Extended Analysis of Binary Adjustment Algorithms

Sergey Gorinsky Harrick Vin

Technical Report TR2002-39
Department of Computer Sciences
The University of Texas at Austin
Taylor Hall 2.124, Austin, TX 78712, USA
{*gorinsky,vin*}*@cs.utexas.edu*

August 11, 2002

***Abstract* – Congestion control in the Internet relies on binary adjustment algorithms. For example, Transmission Control Protocol (TCP) in its congestion avoidance mode behaves similarly to Additive-Increase Multiplicative-Decrease (AIMD) algorithm. Chiu and Jain offer a theoretical justification for choosing AIMD: among stable linear algorithms, AIMD ensures the quickest convergence to maxmin-fair states. Whereas Chiu-Jain model rests on a well-known unrealistic assumption of uniform feedback, more precise analytic characterizations of TCP behavior are developed and validated. In particular, the advanced theory and experiments agree that TCP congestion control does not converge to maxmin fairness. However, despite the recent progress in TCP feedback modeling, it is still common to use Chiu-Jain model for comparison of binary adjustment algorithms. This paper argues against such practice. We provide evidence that due to the incorrect assumption of uniform feedback, Chiu-Jain model is not suitable for trustworthy conclusions about properties of an adjustment algorithm. We emphasize that until algorithms are analyzed with a more realistic feedback model, optimal choice of a binary adjustment algorithm will remain an open problem.**

I. Introduction

In such a complex distributed system as the Internet, it is all but impossible to provide each user with an exact up-to-date value for its fair and efficient load on the network. Instead, users control congestion with *binary adjustment algorithms*: a user adjusts its load in response to binary signals that indicate whether the user must decrease or can increase the load. For example, Transmission Control Protocol (TCP) exercises binary congestion control – the TCP sender steps up its transmission after receiving a new acknowledgment; the sender reduces its load upon a retransmission timeout or after receiving three duplicate acknowledgments [1], [5]. Until the first indication of congestion, each TCP connection raises its load in a manner resembling the Multiplicative-Increase (MI) algorithm [4]. This reliance on MI is supposed to enable quick convergence to efficient states. Once efficiency is achieved, the TCP connection switches to the congestion avoidance mode and adjusts the load similarly to Additive-Increase Multiplicative-Decrease (AIMD) algorithm [4]. The choice of AIMD is supposed to provide stability, i.e., convergence to fair efficient states.

Chiu and Jain provide a theoretical justification for favoring AIMD: according to their analysis of linear adjustment algorithms for a simple feedback model, AIMD yields the quickest convergence to maxmin-fair states [4]. For simplicity, Chiu-Jain model assumes *uniform feedback* – all users receive identical feedback. In reality however, the probability to receive a congestion indication is higher for the user with a larger load. Subsequent analytical studies of TCP congestion control represent feedback more realistically and predict the transmission rate for a TCP connection more accurately [2], [10], [12], [13]. Furthermore, experiments and more realistic models with *proportional negative feedback* agree that bandwidth allocation under TCP does not converge to maxmin fairness [13], [16].

While reliance of TCP on AIMD does not attain the original goal of convergence to maxmin fairness, it is logical to reexamine the presumed superiority of AIMD over alternative algorithms. In fact, new algorithms have been proposed to improve upon various features of AIMD congestion control [3], [7], [8], [9]. However, despite the recent advances in TCP feedback modeling, it is still common to use Chiu-Jain model for comparison of binary adjustment algorithms. This paper argues against such practice. We provide evidence that due to the incorrect assumption of uniform feedback, Chiu-Jain model is not suitable for trustworthy conclusions about properties of an adjustment algorithm. In particular, we show that albeit the scalable MIMD (Multiplicative-Increase Multiplicative-Decrease) algorithm is not stable under uniform feedback, MIMD does converge to fair states under the more realistic assumption of proportional negative feedback. Our findings suggest that until algorithms are analyzed with a realistic feedback model, optimal choice of a binary adjustment algorithm will remain an open problem.

The rest of our paper is structured as follows. Section II discusses the issue of stability. Section III examines the speed of convergence to fair states. Section IV studies the impact of different RTTs. Finally, Section V sums up our conclusions.