure 5(b) is almost identical to the case of no background traffic. This means that these protocols do not affect the foreground Web traffic, satisfying one of the

objectives of the background data transfer. Furthermore, Figure 5(a) shows that the average throughput of the ImTCP-bg connection is the closest to the available bandwidth. Therefore, ImTCP-bg is determined to have the most ideal characteristics for background data transfer, which satisfies objectives (1) and (2) in Section 2.


4.2.2 Case of multiple connections

Finally, the results for the case in which two or more background data transfer connections are established are shown, and the effect of multiple background TCP connections is evaluated. In the first simulation, five background TCP connections join the network at 0, 50, 100, 150, and 200 seconds, and end data transmissions at 500, 450, 400, 350, and 300 seconds. This means that the number of active background TCP connections in the network is as follows: 1, 2, 3, 4, 5, 4, 3, 2, and 1. Table 1 shows the average queue length at the output link buffer of the bottleneck router and Figure 6 shows the throughput of background data transfer connection as functions of time. Here, N_{web} is set to 20.

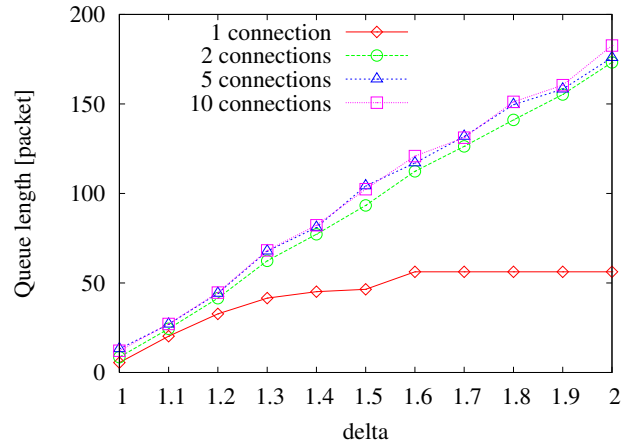
TCP Reno shows the worst behavior for the background data transfer, in terms of large queue length at the bottleneck link (Table 1) and over-utilization of the available bandwidth of the network (Figure 6(a)). Table 1 also shows that the average of queue length of TCP Nice is the smallest among the four variants, meaning that TCP Nice is the best choice for satisfying objective (1), described in Section 2. However, Figure 6(b) shows that the throughput of the background data transfer is the lowest among the four variants, especially when the number of connection is small.

Figure 6(c) shows that when TCP-LP is used for background data transfer, packet losses occur immediately after a new connection is established. This is because TCP-LP needs the maximum RTT to control its congestion window size. TCP Nice and TCP-LP detect network congestion by using the minimum and the maximum RTT (or one-way packet delay). However, essentially, monitoring the maximum RTT by background TCP is difficult because these TCP variants decrease the congestion window size at an early stage of network congestion. Therefore TCP-LP intentionally continues to increase its congestion window size until packet losses occur at the initial slow start phase to determine the maximum RTT value. Consequently, the TCP-LP connection cannot avoid to occur packet losses in the beginning of the data transfer. Furthermore, in Figure 6(c) we can see that the throughput of TCP-LP connections are shown to be quite low for some time after the packet loss. This is because the fast retransmission and fast recovery mechanism of TCP Reno is activated.

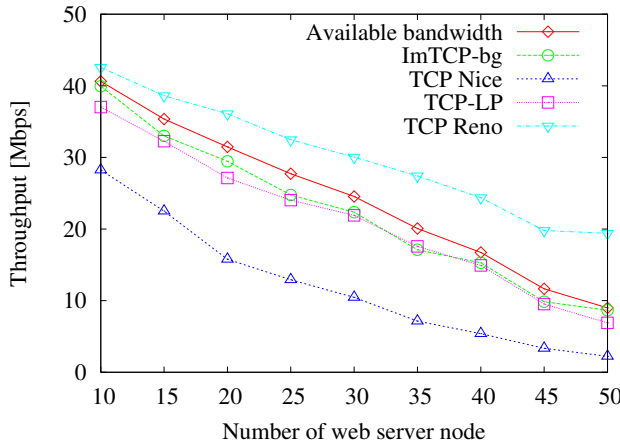
Figure 6(d) shows that ImTCP-bg connections can utilize the available bandwidth of network path effec-



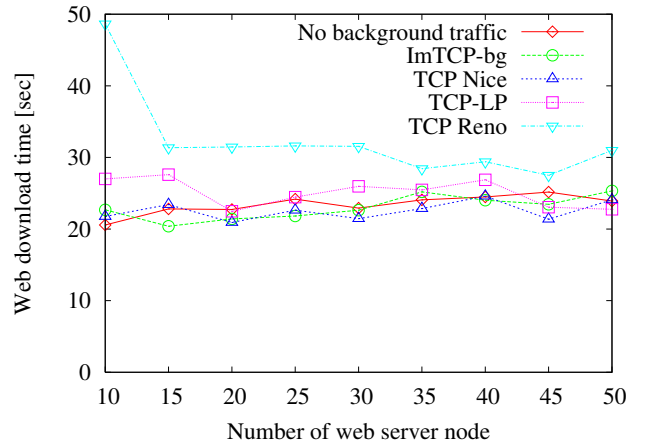
(a) Average utilization of available bandwidth



(b) Average queue length

Fig. 4 Effect of parameter δ


(a) Average of throughput



(b) Average of download time

Fig. 5 Results of one connection case

Table 1 Average of queue length in the case of five connections

	TCP Reno	TCP Nice	TCP-LP	ImTCP-bg
Average queue length [packets]	583.34	8.44	64.63	44.11

tively even when only one connection exists. This is because ImTCP-bg controls its congestion window size appropriately using the results of inline network measurement. Furthermore, when multiple connections exist in the network, ImTCP-bg connections can maintain high utilization of the available bandwidth and the change in the throughput of each ImTCP-bg connection is stable compared with other background TCP variants. That is because ImTCP-bg dynamically changes the degree of decrease congestion window size according

to the change in the RTT. From these simulation results, the bandwidth-based algorithm with inline measurement and the RTT-based algorithm are determined to co-exist well in ImTCP-bg to realize background data transfer.

We then check the utilization of the available bandwidth in more detail. In the next simulation, K background TCP connections (C_1, C_2, \dots, C_K) exist in the network. C_1 joins the network at 0 second and ends data transmission at 500 seconds. C_i ($i = 2, 3, \dots, K$)




Fig. 6 Change of throughput with five connections

starts at $\frac{200}{i-1}$ seconds and ends at $500 - \frac{200}{i-1}$ seconds. This means that the number of active background TCP connections in the network changes as time passes similarly to the previous simulation (Figure 6). Figure 7 shows the change of the utilization of the available bandwidth as a function of the available bandwidth when the number of background TCP connections is 1 (Figure 7(a)), 2 (Figure 7(b)), 5 (Figure 7(c)), and 10 (Figure 7(d)). Here, available bandwidth is changed by changing the number of servers, N_{web} .

Figure 7 shows that the bandwidth utilization of TCP Reno connections greatly exceeds the available bandwidth, meaning that it makes severe effect on the foreground traffic. When using TCP Nice, the less the available bandwidth is, the less the utilization of available bandwidth becomes. This is because when the available bandwidth is small, TCP Nice frequently de-

fects the network congestion and decreases the congestion window size in order not to affect the foreground traffic. For TCP-LP, the utilization of the available bandwidth exceeds 1.0 when the available bandwidth is small. This means that it does not satisfy the purpose of background transfer, that is, TCP-LP affects the foreground traffic. Furthermore, when the available bandwidth is large and many background TCP connections exist, the utilization of the available bandwidth becomes small. This is because packet losses occur after the new TCP-LP connection is established, as Figure 6(c) shows. On the other hand, ImTCP-bg shows the best behavior regardless of the available bandwidth and the number of background TCP connections. Therefore, we conclude that ImTCP-bg can utilize effectively the available bandwidth in any cases.




Fig. 7 Utilization of Available bandwidth with multiple connections

5. Conclusions

In the present paper, ImTCP-bg, a new background TCP data transfer mechanism that uses an inline network measurement technique, was proposed. ImTCP-bg provides a background data transfer without interfering with the foreground traffic by setting the upper limit of its congestion window size based on the results of the inline network measurement. ImTCP-bg also employs an enhanced RTT-based mechanism, which dynamically determines the control parameters. ImTCP-bg can detect and resolve network congestion even when reliable measurement results cannot be obtained. Through simulation evaluations, the effectiveness of ImTCP-bg in terms of the degree of interference with foreground traffic and utilization of the available bandwidth was confirmed. In future studies, ImTCP-bg will be imple-

mented and its performance will be evaluated in an actual network. Besides to this, we will consider about the parameter settings (smoothing parameter γ and RTT threshold δ) by using results of actual network experiments.

References

- [1] J.B. Postel, "Transmission control protocol," RFC 793, Sept. 1981.
- [2] W.R. Stevens, TCP/IP Illustrated, Volume 1: The Protocols, Addison-Wesley, 1994.
- [3] B. Krishnamurthy, C. Wills, and Y. Zhang, "On the use and performance of content distribution networks," Proceedings of ACM SIGCOMM 2001 Internet Measurement Workshop, Nov. 2001.
- [4] S. Saroiu, K. Gummadi, R.J. Dunn, S.D. Gribble, and H.M. Levy, "An analysis of Internet content delivery systems," Proceedings of 5th Symposium on Operating Systems Design and Implementation (OSDI 2002), Dec. 2002.

- [5] G. Pierre and M. van Steen, "Design and implementation of usercentered content delivery network," Proceedings of 3rd IEEE Workshop on Internet Applications, June 2003.
- [6] Akamai Home Page. available at <http://www.akamai.com/>.
- [7] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and Zipf-like distributions: Evidence and implications," Proceedings of IEEE INFOCOM 1999, March 1999.
- [8] M. Crovella and P. Barford, "The network effects of prefetching," Proceedings of IEEE INFOCOM 1998, March 1998.
- [9] A. Venkataramani, P. Yalagandula, R. Kokku, S. Sharif, and M. Dahlin, "The potential costs and benefits of long term prefetching for content distribution," Computer Communication Journal, vol.25, no.4, pp.367-375, March 2002.
- [10] Microsoft Corporation, Background Intelligent Transfer Service in Windows Server 2003, Sept. 2002. available at <http://www.microsoft.com/windowsserver2003/techinfo/overview/bits.msp>.
- [11] F. Dabek, M.F. Kaashoek, D. Karger, R. Morris, and I. Stoica, "Wide-area cooperative storage with CFS," Proceedings of 18th Symposium on Operating Systems Principles (SOSP 2001), Oct. 2001.
- [12] A. Rowstron and P. Druschel, "Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility," Proceedings of 18th Symposium on Operating Systems Principles (SOSP 2001), Oct. 2001.
- [13] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services," RFC 2475, Dec. 1998.
- [14] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group," RFC 2597, June 1999.
- [15] R. Kokku, P. Yalagandula, A. Venkataramani, and M. Dahlin, "NPS: A non-interfering deployable web prefetching system," Proceedings of 4th USENIX Symposium on Internet Technologies and Systems, March 2003.
- [16] P. Key, L. Massoulie, and B. Wang, "Emulating low-priority transport at the application layer: A background transfer service," ACM SIGMETRICS Performance Evaluation Review, vol.32, pp.118-129, June 2004.
- [17] A. Venkataramani, R. Kokku, and M. Dahlin, "TCP Nice: A mechanism for background transfers," Proceedings of 5th Symposium on Operating Systems Design and Implementation (OSDI 2002), Dec. 2002.
- [18] A. Kuzmanovic and E.W. Knightly, "TCP-LP: A distributed algorithm for low priority data transfer," Proceedings of IEEE INFOCOM 2003, April 2003.
- [19] C.L.T. Man, G. Hasegawa, and M. Murata, "A new available bandwidth measurement technique for service overlay networks," Proceedings of 6th IFIP/IEEE MMNS 2003 E2EMON Workshop, Sept. 2003.
- [20] C.L.T. Man, G. Hasegawa, and M. Murata, "Available bandwidth measurement via TCP connection," Proceedings of 7th IFIP/IEEE MMNS 2004 E2EMON Workshop, Oct. 2004.
- [21] R.L. Carter and M.E. Crovella, "Measuring bottleneck link speed in packet-switched networks," International Journal on Performance Evaluation, vol.27-28, pp.297-318, Oct. 1996.
- [22] B. Melander, M. Bjorkman, and P. Gunningberg, "A new end-to-end probing and analysis method for estimating bandwidth bottlenecks," Proceedings of IEEE GLOBECOM 2000, Nov. 2000.
- [23] M. Jain and C. Dovrolis, "End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput," Proceedings of ACM SIGCOMM 2002, Aug. 2002.
- [24] C. Dovrolis and D. Moore, "What do packet dispersion techniques measure?," Proceedings of IEEE INFOCOM 2001, April 2001.
- [25] V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell, "pathChirp: Efficient available bandwidth estimation for network paths," Proceedings of Passive and Active Measurement Workshop (PAM 2003), April 2003.
- [26] The VINT Project, "UCB/LBNL/VINT network simulator - ns (version 2)." available at <http://www.isi.edu/nsnam/ns/>.



Tomoaki TSUGAWA is now a M.E. candidate at Graduate School of Information Science and Technology, Osaka University. His research work is in the area of transport architecture for future high-speed networks. He is a member of IEICE.



Go HASEGAWA received the M.E. and D.E. degrees in Information and Computer Sciences from Osaka University, Osaka, Japan, in 1997 and 2000, respectively. From July 1997 to June 2000, he was a Research Assistant of Graduate School of Economics, Osaka University. He is now an Associate Professor of Cybermedia Center, Osaka University. His research work is in the area of transport architecture for future high-speed networks. He is a member of the IEEE and IEICE.



Masayuki MURATA received the M.E. and D.E. degrees in Information and Computer Sciences from Osaka University, Japan, in 1984 and 1988, respectively. In April 1984, he joined Tokyo Research Laboratory, IBM Japan, as a Researcher. From September 1987 to January 1989, he was an Assistant Professor with Computation Center, Osaka University. In February 1989, he moved to the Department of Information and Computer Sciences, Faculty of Engineering Science, Osaka University. From 1992 to 1999, he was an Associate Professor in the Graduate School of Engineering Science, Osaka University, and from April 1999, he has been a Professor of Osaka University. He moved to Advanced Networked Environment Division, Cybermedia Center, Osaka University in 2000, and moved to Graduate School of Information Science and Technology, Osaka University in April 2004. He has more than two hundred papers of international and domestic journals and conferences. His research interests include computer communication networks, performance modeling and evaluation. He is a member of IEEE, ACM, The Internet Society, IEICE and IPSJ.