In Proceedings of ACM SIGCOMM'91

Comparison of Rate-Based Service Disciplines*

Hui Zhang and *Srinivasan Keshav*
Computer Science Division
University of California at Berkeley
Berkeley, CA 94720

Abstract

This paper compares six new queue service disciplines that can be implemented at the output queues of switches in a connection-oriented packet switched data network. These are Virtual Clock, Fair Queueing, Delay-Earliest-Due-Date, Jitter-Earliest-Due-Date, Stop-and-Go and Hierarchical Round Robin. We describe their mechanisms, their similarities and differences, and some implementation strategies. In particular, we show why each discipline can or cannot provide bandwidth, delay and delay jitter guarantees. This leads to some interesting conclusions about the relative strengths and weaknesses of each approach.

1 Introduction

High speed networking introduces opportunities for new applications that have stringent performance requirements in terms of throughput, delay, delay jitter [1] and loss rate [3]. Conventional packet switching data networks with window-based flow control and first-come-first-served service discipline cannot provide services with strict performance guarantees. Thus, new rate-based flow control and rate-based service disciplines have been proposed in the context of a connection-oriented network architecture with explicit resource allocation and admission control policies [16].

A rate-based service discipline is one that provides a client with a minimum service rate independent of the traffic characteristics of other clients (though it may serve a client at a rate faster than this minimum). Such a discipline, operating at each switch[2] in the network, manages the following resources: bandwidth, service priority and buffer space. In conjunction with appropriate admission policies, such disciplines allow clients to get performance guarantees in terms of throughput, delay, delay jitter and loss rate.

Several rate-based scheduling disciplines have been proposed recently: Delay-EDD [4, 6], Virtual Clock [16], Fair Queueing[3][2], the Stop-and-Go protocol [7, 8], the Hierarchical Round Robin Discipline (HRR) [11], and Jitter-EDD [14]. We feel that it is instructive to compare and contrast these disciplines since they are closely related, yet they have some important differences.

In particular, we wish to answer the following questions:

- What are the similarities and differences between the underlying mechanisms?
- What kinds of performance guarantees can be provided?
- What are the buffer space requirements?
- What are the associated admission control policies?
- What are the implementation issues?

We do not intend to select any particular discipline as the 'best'; rather, our intention is to explore the differences between the disciplines. The choice of which discipline to implement depends upon a number of factors, such as the performance requirements, the workload characteristics, network configuration and implementation complexity. Since these are site-specific variables, we do not feel that it is appropriate to declare any discipline the 'winner'.

A comparison between Virtual Clock and Leaky Bucket [13] is presented in [15]. However we feel that comparing Leaky Bucket and Virtual Clock is inappropriate since Leaky Bucket is not a service discipline: it just controls the rate of one channel and does not specify the service order of packets from different channels.

Another similar study [17] compares Virtual Clock, EDD, Stop-and-Go and HRR. It uses TDM as the baseline for comparison and takes a slightly different perspective from ours.

---

*This research was supported by the National Science Foundation and the Defense Advanced Research Projects Agency (DARPA) under Cooperative Agreement NCR-8919038 with the Corporation for National Research Initiatives, by AT&T Bell Laboratories, Hitachi, Ltd., the University of California under a MICRO grant, and the International Computer Science Institute. The views and conclusions contained in this document are those of the authors, and should not be interpreted as representing official policies, either expressed or implied, of the U.S. Government or any of the sponsoring organizations.

[1] Delay jitter is defined to be the maximum difference between end-to-end delays experienced by any two packets [14].

[2] In the literature, the term 'switch' is used in the context of ATM networks, while 'gateway' or 'router' is used in the Internet environment. In this paper, we will uniformly call all switching elements as 'switches'.

[3] Fair Queueing was first proposed in datagram networks. However, when used in a connection-oriented network architecture with resource allocation and admission control, as considered in this paper, it can also provide bandwidth guarantees as mentioned in [2]

Our study is more complete in the sense that we also compare Fair Queueing, and discuss buffer space requirements, admission control policies and implementation issues.

The guarantees discussed in this paper are deterministic guarantees as defined in [3]: that is, all packets should meet a guarantee even in the worst case. Static allocation of buffers is assumed to prevent packet loss even in the worst case. While statistical guarantees can also be provided by controlled overbooking resources, as proposed in [4, 9], we do not discuss those guarantees in this paper.

2 Background

Guaranteed service requires that the network *protect* clients from two sources of variability: misbehaved users and network load fluctuations. Misbehaved users can send packets at a higher rate than the bandwidth allocated to them and network load fluctuations may cause a higher instantaneous arrival rate from a channel[4] at some switch, even though the channel satisfies the bandwidth allocation constraint at the entrance to the network. Traffic distortion due to network fluctuations suggests that protection should be implemented within the network using a rate-based service discipline. Such a discipline achieves protection by ensuring that each channel gets a guaranteed rate of service regardless of the behavior of other channels.

What should happen when a channel sends packets at a higher rate than the bandwidth allocated to it, but this higher rate will not affect the performance of other channels? Rate-based service disciplines can be classified into two categories depending on the policy they adopt:

- Rate allocating service disciplines will serve packets at the higher rate as