

Stability of FIFO networks under adversarial models: State of the art [☆]

Vicent Cholvi ^{*}, Juan Echagüe

Universitat Jaume I, Campus del Riu sec, 12071 Castellón, Spain

Received 28 October 2006; received in revised form 23 April 2007; accepted 2 July 2007

Available online 12 July 2007

Responsible Editor: E. Altman

Abstract

Network stability is an important issue that has attracted the attention of many researchers in recent years. Such interest comes from the need to ensure that, as the system runs for an arbitrarily length of time, no server will suffer an unbounded queue buildup.

Over the last few years, much research has been carried out to gain an understanding of the factors that affect the stability of packet-switched networks. In this paper, we attempt to review the most noteworthy results in this area. We will focus on networks where the scheduling policy is of the FIFO type, which is, by far, the most widely adopted policy. We gather these results and present them in an organized manner. Furthermore, we also identify some directions open to future research.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Network stability; FIFO scheduling; Adversarial queuing theory

1. Introduction

A growing number of networking applications today have constraints in terms of their maximum allowable end-to-end delay, packet loss rate, bandwidth, availability, and so forth. Therefore, it is becoming increasingly important to find out the conditions under which a given communication net-

work guarantees performance bounds when dealing with emerging real-time-oriented networking scenarios. In spite of the significant advances in the complexity of communication networks, much work still needs to be done in that direction.

In order to characterize the performance of a network, one crucial issue is that of *stability*, which has become a major topic of study in the last decade. Roughly speaking, a communication network system is said to be *stable* if the number of packets waiting to be delivered (backlog) is finitely bounded at any one time. The importance of such an issue is obvious, since if one cannot guarantee stability, then one cannot hope to be able to ensure deterministic

[☆] This work was partially supported by the Spanish Ministry of Science and Technology under Grant No. TSI2004-02940.

^{*} Corresponding author. Tel.: +34 964728332; fax: +34 964728435.

E-mail address: vcholvi@uji.es (V. Cholvi).

guarantees for most of the network performance metrics.

In a paper dating back to 1975 [1], Kelly proved that in stochastic networks where packet sizes and packet inter-arrival times are exponentially distributed, if the service time at servers follows the same distribution then the well-known scheduling policy FIFO (which is by far the most widely adopted policy) is stable. After that and for many years, the common belief was that only overloaded queues¹ could generate instability, while underloaded ones could only induce delays that are longer than desired, but always remain stable. This general wisdom goes back to the models of packet-switching networks originally developed by Kleinrock [2], and based on Jackson queuing networks [3]. Stability results for more general classes of queueing networks [4,5] also confirmed that only overload generates instability. This belief was shown to be wrong when it was observed that, in some networks, the backlogs in specific queues could grow indefinitely even when such queues were not overloaded [6,7]. It was later observed that instability could also occur, even when the ratio between arrival rate and service rate is arbitrarily small, under the FIFO scheduling policy, both by using probabilistic assumptions [8,9] and by considering deterministic ones [10]. However, the above mentioned counterexamples required that the time needed to process a packet be different from one to another. Clearly, this is not, in general, a valid assumption in packet-switched networks, where servers generally have the same service rates for all packets (such networks are usually said to be of the *Kelly type* [5]). Nevertheless, shortly afterwards it was shown that instability could also arise in some types of Kelly networks, including networks using the FIFO scheduling policy [11,12]. These later results aroused an interest in understanding the stability properties of packet-switched networks.

This paper provides a review and synthesis of the most important results concerning stability of networks using the FIFO scheduling policy, which have usually appeared in a continuous but dispersed form. Here, we bring these results together and present them in an organized manner. Furthermore, throughout this survey we also identify a number of directions open to future research.

The paper has two clearly differentiated parts. In the first, which comprises Sections 2–4, we talk about network stability and discuss the different dimensions from which stability can be investigated. We also characterize the model of input traffic we will use and formally define what we mean by stability in that model. In the second part, which comprises Sections 5–7, we review the state of the art concerning stability in networks with a FIFO scheduling policy, and present the results obtained when considering several different points of view; whereas in Section 5 we present some instability results, in Section 6 we present the most relevant directions used to approach stability and in Section 7 we tackle the problem of deciding which networks are stable. We finish with some concluding remarks in Section 8.

In order to keep a historical perspective of the results, in each reference we include the first version of the paper (mostly conference versions) as well as the final one (mostly journal papers).

2. Network scenario

Network stability can be approached from three dimensions (G, A, P) : the *network topology* G , the *input traffic pattern* A and the *packet scheduling protocol* P .

2.1. Network topology

The network topology constitutes the underlying infrastructure by means of which packets travel from their source to their destination. It is composed of network switches (also referred to as routers, servers or nodes), which are interconnected by means of unidirectional or bidirectional links. Each node contains a server for each outgoing link. Servers may have different service rates (link bandwidths), measured in packets per unit of time. However, in general and for the sake of simplicity, it is assumed that each link can transmit a single packet in each time step (i.e., they have a normalized service rate of 1). Furthermore, there is a propagation delay associated to each link. Each server schedules the packets that must cross the link using a nonpreemptive scheduling policy (which may be different at each server). Packets are forwarded in a store-and-forward manner.

We will represent networks by means of directed graphs $\mathcal{G} = (V, E)$, where vertices V represent the nodes and edges E are the links between servers,

¹ A queue is considered to be overloaded when the total arrival rate at any server is greater than the service rate.

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات