

Metastable Regimes for multiplexed TCP flows

François Baccelli*, M. Lelarge[†] and D. R. McDonald[‡]

October 1, 2004

Abstract

Consider the mean field limit of a model for multiple HTTP sources multiplexed through a drop-tail router. The limit may exhibit two stationary regimes. In the fluid regime the flows are independent, there are no packet losses and the average throughput is high. In the turbulent regime the flows are synchronized, there are periodic congestion epochs with packet losses and the average throughput is reduced.

In the prelimit with a finite number of sources the above regimes become metastable in the sense that we observe periodic fluctuations between the fluid and turbulent regimes. This paper outlines a general framework for describing these metastable regimes.

Key words and phrases: Mean-field, HTTP, Metastable, Large Deviations, tunnelling .

MSC 2000 subject classifications: Primary 60J25, 60K35; secondary 94C99.

1 Introduction

There have been many recent developments on the emerging science of spontaneous order, [5]. Spontaneous order or "synch" may occur among a collection of (stochastic) dynamical systems or particles due to a "coupling" between any one particle and the ensemble. This phenomenon of "synch" is quite common when observing multiple TCP/IP connections (each connection is a dynamical system) routed through a common tail-drop bottleneck router. The coupling is provided by the spurt of packet losses and the resulting rate reductions caused when the total transmission rate exceeds the link rate (i.e. the ensemble average exceeds a threshold). Each connection suffering a loss reduces its transmission rate by half thereby synchronizing a number of connections with a relatively low transmission rate. These connections then increase their transmission rates together according to the rules of TCP (i.e. linearly) until the next spurt of losses. After a while the total transmission rate looks like the familiar saw-tooth rising up to the link rate and then falling abruptly with many connections having synchronized transmission rates.

In [1] we studied HTTP connections routed through a bottleneck router tail-drop router at some popular web site. In contrast to a long lived TCP connection transmitting a large file, HTTP connections tend to be alternate between busy and silent periods. Once a connection to a remote web site is established, the user may click on a link which causes a busy period while a page is transmitted through the bottleneck router. The silent period follows while the user reads the page. When the user is finished with the page he might click on another link to restart the busy period.

*INRIA-ENS

[†]INRIA-ENS and IBM T.J. Watson Research Center

[‡]Department of Mathematics, University of Ottawa, dmdsg@mathstat.uottawa.ca, Research supported in part by NSERC grant A4551

