Management and Control of Resources in Broadband Networks with Quality of Service Guarantees

Aurel A. Lazar and Giovanni Pacifici

Department of Electrical Engineering
and
Center for Telecommunications Research
Columbia University, New York, NY 10027-6699

Abstract

Resource sharing policies and their evaluation for broadband networks with quality of service guarantees are presented. The networks under consideration guarantee quality of service based on the principle of Asynchronous Time Sharing. Scheduling, distributed scheduling and admission control strategies and the relationship among them are discussed. The performance of the algorithms is evaluated based on the schedulable and admissible load regions. Finally, a reference model that defines the main primitives of the network control architecture is presented.

1. Introduction

Integrated telecommunication networks carry traffic of several different classes, including one or more real-time (isochronous) traffic classes, each with its own set of traffic characteristics and performance requirements. Two different approaches have been advanced to deal with this phenomenon. In the circuit-switched approach (synchronous transfer mode), sufficient resources are allocated to each call to handle its maximum utilization; this guarantees that the call will get the quality of service (QOS) it requires, but may be wasteful of system resources. In the packet-switched approach (*e.g.*, asynchronous transfer mode), traffic from all sources is packetized, and statistical multiplexing techniques are used to combine all network traffic through a single switching fabric. This allows higher network utilization, but requires more sophisticated controls to ensure that the appropriate QOS is provided.

In broadband networks, the internodal propagation delays are more significant than node processing delays. This observation has been used in the published literature to argue that dynamic, adaptive feedback, reactive control algorithms, operating within the network, are not suitable for broadband networks [1]. The argument typically states that because of high transmission speeds, by the time a downstream node detects a congestion condition and attempts to signal its neighbouring upstream nodes to adjust their behaviours, the large number of cells in-transit could not possibly be affected by the closed-loop feedback controls.

Instead the advocated control algorithms have been network-edge, congestion avoidance, preventive algorithms. In this case, it is assumed that there is a pre-negotiated contract between the network sources (terminating entities) and the network control architecture, characterizing traffic peak rates, average rates and the burstiness of the traffic stream that each source is allowed to transmit. Mechanisms are put in place in the control architecture to police the actual traffic behaviour of a network source to ensure that it does not exceed the limits set forth in the negotiated contract [2]. Other proposed refinements to traffic flow enforcement include mechanisms that mark and discard excess traffic in the presence of network overload and congestion [3].

Not withstanding this preventive control approach, there have been published results [4] indicating that, whatever the degree of sophistication of the network-edge preventive algorithms, additional reactive controls may be necessary within the network fabric to adequately handle the complex dynamic fluctuations in high speed traffic interactions. Scheduling and buffer management are examples of such controls.

This paper deals with the design principles of resource control algorithms, together with their interaction and cooperation in a wide area network environment, and presents a framework for evaluating the overall performance of the system. We will focus in particular on two levels of control: scheduling, that mediates the low level competition for service between cells of different classes, and admission control, that regulates the acceptance or blocking of incoming traffic on a call-by-call basis. The performance of the scheduling algorithms will be evaluated based on the schedulable region. The interaction between scheduling and admission control will be quantified using the admission control region. Both the schedulable region and the admission control region are concepts that have been recently introduced in the literature [5, 6].

The resource control algorithms are based on the Asynchronous Time-Sharing (ATS) design principle [7]. ATS is a set of resource allocation principles for the design of broadband packet-switched networks that guarantee QOS. ATS-based networks are similar to those based on Asynchronous Transfer Mode (ATM) in that all traffic offered to the network is in the form of small, fixed-size cells. The primary distinction of ATS is that several classes of traffic with different QOS requirements are considered explicitly at every level of system design, both at the edge and at the core of the network. Therefore, one of the fundamental requirements on ATS systems is that the *core* of the network makes a distinction between traffic classes.

These design principles have broad applicability, and can help to *efficiently* provide QOS in many different network settings. They have already been used in the design of two high-speed integrated networks: MAGNET II [8], a testbed for MAN applications, and TeraNet [9], a gigabit/s lightwave network. The introduction of traffic classes into ATM networks, although not in the ATM standard at this time, may be accomplished in a fully compatible manner. For example, traffic class information could be carried in the Virtual Channel Identifier field of the cell header.

This paper is organized as follows. In Section 2 the basic concepts of the ATS framework are presented, along with an overview of an architecture for joint scheduling and admission control. The scheduling problem is introduced in Section 3, while in Section 4 the extension to the networking environment is discussed. In Section 5 the admission control problem and its interplay with scheduling is formulated. In Section 6 a reference model for broadband networks is briefly presented. The emphasis in this section is on the overall network control architecture.

2. Problem Setting

The generic resource allocation problem presented in this paper was originally motivated by requirements on broadband networks with quality of service guarantees. A class networks based on the concept of Asynchronous Time Sharing was implemented to meet these requirements. The switching architecture of these networks is briefly described in Section 2.1. Four traffic classes are introduced via quality of service constraints. Note that, in order to keep the complexity of the network manageable, the QOS for these classes is defined for the network as a whole, rather than for each individual call. The introduction of traffic classes calls for the introduction of resource allocation algorithms on both the cell and on the call level. The resource allocation problem is introduced in Section 2.2.

2.1 The Architecture and Framework for ATS

At the heart of the distinction between ATS and ATM is a clear definition of traffic classes based on QOS considerations; fundamental to any performance analysis is the set of modeling assumptions on which the analysis is based. This section describes these and other key elements of the ATS approach.

We consider a class of networks that guarantee quality of service based on the Asynchronous Time Sharing principle. The basic architecture of the ATS-based switching node has been recently implemented in a new prototype multihop lightwave network called TeraNet [9]. The architecture of the network interface units (switching nodes) is shown in Figure 1.

Each network interface unit consists of three input links, three output links and a bus based non-blocking switch fabric. Congestion may arise only at the output links. The architecture supports

for finding the region in the space of loads for which the average time delay is finite. In our case the set of constraints that determine the schedulable region is defined by the QOS constraints at the cell level. Examples of constraints were given in Section 2.1 and include: hard time delay constraints, probability of blocking and average gap constraints, average throughput and average time delay constraints. Note that the schedulable region might be finite even for the case of a queueing system with finite buffer size. This is because the QOS constraints at the cell level might restrict the loading on the system before the finite buffer size does.

In Asynchronous Time Sharing (ATS), transmission resources are time-shared between traffic classes according to a cycle scheme [7]. MAGNET II Real-time Scheduling (MARS) algorithm [5] is a mechanism for adaptively setting the parameters which govern this cycle scheme, based on observations of cell arrivals and departures. The scheduling algorithm is based on the intuition that in order to achieve high throughput, each cycle should serve only the cells whose transmission cannot be further delayed to satisfy the QOS requirements.

Figure 3 depicts the schedulable region for the MARS algorithm.

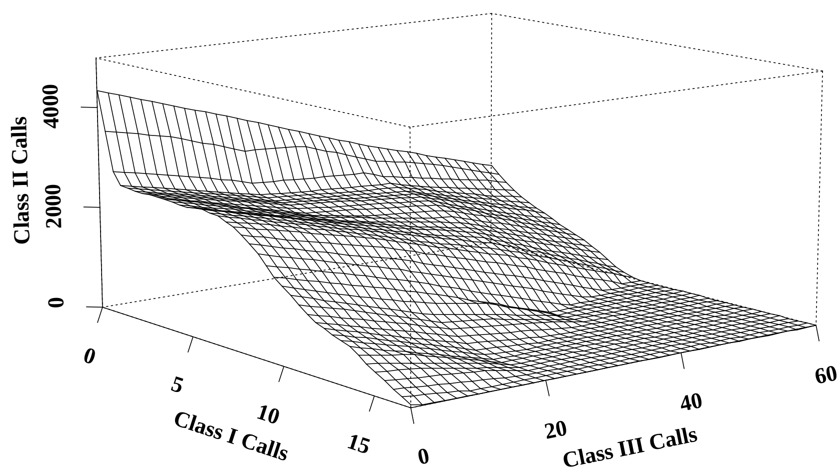

Figure 3: Schedulable region for MARS, with QOS=[2 ms, 4 ms, 0.001, 5.0, 8 ms].

4. Distributed Scheduling in Broadband Networks

In the previous section we discussed the role of resource allocation algorithms for an ATS node taken in isolation. However, to efficiently cope with congestion in a wide area network environment, an interaction among the different resource allocation algorithms is required. The design principles of cooperative distributed algorithms for wide area integrated networks, with substantial delays on the communication links, are the object of this section.

With cooperative distributed algorithms the quality of service, for all the network nodes, is met through *coordination*. This coordination involves the following actions. Each node predicts the traffic streams of any upstream neighbours on a horizon equal to *twice the propagation delay* between the node and each of its upstream neighbours. Feedback signals to upstream nodes are triggered by comparing the quality of service parameters of *estimated queueing dynamics,* derived from the traffic prediction entities, with threshold values. By using traffic prediction a node can anticipate (local) network overload and congestion and still have ample time to send feedback signals to affect upstream cell transmissions at the times when congestion is expected to occur. Whenever an upstream node receives an adaptive feedback signal, it discards cells of some of its traffic classes up to the limits that

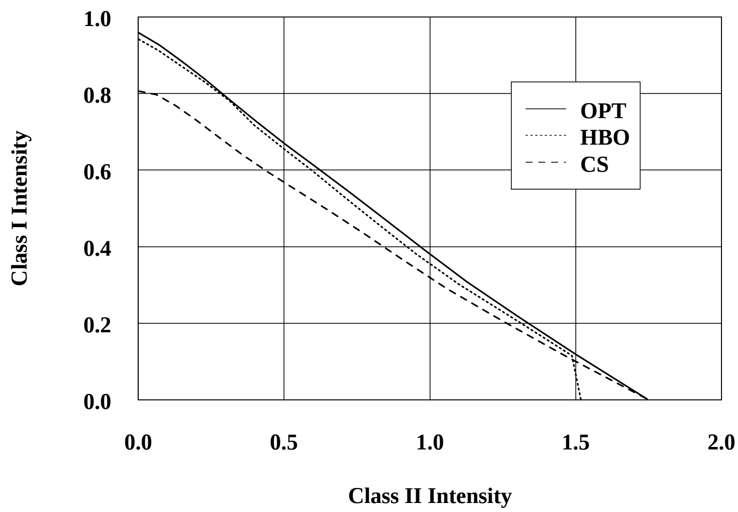

Figure 6: Admissible load regions for the optimal, HBO, CS admission control policies, with call blocking constraints $\kappa^I = \kappa^{II} = 0.1$

6. The Network Control Architecture

In the previous sections we have discussed the design principle of scheduling and admission control, emphasizing the problem of how to distribute these resource allocation algorithms and how to ensure effective communication among them. In order to more formally identify the issues arising in congestion control of high speed networks and to show how the ideas presented here for two classes of control strategies can be generalized, an Integrated Reference Model for broadband networks is briefly presented.

The network architecture contains the primitives for controls, communications and management. These are organized in the Traffic Control, Information Transport and Management Architectures, respectively. The subdivision of the IRM into the TCA and ITA is based on the principle of separation between communications and controls [12]. The separation between the MA and TCA is primarily due to the different time scales on which these architectures operate. Note, however, that in addition to control functions management also includes other tasks such as, e.g. , fault management. The Traffic Control and the Information Transport Architectures are logically divided into a set of vertical planes and a number of horizontal layers and modules (see Figure 7). The purpose of this division is to facilitate the identification of the main issues when implementing the network architecture. The vertical subdivision corresponds to the main control and communications tasks. The control and communications tasks are originated, respectively, in the resource management and control (M)-, resource monitoring and management (D)-, connection management and control (C)-, and user transport (U)-planes. The first three planes are part of the Traffic Control Architecture. The latter plane is part of the Information Transport Architecture. A plane is characterized by a set of entities and their relationships. The (M)-plane has the entities and algorithms responsible for resource management and control. The (D)-plane contains the entities and algorithms for monitoring and management. (The data about the network is stored in a Knowledge Database.) The (C)-plane contains the entities and algorithms responsible for connection management and control. The (U)-plane models the user transport of information. All entities and algorithms that support or are part of information transport are organized in this plane. The (U)-and (C)-planes are horizontally layered. The horizontal subdivision corresponds to the layering concept originally introduced by the OSI RM. Recursive application of the OSI Service Model consisting of a service provider and multiple service users is the basis for layering the (U)-and (C)-planes. The (D)- and (M)-planes consist of a number of objects or modules.

[8] A. A. Lazar, A. T. Temple, and R. Gidron, "MAGNET II: A metropolitan area network based on asynchronous time sharing," *IEEE Journal on Selected Areas in Communications*, vol. SAC-8, pp. 1582–1594, October 1990.

[9] R. Gidron and A. T. Temple, "TeraNet: A multihop multichannel lightwave network," in *Proceedings of the IEEE International Conference on Communications*, (Denver, CO), pp. 602–608, June 1991.

[10] A. A. Lazar, G. Pacifici, and J. S. White, "Real-time traffic measurements on MAGNET II," *IEEE Journal on Selected Areas in Communications*, vol. SAC-8, pp. 467–483, April 1990.

[11] J.-T. Amenyo, A. A. Lazar, and G. Pacifici, "Cooperative distributed scheduling for ATS-based broadband networks," CTR Technical Report 240–91–21, Center for Telecommunications Research, Columbia University, New York, August 1991.

[12] A. A. Lazar, "An architecture for real-time control of broadband networks," in *Proceedings of the IEEE Global Telecommunications Conference*, (Phoenix, AZ), pp. 289–295, December 2–5 1991.