Packet Scheduling for Multimedia Traffic in a Resource-Partitioned Network

Robert Simon
Department of Computer Science
George Mason University
Fairfax, VA 22030

Abstract

We present a general purpose model and schedulability condition for supporting multiple service classes using bandwidth resource partitioning. The model supports guaranteed communication performance for real-time multimedia traffic via an Earliest Deadline First (EDF) packet transmission policy. We show how the EDF model can incorporate the effect of multiple service classes sharing the same transmission link.

1 Introduction

Future integrated service networks will be able to simultaneously support multiple traffic classes, including real-time, multimedia and best-effort traffic. Packets from different traffic classes may be scheduled for transmission according to different service policies. However, when a single communication link is shared by packets from multiple classes then the scheduling algorithms must take into account the effect of multiplexing multiple service classes on the same link.

This paper describes a link-based packet scheduling model supporting performance guarantees for delay and throughput for real-time and multimedia communication traffic. The model is based on the assumption that the output link is shared by multiple classes of traffic. Packets from each of these traffic classes are supported by potentially different scheduling policies and admission control procedures. We first explain the network model, and then present a schedulability result.

2 Traffic and Link Model

Our link model for packet transmission allows multiple types of service classes to coexist. This flexibility if important since many Integrated Service Network models propose the simultaneous deployment of multiple service classes in order to provide network clients a range of service options [2]. The result is increased network utilization and fairer pricing structures.

All packets arriving for transmission at the link are tagged as belonging to a specific service class. Packets from each service class share the same links and may have their own set of transmission scheduling mechanisms and admission control tests. We refer to the traffic class which supports guaranteed timing and throughput as the guaranteed class. Other service classes may offer statistical guarantees service, best-effort, etc. Guaranteeing communication with multiple service classes and traffic types is possible by enforcing *resource partitioning* between the guaranteed service class packets and the other classes. Partitioning ensures the guaranteed service class a known and predictable amount of bandwidth. The partitioning model is flexible, and does not specify exactly how the allocation is implemented.

Each multimedia traffic request requiring deterministic support is called a *guaranteed multimedia channel*, or *gmc*. A *gmc* m is characterized a per-link delay bound, δ_m, as well as a packet workload function, $U_m(\tau)$. δ_m is the maximum time delay that a packet from from *gmc* m can experience. $U_m(\tau)$ represents an upper bound on the number of packets from *gmc* m which the link may receive for transmission over an arbitrary time τ. The values of δ_m and $U_m(\tau)$ are determined by the Quality-of-Service needs for *gmc* m. The actual determination of these values is outside the scope of this paper; see [3] for an example of how they could be calculated. A new *gmc* channel is accepted if the link scheduling policy can *guarantee* that, over any arbitrary time period τ, a total of $U_m(\tau)$ packets can be transmitted, such that that maximum delay that any packet experiences before transmission is no more than δ_m.

The switching node scheduling architecture is shown in Figure 1. Packets are placed in a service queue based on their service class. There are two scheduling policies which must be modeled – inter-queue scheduling and intra-queue scheduling for the guaranteed service class. The inter-queue scheduling policy determines when the packets from a service queue are eligible to be transmitted. The guaranteed service class intra-queue scheduling policy determines which *gmc* packets are selected for transmission and when these packets are transmitted.

The inter-queue scheduling policy ensures the guaranteed service class is allocated a fixed amount of bandwidth, under worst case circumstances of maximum congestion, through resource partitioning. The allocation partitioning strategy works as follows: Let P_i be the partition value for the guaranteed class for an arbitrary link, identified as link i. P_i is in the range

$$0 \leq P_i \leq 1$$

Let r_i be the transmission speed of link i in bits per second. Suppose d is the largest packet size in bits for any service queue. Then the *maximum* time D_i to transmit one packet from any service queue is given by

$$D_i = \frac{d}{r_i} \tag{1}$$

The total number of maximum-sized packets which can be transmitted per second, C_i, is given by

$$C_i = \left\lfloor \frac{1}{D_i} \right\rfloor$$

Resource partitioning ensures that the guaranteed service class is allocated G_i packets per second, where G_i is given by

$$G_i = \lfloor P_i \times C_i \rfloor \tag{2}$$

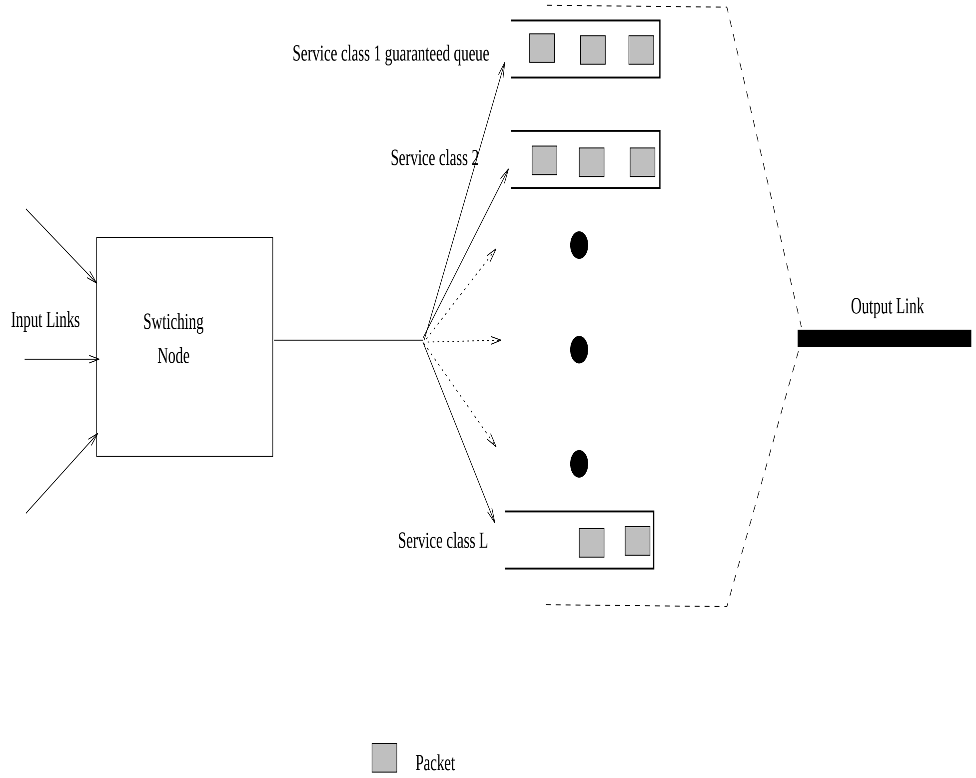

Figure 1: Switching node scheduling architecture.

The only inter-class scheduling policy we impose is that relationship (2) is obeyed. G_i represents a lower bound on the bandwidth the guaranteed service queue receives under worst-case conditions of maximum congestion, when all service queues are being utilized at full capacity. Under this general policy a variety of actual inter-queue scheduling policies can be used. For instance, inter-queue scheduling could be done after each packet finishes transmitting. If a packet were available on the guaranteed queue then it could be eligible for immediate transmission. In this case $P_i = 1$. P_i can be less than 1 under a variety of other inter-queue scheduling policies, such as Time Division Multiplexing. Because of packet-switched architecture, other service classes may use any underutilized bandwidth from the guaranteed service class.

We now explain the intra-queue scheduling discipline for the guaranteed service class. The model for scheduling *gmc* packets on the guaranteed service class is based upon a real-time priority scheduling algorithm, originally proposed by Liu and Layland, called Earliest Deadline First (EDF) [1]. The scheduling algorithm is known to yield a feasible schedule if one exists. Our technique can easily be extended to account for other types of scheduling algorithms which provide guaranteed service, such as static priority or rate monotonic.

Within our ISN model EDF scheduling is explained as follows: Consider the case of K *gmc* connections accepted for service at link i. Suppose *gmc* m, $1 \leq m \leq K$, has a transmission timing deadline of δ_m and a maximum number of packets over that time, which require a total time of U_m to transmit.[1] The timing deadline represents the maximum delay which the packet may experience before being transmitted. For

---

[1] This is equivalent to a task's *computational time* in the Liu and Layland's formulation.

*gmc*s this value is specified at call establishment time by the routing assignment table. We explain below how U_m is calculated for an arbitrary *gmc* m. Under EDF at any single point in time the packet from the *gmc* with the closest deadline is selected for transmission. In Liu and Layland's work a sufficient condition to guarantee the schedulability of K *gmc*s is that the following holds:

$$\sum_{m=1}^{K} \frac{U_m}{\delta_m} \leq 1 \tag{3}$$

The schedulability condition as expressed in Equation 3 makes several assumptions. For instance, it assumes that packet transmission is preemptive. However, packet transmission in high-speed networks is typically non-preemptive. Further, the schedulability test as expressed in equation 3 does not take into account the effect that multiple service queues with resource partitioning. Both the policy of not preempting packet transmission and the existence of multiple priority queues may lead to priority violations, whereby a higher priority packet's transmission may be delayed by the transmission of lower priority packets. It is therefore necessary to derive a new schedulability condition for guaranteed traffic. There are two potential sources of priority violations in determining guaranteed queue schedulability. One is from the non-preemptive EDF scheduling policy and the other from the unspecified inter-queue scheduling policy. The new schedulability test combines these two sources of priority violations into a single equation, which is expressed below.

To derive a schedulability test we can construct a scheduling model equivalent to a non-preemptive single service EDF system whose schedulability condition can be established using equation 3. The equivalent preemptive system is constructed so that packets are transmitted in the same order as in our link model. This is done by considering the scenarios under which packets from lower priority *gmc*s or from other service classes are transmitted even when packets from higher priority *gmc*s are available. Under this circumstance the higher priority packets are said to be *blocked*. There are two cases which may cause blocking. The first is when higher priority *gmc* packets arrive on the guaranteed queue during the time a lower priority guaranteed queue packet is being transmitted. The second case occurs when *gmc* packets arrive or are available for transmission during the time the inter-queue scheduling mechanism schedules any packets from the non-guaranteed queuing class.

In order to construct a preemptive system which transmits packets in the same order as our link model we must account for the the blocking effect. This is done by using *shadow activities*. The basic idea in using a shadow activity is to introduce conditions into the system used to block packet transmission from higher priority *gmc*s until the lower priority activities finish transmission. The shadow activities are therefore given the highest priority in the system. Assume that δ_1 is the smallest delay value of any *gmc* m. Next, let $\mathcal{S}_1$ represent the amount of time spent transmitting packets as a result of non-preemptive EDF scheduling which can result in priority violations, and let $\mathcal{S}_2$ represent the maximum amount of time the non-guaranteed service classes have to transmit. Now assume that the scheduling policy is preemptive. Then the schedulability condition for this preemptive system is given by:

$$\sum_{m=1}^{K} \frac{U_m}{\delta_m} + \frac{\mathcal{S}_1}{\delta_1} + \frac{\mathcal{S}_2}{\delta_1} \leq 1 \tag{4}$$

Notice that $\mathcal{S}_1$ models the case of a priority violation caused by lower priority packets from a *gmc* blocking the transmission of higher priority packets. Since the EDF intra-queue scheduler makes decisions after each packet transmission, this is given by:

$$\frac{\mathcal{S}_1}{\delta_1} = \frac{D_i}{\delta_1}$$

$\mathcal{S}_2$ may be modeled in a variety of ways, depending on the interqueuing policy. The only restriction that is maintained is that equation (2) is not violated. For instance, if the EDF queue always has priority and $P_i = 1$ than $\mathcal{S}_2 = 0$.

3 Future Work

We have presented a model for link-based EDF packet scheduler which incorporates the effect of other packet transmission policies. We are currently working on inter-queuing service policies which attempt to maximize the ability to the link to service guaranteed traffic without unduly penalizing the other traffic classes.

References

[1] Liu, C., and Layland, J. “Scheduling algorithms for multiprogramming in a hard real-time environment,” *JACM* 30(1), January 1973, pp. 46-61.

[2] Shenker, S., and Partridge, C., “Specification of guaranteed quality of service,” Internet draft draft-ietf-intserv-guaranteed-svc-03.txt, November 1995.

[3] Simon, R., *An Integrated Communication Architecture for Distributed Multimedia Applications*, Ph.D. thesis, Department of Computer Science, University of Pittsburgh, 1996.