Minimizing Session Delays in a Network of Schedulers

Nandita Dukkipati and David Gutierrez

I. Introduction

Provision of Quality-of-Service is an important issue in the Internet. Currently, the Internet just provides a *Best Effort* service and does not have any mechanism to provide QoS measures such as delay, throughput, packet loss and delay jitter. The two fundamental pieces for QoS implementation are:

1. QoS within a single network element (for example, buffer management, scheduling and traffic shaping)
2. Techniques for coordinating QoS at each of the network elements to provide end-to-end QoS

Schedulers at the network nodes are an integral part in the provision of QoS. Scheduling algorithms are used to select packets for transmission on a link and hence control the the bandwidth allocation among the different flows at a node. The schedulers at different nodes in the network have to coordinate in order to provide the *end-to-end* QoS. It is often easy to analyze the scheduling disciplines at a single node. Analyzing the scheduling algorithms in a network scenario is practically more important and difficult as well.

II. Proposed Problem

A. Motivation

Most of the flows in the Internet today are *short* flows (also known as *mice*) and have about 30 packets on an average. An important measure of QoS for such flows is the response time as seen by a user. The response time as seen by a user is the time interval between session initiation and the time when the last packet of the session is received. Henceforth, we will refer to this measure as the *session delay*.

It is of interest to minimize the average session delay for the sessions in the network. It is known that in case of a single node, the Shortest Job First (SJF), scheduling algorithm provides the minimum average session delay. It is however, not an optimal discipline in the network case. There is no known optimal scheduling discipline to minimize the average session delays in the network scenario. So, at this point the natural question that arises is: How do the different scheduling disciplines proposed in literature, perform relative to each other, with respect to the average session delay in the network case?

This motivates us to the following proposed problem.

B. Problem

Assumptions:
1. The sessions in the network are (σ, ρ) characterized
2. There is per-flow queueing at all the nodes in the network
3. The network could be of any arbitrary topology

The problem proposed is:
1. Find the rate allocation for sessions at every node along their path, in a network of Generalized Processor Sharing schedulers, in order to minimize the average session delay.
2. How does the average session delay in (1) above compare with the average session delay when SJF schedulers are used at all the nodes in the network. (Note: there is no rate allocation in the case of SJF schedulers).

We intend doing simulations to obtain an intuition about how the two schedulers perform with respect to the measure of average session delay and then a deterministic analysis to support the results.

C. Related Work

Most of the previous work in literature deals with the problem of obtaining a bound on delay and backlog for sessions in GPS schedulers, *given* a particular rate allocation. The problem being addressed here is different because we have to find the rate allocations to minimize the session delays. Further, most of the previous work was interested in obtaining bounds on per packet delay. In our case we are interested in the total session delay i.e. the departure time of the last packet of the session minus the arrival time of the session to the network.

Parekh and Gallger proposed GPS in their papers [1] and [2].

REFERENCES

[1] A.K Parekh and R.G. Gallager, A Generalized Processor Sharing Approach to Flow Control in Integrated Services Packet Networks: The Single Node Case, IEEE/ACM Transactions on Networking, vol.1, No. 3, pp. 344-357, June 1993

[2] A.K Parekh and R.G. Gallager, A Generalized Processor Sharing Approach to Flow Control in Integrated Services Packet Networks: The Multiple Node Case, IEEE/ACM Transactions on Networking, vol.2, No. 2, pp. 137-150, April 1994