TRAFFIC MANAGEMENT IN INTEGRATED SERVICES NETWORKS: SCHEDULING AND RESOURCE PARTITIONING

by
RAHUL GARG
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

Submitted
in fulfillment of the requirements of the degree of
Doctor of Philosophy

to the

INDIAN INSTITUTE OF TECHNOLOGY, DELHI
HAUZ KHAS, NEW DELHI-110016, INDIA
July 1999

Certificate

This is to certify that the thesis titled **"Traffic Management in Integrated Services Networks: Scheduling and Resource Partitioning"** being submitted by Rahul Garg to the Department of Computer Science and Engineering, Indian Institute of Technology, Delhi, for the award of the degree of Doctor of Philosophy, is a record of bona-fide research work carried out by him under my supervision, and in my opinion, it has reached the standard fulfilling the requirements of the regulations relating to the degree.

The results obtained in this thesis have not been submitted to any other university or institute for the award of a degree or a diploma.

Dr. Huzur Saran
Professor
Department of Computer Science and Engineering
Indian Institute of Technology
New Delhi 110016

To my parents

Acknowledgements

The journey towards my Ph.D. was long and had many diversions. This could not have been carried out without the constant support and encouragement which I received from my parents. I dedicate this thesis to them.

There were many people who helped me in a number of ways in this pursuit and need a special mention.

Firstly, I thank Dr. Srinivasan Keshav who opened my eyes to the intricacies of computer networks. His undergraduate course in networks, which he taught while he was visiting IIT-Delhi from Bell Labs, got me interested in the topic and helped me build a solid foundation. Secondly, I thank Dr. Xiaoqiang Chen whom I met during my summer internship at Bell Labs. Chen gave me support and encouragement in my most difficult times. The work with him forms the most important part of this thesis. I thank my advisor Prof. Huzur Saran for guiding the later parts of my thesis. He provided me support in difficult times and helped me in a number of ways. My work with him forms a significant contribution of this thesis. I thank Prof. Raphael Rom, who was a Sun Microsystems Labs, for guiding my way during the later stages. Prof. Subhashis Banerjee was especially helpful by creating an open and healthy work environment at IIT Delhi. I thank him for taking pains to make sure that students get "unrestricted" access to most of the computing facilities in the CSE Department of IIT Delhi. Finally, I thank Prof. S.N. Maheshwari, whose encouragement right from the undergraduate days prompted me to pursue a Ph.D.

I thank my friend Vishal for providing cheerful company during the dull and depressing days at IIT-Delhi. Vivek who was physically at Santa Cruz, also made his presence felt to me in Delhi, thanks to the newer communication media. I thank all my other friends at IIT-Delhi including Shiv, Sachin, Sharad, Avinash, for their excellent company.

Finally I thank all my newly formed friends at IBM, especially Vinayaka, Johara, Tanveer and Raghavendra who also read the final draft of my thesis and pointed out several "bugs".

Abstract

Traffic characterization, admission control and packet scheduling are important traffic management functions needed for provisioning QoS in Integrated Services Networks. In this thesis, we examine each of these functions in the context of generic networks and virtual (private) networks providing integrated services.

We characterize burstiness of several long traces of video at different time scales using burstiness function. We examine video compressed using JPEG, MPEG and NV's compression algorithms. JPEG video has low short term burstiness and a large long term burstiness, MPEG video has burstiness at short as well as long time scale whereas NV's traffic has only short term burstiness. We discuss the implications of burstiness at different time scale on bandwidth and buffer allocation.

We then focus on deterministic guarantees (which provide a bound on worst case behavior) and compare the performance of leaky bucket traffic model with $X_{min}, X_{avg} I$ traffic model while using optimal admission control tests. We found that the leaky bucket model outperforms the X_{min}, X_{avg}, I model. We show that using good traffic models and optimal admission control tests, even deterministic guarantees provide reasonable network utilization.

An important contribution of this thesis is a new scheduling algorithm called the Recursive Round Robin (RRR) scheduler. The proposed scheduler operates on fixed sized packets and needs simple bit manipulation operations to make scheduling decisions. Thus it can operate at very high speeds. We discuss several variants of the basic scheduler, including the variable size packet scheduler. The delay and fairness properties of the scheduler are analytically derived and discussed in detail. Efficient hardware and software implementations of the scheduler are also suggested. The scheduler has applications in ATM network interface cards, ATM switches and IP routers.

We next show why traditional schedulers are inappropriate for providing bounded delay services in virtual (private) networks. We define the concept of output burstiness and show how generic latency-rate schedulers with bounded output burstiness can provide good end-to-end delay bound in virtual networks. We suggest two variants of the RRR schedulers for virtual networks.

The output burstiness constraint reduces the amount of sharing possible in virtual networks. We propose a capacity resizing approach to improve sharing in virtual networks. The capacity increase requests which are sent at a granularity of session arrivals, are admission controlled, using a new technique called Stochastic Fair Sharing (SFS), which fairly redistributes the free link capacity among different traffic classes. SFS uses the concept of trunk reservation to give high priority to traffic classes with low normalized usage. SFS can be used to achieve better link sharing in virtual networks. We present simulation results for a single link and a twelve node network. SFS achieves fair sharing on a single link and max-min fair sharing in a network with fairness index between 0.96 and 0.99, as compared to an index between 0.66 and 0.97 for simple schemes like FCFS. The bandwidth penalty for using SFS was less than 5% and the signaling load for SFS in a network was less than 100 messages per second per router.

Contents