

A Fair and Efficient Congestion Avoidance Scheme Based on the Minority Game

Hiroshi Kutsuna* and Satoshi Fujita*

Abstract—In this paper, we propose a new congestion control scheme for high-speed networks. The basic idea of our proposed scheme is to adopt a game theory called, “Minority Game” (MG), to realize a selective reduction of the transmission speed of senders. More concretely, upon detecting any congestion, the scheme starts a game among all senders who are participating in the communication. The losers of the game reduce the transmission speed by a multiplicative factor. MG is a game that has recently attracted considerable attention, and it is known to have a remarkable property so that the number of winners converges to a half the number of players in spite of the selfish behavior of the players to increase its own profit. By using this property of MG, we can realize a fair reduction of the transmission speed, which is more efficient than the previous schemes in which all senders uniformly reduce their transmission speed. The effect of the proposed scheme is evaluated by simulation. The result of simulations indicates that the proposed scheme certainly realizes a selective reduction of the transmission speed. It is sufficiently fair compared to other simple randomized schemes and is sufficiently efficient compared to other conventional schemes.

Keywords—Congestion Control, AIMD, Minority Game

1. INTRODUCTION

According to the recent advancement of network technologies, (ultra) high-speed networks have become available in many places (e.g., IEEE802.3ae has adopted 10Gigabit Ethernet (10GbE) as a standard specification for high-speed networks and InfiniBand has been widely adopted as a high-speed network interface for attaining a transmission speed of 30 Gbps [1]). Such high-speed networks could provide a broad communication bandwidth to many users, which opens up several new application fields such as on-demand video-streaming of HDV (High Definition Video) quality and remote medical treatment by sharing three-dimensional ultrasonic wave scanner video among distant medical institutions. In order to fully utilize such a broad bandwidth of high-speed networks, we have to control the traffic flow issued by each sender in such a way that the packets are transmitted with a *nearly maximum* transmission speed, while bounding the occurrence of congestions as much as possible.

The current version of TCP adopts AIMD (Additive Increase/Multiplicative Decrease)[3] as a congestion control scheme. However, in this scheme, it is difficult to fully utilize the broad bandwidth of high-speed networks for the following reasons [10]: in AIMD, upon detecting congestion, *all* senders reduce their transmission speed to *half* in order to release shared re-

Manuscript received April 5, 2011; accepted June 14, 2011.

Corresponding Author: Satoshi Fujita

* Dept. of Information Engineering, Hiroshima University, Higashi-Hiroshima, 739-8527, Japan (fujita@se.hiroshima-u.ac.jp)

sources for future use (e.g., multiplicative decrease). Then, all senders will share this released resource again by gradually increasing the transmission speed with the same increasing speed (e.g., additive increase). Thus, AIMD could realize a “fair” sharing of resources while utilizing at least half of the given communication bandwidth. However, in high-speed networks with a *broad* communication bandwidth, it takes a very long time before fully utilizing the resources that have been released by the senders. Although this kind drawback for AIMD could be resolved by tuning the multiplicative decreasing factors like Scalable TCP [11], High-Speed TCP [10], and Bimodal [8], they still follow the same idea as AIMD where *all* senders *uniformly* reduce their transmission speed with an adequate multiplicative decreasing factor.

In this paper, we propose a new congestion control scheme for high-speed communication networks, which attains a high communication throughput while keeping the fairness of the bandwidth assignment to the senders. The basic idea of the scheme is to adopt a game theory called “Minority Game” (MG, for short) to realize a *selective reduction* of the transmission speed. More concretely, all senders play a game upon detecting congestion, and the transmission speed of each sender is reduced according to the results of the game. MG is a game that has recently attracted considerable attention in the field of complex systems, and it is known to have the remarkable property where the number of winners decreases to half of the number of players in spite of the selfish behavior of the players to increase its own profit. In this paper, we propose two concrete schemes to associate game wins with the profits given to the senders. The basic idea of the proposed scheme is to adopt a *fair share* of the communication link [8] as a reference for the initial transmission speed after detecting any congestion. The performance of the proposed scheme is evaluated by simulation. The results from the simulations indicate that the proposed scheme significantly improves the throughput of AIMD. In addition, the proposed scheme certainly realizes a sufficiently fair reduction of the transmission speed, compared with other schemes including AIMD and simple randomized selection schemes.

The remainder of this paper is organized as follows: in Section 2, we overview the existing congestion control schemes in packet communication networks. Section 3 describes an outline of MG. Section 4 proposes a new congestion control scheme based on MG, whose performance is experimentally evaluated in Section 5. Finally, Section 6 concludes the paper and discusses future work.

2. CONGESTION CONTROL

2.1 TCP

In TCP, each sender detects an occurrence of congestion either by receiving a duplicate ACK, or by detecting a time out. In addition, each sender adaptively controls its transmission speed by changing the size of the congestion window, which is denoted as *cwnd* in the paragraphs below.

If TCP detects congestion by time out, it moves to the slow-start phase. In this phase, after setting the threshold to half of the current *cwnd*, it decreases *cwnd* to the minimum value, and then restarts the packet transmission. During the packet transmission, the sender (exponentially) increments *cwnd* until it reaches the threshold (e.g., it increments *cwnd* by one for each receipt of a correct ACK packet). Note that since the number of ACKs coincides with the number of transmitted packets, if it receives all ACKs for the packets transmitted when $cwnd = n$, the value of *cwnd* becomes $2n$ at the next transmission (e.g., it increases twice). After reaching the thresh-

old, the protocol moves to the congestion avoidance phase, and continues an additive increase of the window size as described below.

The current version of TCP adopts AIMD (Additive Increase / Multiplicative Decrease) as the congestion control scheme in the congestion avoidance phase [3]. In AIMD, upon detecting congestion by Duplicated ACK, each sender decreases $cwnd$ to $(1-\beta)cwnd$ by using the parameter of β , where in general, the value of β is set to $1/2$. It then (linearly) increases the transmission speed by using the additive increasing factor of α , as long as no additional congestions occur (e.g., after receiving a correct ACK, it increases $cwnd$ by $\alpha/cwnd$, where in general, the value of parameter α is set to one). Note that if it receives all ACKs for the packets transmitted when $cwnd = n$, the value of $cwnd$ becomes $n + \alpha$ at the next transmission.

2.2 Related Work

Although it could realize a fair sharing of resources, AIMD has a drawback in that it is difficult to fully utilize the broad bandwidth of high-speed networks, since all senders always reduce their transmission speed to half. In order to overcome the drawback of the original AIMD, there have been several proposed extensions for AIMD in the literature. In this subsection, we will overview those schemes in detail.

2.2.1 Scalable TCP

Scalable TCP is an extension of the TCP that has been proposed for high-speed networks [11]. More concretely, scalable TCP extends the procedure in the congestion avoidance phase in the following manner: the first idea is to adopt a small multiplicative increasing factor of 1.01. Hence, if it receives all ACKs in the packets transmitted when $cwnd = n$, the window size at the next transmission (exponentially) grows to $1.01n$. It is shown that by adopting such a multiplicative increasing factor, the length of intervals between two consecutive congestion detections becomes (almost) independent of the network bandwidth [11] (recall that in AIMD, the time interval for the next congestion increases as the bandwidth of the underlying network is increased). The second idea of the scheme is to adopt a small multiplicative decreasing factor of 0.125, which is also effective for improving the utilization ratio of the broad communication bandwidth.

2.2.2 High-Speed TCP

High-Speed TCP (HS-TCP) is another extension of TCP [10], which modifies the original AIMD in the following manner: an increase of the window size is conducted as $cwnd := w + a(w)/w$, and a decrease of the window size is conducted as $cwnd := w(1 - b(w))$, where w denotes the “current” congestion window size. Functions $a(w)$ and $b(w)$ correspond to parameters α and β in the original AIMD, respectively, and in HS-TCP, those functions are defined in such a way that: 1) $a(w)$ quadratically increases as the window size w is increased, and 2) $b(w)$ logarithmically decreases as w is increased, to realize an efficient utilization of the broad bandwidth of high-speed networks.

2.2.3 Bimodal

Bimodal [8] extends AIMD in such a way that congestions due to the arrival of new senders are distinguished from those caused by the existing senders. In Bimodal, such a distinguishment

is realized by using the notion of, “fair share.” The fair share of a link is defined as B/N , where B is the maximum transmission speed of the link and N is the number of senders. (Note that an ultimate goal of congestion control schemes is to equally assign a transmission speed of B/N to each sender.) If a sender detects a congestion when it is slower than the fair share, it could judge that the congestion is caused by the arrival of new senders. On the other hand, if it is not slower than the fair share, it could judge that the existing senders cause the congestion. A key technique proposed in Bimodal is how to calculate this kind of fair share, since each sender cannot directly acquire the precise values of N and B .

Let $a_i(t)$ denote the value of $cwnd$ in sender i at time t , and let $A(t) = \sum a_i(t)$. Let t_b be the time immediately after a congestion detection, and t_e be the time immediately before the next congestion detection. Each sender linearly increases the value of $cwnd$ during the time interval $[t_b, t_e]$ as in AIMD. Let F_i denote the fair share calculated by sender i . If the first congestion occurs when $cwnd \geq F_i$, it updates the fair share as $F_i = cwnd$, and then, decreases the window size to $(1 - \gamma)cwnd$ for some fixed $\gamma (< \beta)$. On the other hand, if the first congestion occurs when $cwnd < F_i$, it decreases the window size to $(1 - \beta)cwnd$, and updates F_i at the time of the next congestion detection, in the following manner: Since all senders reduce their window size to $(1 - \beta)$ times after the first congestion, it holds $A(t_b) = (1 - \beta)B$. In addition, it holds $A(t_e) = B$ since it fully utilizes the bandwidth immediately before the next congestion. Thus, we have

$$A(t_e) - A(t_b) = \beta B.$$

By letting Δ be an increase of the window size during time interval $[t_b, t_e]$, since it is identical for all senders, we have

$$A(t_e) - A(t_b) = \Delta \times N.$$

Thus, the fair share B/N can be calculated as $B/N := \Delta/\beta$ by using the two known values of Δ and β .

3. MINORITY GAME

In 1997, Challet and Zhang proposed a game called Minority Game (MG) [2, 4] inspired by the El Farol’s bar problem devised by Arthur in 1994 [12]. MG is a simplified model of a financial market, in which the objective of each player is to maximize the profits received from the market. MG has attracted considerable attentions in recent years, and has been investigated extensively during the past decade from various viewpoints [7, 6]

3.1 Overview

Let G be a set of an odd number of players. In MG, all players in G synchronously repeat a round, and receive a profit from the field when it becomes a winner of a round. In each round, each player selects one of two groups (e.g., 0 or 1). Then, the minority of players with respect to the selected groups is determined as the winner of the round. The objective of each player is to maximize the total profits received from the field. By definition, in order to maximize the total profits, we have to maximize the number of winners in each round (e.g., it must be as close to

half the total number of players as possible). In fact, under an appropriately designed MG, we could observe an emergence of cooperative behavior of players such that the number of winners converges to half the number of players in spite of the selfish behavior of each player [2].

Thus far, several variants of MG have been examined in the literature [7, 6]. In the following, among them, we will focus on the basic model of MG in which: 1) each player follows a uniform and simple selection strategy, and 2) each player cannot know the selection of the other players.

3.2 Behavior of Each Player

The selection of a group in each round is conducted by referring to the history of winners in the last m rounds and a set of s strategy tables (STs), where the history of winners is the only information shared by all players. Each ST contains 2^m entries, each of which corresponds to the history in the last m rounds (e.g., entry “001” corresponds to the history such that players who selected groups 0, 0, and 1 in this order won in all of the last three rounds). The content of each entry is either 0 or 1, where for example, content 0 of entry “001” indicates that if the history of winners is 0, 0, 1, then it should select group 0 in the current round. (Note that the number of such STs is $22m$, since each of the 2^m entries corresponds to action 0 or 1.)

Before starting a game, each player randomly selects s STs from such $22m$ STs as a collection of candidate STs. In each round, each player selects an ST from those candidates according to an appropriate selection strategy, and uses it to decide the next selection of the player. More concretely, each player selects an ST with the highest score in each round, where a tie is randomly broken. The score of STs is initialized to zero, and is incremented (resp. decremented) by one if it wins (resp. loses) a round by using the ST. After selecting an ST with the highest score, the player follows the action indicated by the entry of the ST corresponding to the history of winners in the last m rounds.

4. PROPOSED METHOD

In this section, we propose a new congestion control scheme for high-speed networks based on the notion of MG described in the last section. As was explained previously, each player participating in MG tries to maximize its own profits without directly referring to the decisions made by the other players. The selfish behavior of players is very similar to the behavior of senders in communication networks. Hence, it is natural to expect that by adopting MG as a core of congestion control schemes, we could realize an emergence of cooperative behavior of senders such that almost a half of them are selected as the winners, and that this kind of cooperative selection of winners could be effectively applied to the realization of a *selective reduction* of the transmission speed in high-speed communication networks.

4.1 Field of the Game

In the proposed scheme, we regard each sender as a player and the time interval between two consecutive congestions as a round in MG. The profit given to each sender is an assignment of a (broad) bandwidth. More concretely, a winner of a round is given the privilege to start the next transmission with a higher initial transmission speed than the losers (a concrete way for control-

ling the initial transmission speed is described later).

Each router determines the winner of a round. Recall that after receiving a packet from a neighbor, each router first puts it into its local queue, and forwards it to the next router by referring to the header of the packet. A discard of packets at a router occurs when the local queue of the router overflows, and this is the reason why congestions occurs in a router. Thus, congestions occurring in a router can be definitely detected by monitoring the status of all packets passing through the router. In addition, the winner of each round can be identified correctly by each router, if each packet transmitted by a sender is attached to the group that is selected by the sender in the same round (a detailed way for this kind of identification is described below).

4.2 Procedures

At first, each sender randomly selects s candidate STs. Then, at the beginning of each round, each sender selects a group according to the history of winners and an ST with the highest score, where once it selects a group cannot change the selection until it detects the next congestion. The selected group is recorded into the header of all packets transmitted by the sender, and after starting a packet transmission, each sender linearly increases its window size the same way as in AIMD. Packets transmitted by a sender are delivered to the receiver through routers. In each round, each router records the sender of packets passing through the router, and counts the number of senders corresponding to each group. Upon detecting any congestion, it determines the winner of the round by comparing the number of senders corresponding to each group, and sends back the game results to all of the senders, in the following manner:

Feedback to Winners: the feedback to a winner is realized by notifying the fact to the receiver of the winner's packets. More concretely, the router records the fact of winning to (the header of) the packets transmitted by the winner, which will be correctly delivered to the receiver as long as no additional congestions occur. If the router has no packets transmitted by the winner at that time, it waits for the arrival of such packets for a predetermined time period, and if it cannot receive it within the time period, it considers that the winner has terminated the packet transmission, and skips notifying the winner. After receiving a packet notifying that the sender of the packet is a winner in the round, the receiver notifies the fact to the sender by attaching it to the ACK packet that was returned to the sender.

Feedback to Losers: the feedback to a loser is realized by (selectively) discarding packets received from the loser. Each sender recognizes that it becomes a loser of a round if it receives a Duplicate ACK without receiving an ACK from its receiver with a winning notification. Here, the number of packets discarded by the router must be carefully determined in practice, since: 1) if it is too small, it would cause another congestion before completing the notification to the senders, and 2) if it is too large, it would cause an unexpected time out, which wastes a lot of time in the slow-start phase.

After receiving the results of a round, each sender updates the history of winners, and determines the actions for the next round.

4.3 Control of Transmission Speed

Next, we explain how to reflect the game results of a round on the transmission speed in the next round. Recall that upon detecting congestion, (at least one) senders must decrease its

transmission speed by appropriately reducing the window size. A key point in the proposed scheme is to realize a selective reduction of the window size according to the game results. In this paper, we propose the following two schemes to realize such a selective reduction:

- 1) In the first scheme, losers decrease $cwnd$, whereas winners continue to linearly increase $cwnd$ (i.e., each winner is given a profit such that, “it is exempted to reduce the window size.”)
- 2) In the second scheme, all senders decrease $cwnd$ according to the results of the last round (i.e., each winner is given a profit such that “it can restart a transmission with a larger initial window size.”)

Here, it should be noted that we could not “fix” the initial window size in the second scheme in order to realize an adaptive congestion control that is applicable to various situations. In addition, although it would be possible to take an approach such that the router that detected a congestion calculates the initial window size of each sender, it does significantly increase the load of the router, which degrades the overall performance of the routing scheme. Thus, in the following, we will take an approach so that each sender autonomously calculates a reference value of the window size, and uses it to determine the initial window size in the next round. Let Δ denote an increase of $cwnd$ after detecting the last congestion. Then, upon detecting the current congestion, each sender calculates the reference value \tilde{w} as follows:

$$\tilde{w} = \frac{2\Delta}{\beta + \gamma} \quad (1)$$

where γ and β are decreasing factors for winners and losers, respectively (e.g., the winners of the round update the window size as $cwnd := \tilde{w} \times \gamma$, and the losers of the round update the window size as $cwnd := \tilde{w} \times \beta$ and in the following experiments, we fix those parameters to being: $\beta = 0.5$ and $\gamma = 0.2$).

The above definition of the reference value intends to represent the fair share of B/N , which was originally proposed in Bimodal (recall that B denotes the maximum transmission speed and N denotes the number of senders). Let $a_i(t)$ denotes the value of $cwnd$ in sender i at time t , and let $A(t) = \sum a_i(t)$. Let t_b be the time immediately after a congestion detection, and t_e be the time immediately before the next congestion detection. Since it fully utilizes the given bandwidth immediately before *any* congestion detection, by assuming that a half of senders win the game in each round, we have:

$$A(t_b) = \left(1 - \frac{\beta + \gamma}{2}\right) A(t_e)$$

It holds that $A(t_e) - A(t_b) = (\beta + \gamma) B$ since $A(t_e) = B$, and it holds that $A(t_e) - A(t_b) = \Delta \times N$ since $a_i(t_e) - a_i(t_b) = \Delta$ for each i . Hence the fair share of B/N can be calculated as $B/N = 2\Delta/(\beta + \gamma)$ by using the three known values of β , γ , and Δ .

5. EVALUATION

In this section, we evaluate the performance of the proposed scheme by simulation. In the following, we will examine the following two points in detail: 1) how the selective reduction of the transmission speed improves the bandwidth utilization compared with AIMD (Sections 5.2 and 5.3), and 2) how the fairness of bandwidth assignment could be improved by using MG compared with randomized selection schemes (Section 5.4).

5.1 Simulation Model

In the following experiments, we consider a simple network consisting of two routers r_a and r_b connected by a link of 100 Mbps with a delay of 10 ms, and several terminals connected to a router by a link of 1 Gbps with a delay of 2 ms. (See Figure 1 for illustration.) In the following, we suppose that each router is connected with N terminals, and that terminals connected with router r_a behave as senders, and terminals connected with router r_b behave as receivers.

The metrics used in the evaluation are as follows: the throughput is defined as the ratio of the total number of received packets per time to the capacity of the network. The number of arrivals is defined as the number of packets received by each terminal. The fairness of the bandwidth assignment is a distribution of the number of arrivals, which is evaluated by using a fairness metric [9] that is defined as:

$$F = \frac{(\sum x_i)^2}{N \sum x_i^2}$$

where x_i is the number of arrivals to receiver i . (Note that by definition, F is a value in $[0, 1]$, and it is considered to be “fair” if F is sufficiently close to one.)

All experiments are conducted by using NS2. The two parameters of s and m in MG are fixed to three as suggested in [2]. In the simulation, for each combination of parameters, we repeat a “run of 60 sec” 100 times, and take an average of the metrics. By the above results, we can conclude that the throughput of AIMD is maximized when $\beta = 1/3$, and it is sufficiently fair if we do not consider the arrival of new senders. Hence, in the following experiments, we use two versions of AIMD (e.g., AIMD (1/2) with $\beta = 1/2$ and AIMD (1/3) with $\beta = 1/3$, as a reference for evaluating the goodness of the proposed schemes).




Fig. 1. Network structure

5.2 Scheme 1

At first, we evaluate the performance of the first scheme. Recall that in the first scheme, a selective reduction of the transmission speed is realized by reducing the window size of losers. In the following experiments, we compare the performance of the first scheme with two variants of AIMD, by fixing the decreasing factor for the losers to 1/2.

Figure 2(a) illustrates a comparison of the throughput, where the horizontal axis is the number of senders and “Loser” indicates the first scheme. From the figure, we can observe that the first scheme certainly improves the throughput of AIMD (e.g., it improves AIMD [1/2] by 18% and AIMD [1/3] by 16%), which is apparently due to the effect of the *selective reduction* realized in the first scheme. However, such an improvement is achieved by sacrificing the fairness of the scheme. Figure 2 (b) shows the number of arrivals for each receiver, which is sorted in a decreasing order, as before. As shown in the figure, although almost all senders can transmit more packets than AIMD, the first scheme is not fair compared with AIMD. In fact, the fairness metric of the scheme is calculated as $F = 0.990$, which is apparently worse than AIMD.

The unfairness of the bandwidth assignment is due to the following drawbacks of the first scheme: recall that in the first scheme, merely losers decrease their window size to half. However, the one-sided control of the window size does not directly connect winning a game to the profit received from the game. Suppose that a sender wins a game when its transmission speed is low, and another sender loses when its transmission speed is high. Then, if the window size of the loser is greater than a twice that of the winner, the loser can restart its transmission with a larger window size than the winner even though they actually lost the game. That is, it is not guaranteed in the first scheme that a winner of a game can always obtain a higher profit than a loser.

In addition, an amount of profits received from a game strongly depends on the order of wins (e.g., consecutive wins would provide a larger profit compared with a collection of individual wins). Since the first scheme does not take into account the effect of such consecutive wins, the balancing of the number of wins does not directly balance the number of profits received from the game. (Note that a similar claim holds for the distribution of the length of time intervals between two consecutive congestion detections.)

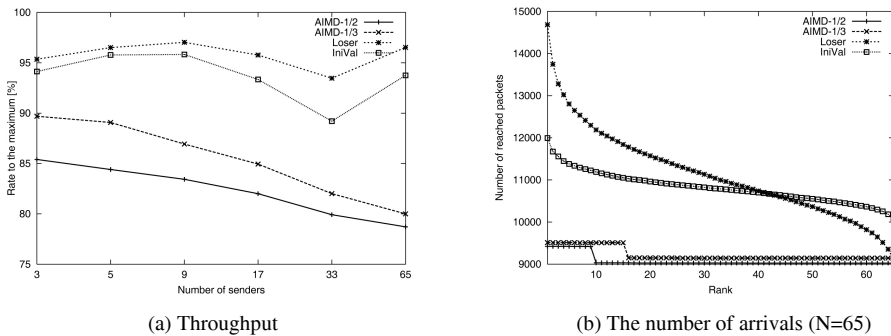


Fig. 2. Performance of the proposed scheme

5.3 Scheme 2

Next, we evaluate the performance of the second scheme. Recall that the basic idea of the second scheme is to use a fair share to calculate the initial window size of each sender. In the following experiments, we fix two multiplicative decreasing factors used in the scheme as $\beta = 0.5$ and $\gamma = 0.2$.

Figure 2 (a) illustrates the throughput of the schemes, where “IniVal” indicates the second scheme. From the figure, we can observe that although the throughput of the second scheme is slightly lower than the first scheme, it still improves the throughput of AIMD (e.g., it improves AIMD [1/2] by 15% and AIMD [1/3] by 14%). In addition, it is worth noting that such an improvement of the throughput does not sacrifice the fairness of the scheme as in the first scheme. In fact, as shown in Figure 2 (b), the second scheme is much fairer compared with the first scheme, and the fairness metric of the scheme is calculated as $F = 0.998$, which is certainly greater than the first scheme and is comparable with the two variants of AIMD. Apparently, the fairness of the second scheme is due to the effect of the fair share used in the calculation of the initial window size. That is, the calculation of the initial window size based on the notion of fair share is not affected by misleading factors such as consecutive wins, and could directly connect the balancing of wins to the balancing of the bandwidth assignment.

5.4 Comparison with Randomized Schemes

Finally, we compare the performance of the proposed schemes with randomized schemes. Note that we can realize a selective reduction of the transmission speed by using a randomization, without using MG as in the proposed schemes. More concretely, we consider the following two randomized schemes in the following experiments: 1) a scheme that randomly selects a “group” in MG without referring to the score of each ST (RAND1), and 2) a scheme that randomly determines the “winner” of each round without playing a game (RAND2).

Figures 3 and 4 illustrate the result. As shown in Figure 3, RAND2 improves the throughput of our second scheme, while RAND1 exhibits almost the same performance with it. However, as for the fairness, both randomized schemes exhibit much worse performance than our second scheme. For example, the fairness metric of RAND2 is calculated as $F = 0.992$, which is as bad as our first scheme, and although RAND1 exhibits almost the same performance with our second scheme for $N = 65$, the distribution of the number of arrivals becomes large when $N = 17$. In fact, the fairness metric of RAND1 for $N = 17$ is calculated as $F = 0.993$, which is much smaller than our second scheme.

A reason for the poor quality of RAND1 for a small number of senders can be explained as follows: first, randomized group selection never converges the number of winners to half of the number of senders as in MG (e.g., for every round, it is far from a half of the number of senders). Since a smaller number of winners implies that it releases a larger amount of resources, it increases the time interval between two consecutive congestion detections that “amplifies” the difference of the profits received from the field. Although such an effect of amplification (due to an increase of the time interval) could be bounded if the number of senders is sufficiently large, it becomes critical when the number of senders is relatively small.

In summary, this kind of randomized selection sacrifices the fairness of the scheme, though it would increase the throughput. Hence, in order to realize a fair sharing of bandwidth with a sufficiently high throughput, we should adopt MG as a means to realize a selective reduction of the

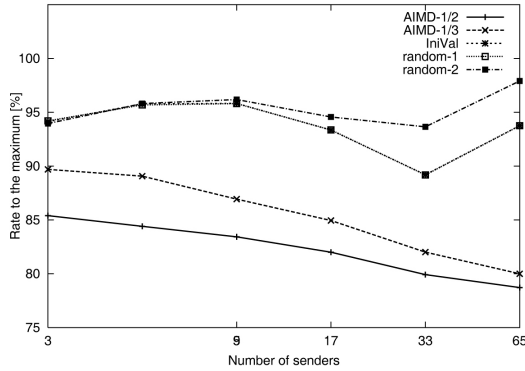
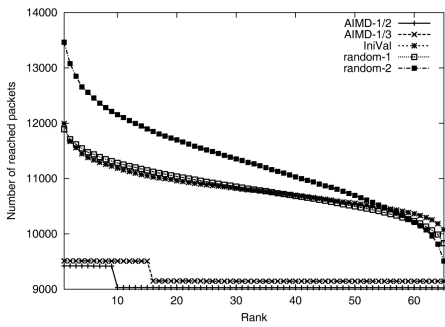
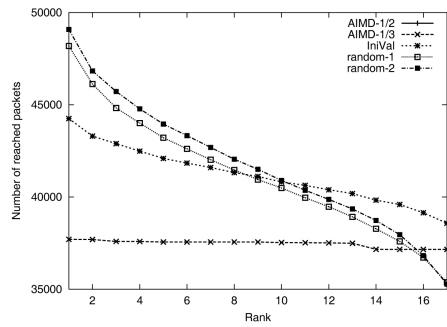


Fig. 3. Throughput of randomized schemes



(a) N=65



(b) N=17

Fig. 4. The number of arrivals

transmission speed.

6. CONCLUDING REMARKS

In this paper, we proposed a new congestion control scheme for high-speed networks. The basic idea of the scheme is to realize a selective reduction of the transmission speed by using the notion of the “Minority Game,” and the results from the experiments indicate that it certainly improves the communication throughput while keeping the fairness of the scheme. An interesting future work is to realize a selective reduction of the transmission speed that reflects the kind of transmitted flows, which will be available in realizing efficient QoS guaranteed routing protocols.

REFERENCES

- [1] URL: <http://www.infinibandta.org>
- [2] D. Challet and Y.-C. Zhang. “Emergence of Cooperation and Organization in an Evolutionary Game”. *PhysicaA*, 1997, pp.246-407.

- [3] D. Chiu and R. Jain. "Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks". *Journal of Computer Networks and ISDN*, Vol.17, No.1, June, 1989, pp.1-14.
- [4] E. Moro. "The Minority Game: an introductory guide". E. Korutcheva and R. Cuerno (Eds.), "Advances in condensed matter and statistical physics". Nova Science Publishers, New York, 2003.
- [5] J. Postel. "Transmission Control Protocol". RFC 793, September, 1981.
- [6] M. Paczuski, K. E. Bassler and A. Corral. "Self-organized Networks of Competing Boolean Agents". *Physical Review Letters*, Vol.84, April, 2000, pp.3185-3188.
- [7] N. F. Johnson, P. M. Hui, R. Jonson and T. S. Lo. "Self-Organized Segregation within an Evolving Population". *Physical Review Letters*, Vol.82, April, 1999, pp.3360-3363.
- [8] P. C. Attie, A. Lahanas and V. Tsoussidis. "Bimodal Congestion Avoidance and Control". *Proceedings of IEEE/ACM Transactions on Networking*, November, 2002.
- [9] R. Jain, D. M. Chiu and H. Hawe. "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems". DEC Research Report TR-301, September, 1984.
- [10] S. Floyd. HighSpeed TCP for Large Congestion Windows. *Internet Draft <draft-floyd-tcp-highspeed-01.txt>*, August, 2002.
- [11] T. Kelly. "Scalable TCP: Improving Performance in High speed Wide Area Networks". *ACM SIGCOMM Computer Communication Review*, Vol.33, 2003, pp.83-91.
- [12] W. B. Arthur. "Inductive Reasoning and Bounded Rationality(The ElFarol Problem)". *Proceedings of American Economic Review*, Vol.84, 1994, pp.406-411.

Hiroshi Kutsuna

He received his BS and MS degrees in Information Engineering from Hiroshima Univ. in 2004 and 2006, respectively. His research interests include game theory and congestion avoidance schemes in communication networks.

Satoshi Fujita

He received his BS, MS, and a Ph.D. degrees in Information Engineering from Hiroshima Univ. in 1985, 1987, and 2000, respectively. He has been a professor at Hiroshima Univ. since 2007. His research interests include parallel algorithms, graph algorithms, communication networks, and resource assignment problem in interconnection networks. He is a member of IEEE Computer Society, SIAM, IPS Japan, and IEICE.