Feedback Modeling in Internet Congestion Control

Sergey Gorinsky

Abstract— **This paper examines the impact of feedback modeling on analytical conclusions about Internet congestion control. In particular, we focus on the problem of convergence of binary adjustment algorithms to fairness. First, we review results for the traditional model that assumes uniform feedback to all users. Then, we demonstrate that under a more realistic feedback model, a larger family of binary adjustment algorithms converges to fairness. We also discuss practical implications of our analytical findings for enhancing Internet congestion control.**

Keywords— **Congestion control, modeling, feedback, adjustment algorithm, fairness, convergence.**

I. Introduction

Congestion control protocols strive to promote efficient and fair utilization of the network capacity. Achieving this goal involves imposing restrictions on communication patterns of distributed applications. For example, instead of allowing an application to blast its data through the network, congestion control can reduce the transmission to a fair rate for the traversed network path. In the Internet, a vast majority of one-to-one communication sessions control congestion using Transmission Control Protocol (TCP) [1], [11]. TCP limits the amount of transmitted but not yet acknowledged data. The limiting window grows when timely acknowledgments confirm lossless delivery. However, if the acknowledgment stream indicates packet loss, TCP reduces the window and thereby curbs the transmission.

As any congestion control design, TCP congestion control consists of two components:

- **Feedback** that notifies a transmitting entity about the congestion status of the network path, and
- **Adjustment algorithm** that regulates the transmission in response to the feedback.

Over the last two decades, numerous enhancements and alternatives have been proposed for both components of TCP congestion control. Proposals of different feedback include DECbit [26], Random Early Detection (RED) [6], and Explicit Congestion Notification (ECN) [25] that enable a congested router to notify the sender about the congestion explicitly without dropping packets, Available Bit Rate (ABR) [16] and eXplicit Control Protocol (XCP) [17] that replace binary feedback with multi-bit reports, Congestion Avoidance using Round-trip Delay (CARD) [13] and Vegas [4] that interpret delay variations as congestion feedback. The set of proposed adjustment algorithms is also plentiful and contains, among others, binomial algorithms [3], Square-Increase Multiplicative-Decrease (SIMD) [15], Additive-Increase Multiplicative-Decrease with Fast Convergence (AIMD-FC) [18], and Ideally-Scalable Congestion Control (ISCC) [19].

Sergey Gorinsky is with Applied Research Laboratory (ARL), Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130-4899, USA. His email address is gorinsky@arl.wustl.edu.

In contrast to the extensive design efforts, modeling of Internet congestion control has been substantially less thorough. Employed models often have little relationship to the Internet reality and lead to incorrect conclusions [7]. This dichotomy is extremely perilous. Simplistic untrustworthy modeling not only undermines understanding of the modern Internet but also stifles widespread adoption of innovative designs because such deployment requires general agreement that the proposed designs constitute an improvement. Hence, convergence to realistic well-understood models is a key for overcoming the Internet ossification.

Inadequacies of existing models are not surprising: congestion control modeling is a hard problem. In particular, it seems all but impossible to construct a single model that, on the one hand, represents all aspects of Internet congestion control accurately and, on the other hand, is simple enough to be useful. A more promising alternative is to develop precise models for specific aspects being investigated [7].

In this paper, we focus on the question of **convergence of binary adjustment algorithms to fairness**. First, Section II presents a traditional model for reasoning about this problem. Then, Section III reviews analytical results for fairness convergence in the traditional model. In Section IV, we identify deviations of the model from the Internet reality. Section V shows that under a more realistic feedback model, a larger family of binary adjustment algorithms converges to fairness. We also discuss practical implications of our analytical findings for enhancing Internet congestion control. Finally, Section VI sums up the paper and suggests directions for future work.

II. Traditional Model

The traditional model for reasoning about fairness convergence of binary adjustment algorithms represents the network as a single resource shared by a set U of distributed users [5]. The model employs a discrete timescale. At every instant t, each user i adjusts its load on the network to a nonnegative value of $x_{i,t}$ (i.e., $x_{i,t} \geq 0$) by executing the investigated adjustment algorithm:

$$\forall\ i \in U \quad x_{i,t} = \begin{cases} I(x_{i,t-1}) & \text{if } y_t = 0, \\ D(x_{i,t-1}) & \text{if } y_t = 1 \end{cases} \tag{1}$$

where I is an increase function, D is a decrease function, and y_t is a binary feedback indicating whether the total load X_{t-1} after the previous adjustment:

$$X_{t-1} = \sum_{i \in U} x_{i,t-1} \qquad (2)$$

exceeds an optimal value C:

$$y_t = \begin{cases} 0 & \text{if } X_{t-1} \leq C, \\ 1 & \text{if } X_{t-1} > C. \end{cases} \qquad (3)$$

The model describes the state of the network at time t as *fair* if all the users at time t impose the same load on the network:

$$\forall\, i, j \in U \quad x_{i,t} = x_{j,t}. \qquad (4)$$

To quantify fairness of the network state at time t, the model uses index F_t from [14]:

$$F_t = \frac{X_t^2}{|U| \sum_{i \in U} x_{i,t}^2} \qquad (5)$$

Possible values of the fairness index cover the range from $\frac{1}{|U|}$ to 1. The index yields its minimal value of $\frac{1}{|U|}$ when only one of the users imposes a positive load (i.e., when $\exists\, j \in U \;\; x_{j,t} > 0$ and $\forall\, i \neq j \;\; x_{i,t} = 0$). The index attains its maximal value of 1 in fair network states.

The model defines that the binary adjustment algorithm *converges to fairness* if for any initial state of the network, the fairness index converges to 1:

$$\Big((\forall\, i \in U \quad x_{i,0} \geq 0) \wedge (t \to \infty) \Big) \Rightarrow (F_t \to 1). \qquad (6)$$

III. Applying the Model

After introducing the above model, [5] applied the model to analyze a family of binary adjustment algorithms that use linear functions for increase and decrease:

$$\forall\, i \in U \quad x_{i,t} = \begin{cases} a_I + m_I x_{i,t-1} & \text{if } y_t = 0, \\ a_D + m_D x_{i,t-1} & \text{if } y_t = 1 \end{cases} \qquad (7)$$

where coefficients a_I, m_I, a_D, and m_D are real numbers. This family includes:

- Multiplicative-Increase Multiplicative-Decrease (MIMD) algorithms with $a_I = 0$, $m_I > 1$, $a_D = 0$, $0 < m_D < 1$,
- Additive-Increase Additive-Decrease (AIAD) algorithms[1] that use $a_I > 0$, $m_I = 1$, $a_D < 0$, $m_D = 1$,
- Additive-Increase Multiplicative-Decrease (AIMD) algorithms where $a_I > 0$, $m_I = 1$, $a_D = 0$, $0 \leq m_D < 1$, and
- Multiplicative-Additive-Increase Multiplicative-Decrease (MAIMD) algorithms that have $a_I > 0$, $m_I > 1$, $a_D = 0$, $0 \leq m_D < 1$.

The analysis in [5] yielded two important conclusions about fairness convergence:

Proposition 1: To converge to fairness, a linear adjustment algorithm should be MAIMD or AIMD.

Proposition 2: For the fastest convergence to fairness, a linear adjustment algorithm should be AIMD.[2]

More recently, the traditional model has been routinely applied to argue that nonlinear adjustment algorithms – such as binomial algorithms [3], SIMD [15], and ISCC [19] – converge to fairness.

IV. Relationship to the Internet Reality

Propositions 1 and 2 are interesting because they serve as a theoretical justification for relying on TCP congestion control in practice. Although TCP executes nonlinear algorithms to control congestion (e.g., upon receiving a non-duplicate acknowledgment in the congestion-avoidance mode, the sender increases the current window by adding a weighted inverse of the window size), linear algorithms can approximate TCP behavior on the scale of round-trip time (RTT): TCP congestion control resembles MIMD in the slow-start mode and behaves similarly to AIMD in the congestion-avoidance mode. Hence, in the light of Proposition 2, the congestion-avoidance mode is supposed to provide TCP with fast convergence to fairness.

The traditional framework for reasoning about fairness convergence makes the following simplifying assumptions about feedback:

- Feedback is **synchronous**, i.e., all the users receive feedback with the same frequency and at the same time.
- Feedback is **uniform**, i.e., feedback is the same for all the users regardless of their individual loads on the network.

These assumptions represent the Internet reality imprecisely. First of all, RTT varies substantially in Internet communications; thus, congestion feedback in the Internet is inherently asynchronous. Second, sessions that transmit more data are more likely to lose a packet during congestion; therefore, even users with the same RTT do not receive uniform feedback from the shared congested link. In the next section, we examine how the deviations of the model affect conclusions about convergence of binary adjustment algorithms to fairness.

V. Significance of the Model Inaccuracies

Whereas the analysis in [5] had an objective to identify binary adjustment algorithms that converge to maxmin fairness [8], [12], [24], experiments have since shown that TCP congestion control does not exhibit such a property: even in stable states, TCP sessions that share a bottleneck link can have different shares of the network capacity. Subsequently, researchers analyzed TCP congestion control under models that represented Internet feedback more realistically, and the derived characterizations of TCP stable states explained the empirical observations [2], [20], [22], [23], [28].

[1] Additive decrease is subject to truncation at 0: $x_{i,t}$ is set to $max(0, a_D + x_{i,t-1})$.

[2] Although, MAIMD can converge faster in rare scenarios [9].

Let us now return from TCP congestion control to the more general context of binary adjustment algorithms and investigate the impact of more realistic feedback modeling on conclusions about fairness convergence. We replace the assumption of feedback uniformity in the traditional model with a more accurate assumption that users with larger loads receive congestion notifications more frequently [23], [27], [28]. To capture this property mathematically, we specify that for a substantially large d:

$$\begin{array}{l} \forall\, i,j \in U \quad \forall\, \tau \geq 0 \\ \qquad (\forall\, t \in [\tau, \tau + d) \ \ x_{i,t} > x_{j,t}) \Rightarrow (n_i > n_j) \end{array} \tag{8}$$

where n_i and n_j represent respectively the number of congestion notifications $y_t = 1$ received by user i and user j during interval $(\tau, \tau + d]$. The revised model still assumes synchronous feedback[3], and each user still adjusts its load at every instant of the discrete timescale. For example, during interval $(\tau, \tau + d]$ user i preforms n_i decreases and $d - n_i$ increases.

Applying the revised model to analyze the network behavior under MIMD yields intriguing results. At time $\tau + d$, user i adjusts its load to:

$$x_{i,\tau+d} = m_D^{n_i} m_I^{d-n_i} x_{i,\tau} \tag{9}$$

where m_D and m_I are respectively the decrease and increase coefficients in the MIMD algorithm. Then,

$$\frac{x_{i,\tau+d}}{x_{j,\tau+d}} = \frac{m_D^{n_i} m_I^{d-n_i} x_{i,\tau}}{m_D^{n_j} m_I^{d-n_j} x_{j,\tau}} = \left(\frac{m_D}{m_I}\right)^{n_i-n_j} \cdot \frac{x_{i,\tau}}{x_{j,\tau}} = k \cdot \frac{x_{i,\tau}}{x_{j,\tau}}$$

with $k = \left(\frac{m_D}{m_I}\right)^{n_i-n_j}$. Suppose that user i imposes a larger load than user j throughout interval $[\tau, \tau + d)$. Then, $n_i > n_j$ according to Property 8. Furthermore, since $0 < m_D < 1$ and $m_I > 1$, we have $0 < k < 1$. Hence, as long as user i imposes a larger load than user j, a series of d adjustments reduces the ratio of the loads closer to 1 by the factor of k. Consequently, each pair of users converges toward imposing equal loads. In this more realistic model, **MIMD converges to fairness**.

The convergence of MIMD to fairness is significant because the traditional analysis declares MIMD unfair and unstable, where *stability* is defined as convergence to efficient fair states – from any initial state, the load of each user under a stable adjustment algorithm converges to and oscillates around the efficient fair share $\frac{C}{|U|}$. Although multiplicative adjustments stabilize the total load around the optimal value C, they do not change the initial ratio of user loads under uniform synchronous feedback. Due to this lack of convergence to fairness, MIMD is not stable in the traditional model. By proving that MIMD converges to fairness under more realistic feedback, we have also demonstrated **stability of MIMD** in the refined model. Therefore, the family of binary adjustment algorithms that converge to stable fair states[4] might be larger in reality than Proposition 1 suggests.[5]

If validated by future empirical studies, the existence of a wider class of stable adjustment algorithms opens exciting avenues for improving Internet congestion control. In particular, AIMD has a disadvantageous property that oscillations of the total load in stable states grow linearly in size as the number of users increases [5], [9]. Experiments with TCP congestion control confirm this property of AIMD: the average loss rate grows linearly with the number of competing TCP sessions and thereby worsens TCP performance [21]. On the other hand, the maximal size of the total load oscillations under MIMD is independent from the number of users [9]. Hence, MIMD can guarantee an upper bound on the average loss rate after convergence. Furthermore, when combined with ECN, MIMD can contain the total load oscillations within the buffer of the bottleneck link and thereby keep the link fully utilized with no packet loss in stable states. These features make MIMD a promising alternative to AIMD adjustments. We plan to design an MIMD-based congestion control protocol and verify experimentally its fairness and stability properties.

VI. Conclusion

This paper examined the impact of feedback modeling on analytical conclusions about Internet congestion control. In particular, we focused on the problem of convergence of binary adjustment algorithms to fairness. First, we reviewed results for the traditional model that assumes uniform feedback to all users. Then, we demonstrated that under a more realistic feedback model, a larger family of binary adjustment algorithms converges to fairness. If supported by future empirical studies, our findings can lead to enhancing Internet congestion control with adjustment algorithms superior to AIMD in terms of their scalability and other properties. We are planning to refine and validate our model for Internet feedback and then apply the model to design promising alternatives to TCP congestion control.

References

[1] M. Allman, V. Paxson, and W. Stevens. TCP Congestion Control. RFC 2581, April 1999.

[3] In an asynchronous version of the traditional model, even AIMD does not converge to fairness [9], [10].

[4] Although, due to the asynchrony of Internet feedback, not necessarily to maxmin-fair states.

[5] This assertion is also supported by the argument in [27] that AIAD converges to fairness under a more realistic feedback model.

[2] E. Altman, K. Avrachenkov, and C. Barakat. A Stochastic Model of TCP/IP with Stationary Random Loss. In *Proceedings ACM SIGCOMM 2000*, August 2000.

[3] D. Bansal and H. Balakrishnan. Binomial Congestion Control Algorithms. In *Proceedings IEEE INFOCOM 2001*, April 2001.

[4] L.S. Brakmo, S.W. O'Malley, and L.L. Peterson. TCP Vegas: New Techniques for Congestion Detection and Avoidance. In *Proceedings ACM SIGCOMM'94*, August 1994.

[5] D. Chiu and R. Jain. Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks. *Journal of Computer Networks and ISDN*, 17(1):1–14, June 1989.

[6] S. Floyd and V. Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

[7] S. Floyd and E. Kohler. Internet Research Needs Better Models. In *Proceedings of HotNets-I*, October 2002.

[8] E. Gafni and D. Bertsekas. Dynamic Control of Session Input Rates in Communication Networks. *IEEE Trans. on Automatic Control*, AC-29(11):1009–1016, November 1984.

[9] S. Gorinsky and H. Vin. Additive Increase Appears Inferior. Technical Report TR2000-18, Department of Computer Sciences, The University of Texas at Austin, May 2000.

[10] S. Gorinsky and H. Vin. Extended Analysis of Binary Adjustment Algorithms. Technical Report TR2002-39, Department of Computer Sciences, The University of Texas at Austin, August 2002.

[11] V. Jacobson. Congestion Avoidance and Control. In *Proceedings ACM SIGCOMM'88*, August 1988.

[12] J. Jaffe. Bottleneck Flow Control. *IEEE Transactions on Communications*, 29:954 – 962, July 1981.

[13] R. Jain. A Delay-Based Approach for Congestion Avoidance in Interconnected Heterogeneous Computer Networks. Technical Report TR-566, DEC, April 1989.

[14] R. Jain, D. Chiu, and W. Hawe. A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems. Technical Report TR-301, DEC, September 1984.

[15] S. Jin, L. Guo, I. Matta, and A. Bestavros. TCP-friendly SIMD Congestion Control and Its Convergence Behavior. In *Proceedings IEEE ICNP 2001*, November 2001.

[16] S. Kalyanaraman, R. Jain, S. Fahmy, R. Goyal, and B. Vandalore. The ERICA Switch Algorithm for ABR Traffic Management in ATM Networks. *IEEE/ACM Transactions on Networking*, 8(1):87–98, Feburary 2000.

[17] D. Katabi, M. Handley, and C. Rohrs. Congestion Control for High Bandwidth-Delay Product Networks. In *Proceedings ACM SIGCOMM 2002*, August 2002.

[18] A. Lahanas and V. Tsaoussidis. Additive Increase Multiplicative Decrease - Fast Convergence (AIMD-FC). In *Proceedings IEEE Networks 2002*, August 2002.

[19] D. Loguinov and H. Radha. Increase-Decrease Congestion Control for Real-time Streaming: Scalability. In *Proceedings IEEE INFOCOM 2002*, June 2002.

[20] J. Mahdavi and S. Floyd. TCP-Friendly Unicast Rate-Based Flow Control. End2end-interest mailing list, January 1997.

[21] R. Morris. Scalable TCP Congestion Control. In *Proceedings IEEE INFOCOM 2000*, March 2000.

[22] A. Mukherjee and J. Strikwerda. Analysis of Dynamic Congestion Control Protocols - A Fokker-Planck Approximation. In *Proceedings ACM SIGCOMM'91*, September 1991.

[23] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP Throughput: A Simple Model and its Empirical Validation. In *Proceedings ACM SIGCOMM'98*, September 1998.

[24] K. K. Ramakrishnan, D. M. Chiu, and R. Jain. Congestion Avoidance in Computer Networks with a Connectionless Network Layer, Part IV- A Selective Binary Feedback Scheme for General Topologies. Technical Report TR-510, DEC, August 1987.

[25] K.K. Ramakrishnan and S. Floyd. A Proposal to Add Explicit Congestion Notification (ECN) to IP. RFC 2481, January 1999.

[26] K.K. Ramakrishnan and R. Jain. A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with Connectionless Network Layer. In *Proceedings ACM SIGCOMM'88*, August 1988.

[27] N. Venkitaraman, T. Kim, K. Lee, S. Lu, and V. Bharghavan. Design and Evaluation of a Suite of Congestion Control Algorithms for the Future Internet. In *Proceedings ACM SIGMETRICS'99*, May 1999.

[28] M. Vojnovic, J. Y. Le Boudec, and C. Boutremans. Global Fairness of Additive-Increase and Multiplicative-Decrease with Heterogeneous Round-Trip Times. In *Proceedings IEEE INFOCOM 2000*, March 2000.