

Comparison of Rate-Based Service Disciplines

Hui Zhang^{1,2}
Srinivasan Keshav²

TR-91-024

March, 1990

Abstract

This paper compares six new queue service disciplines that are 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.

¹This research has been supported in part by AT&T Bell Laboratories, the University of California under a MICRO grant, and the International Computer Science Institute. The views and conclusions in this document are those of the author and should not be interpreted as representing official policies, either expressed or implied, or any of the sponsoring organizations.

²Computer Science Division, University of California at Berkeley and International Computer Science Institute, 1947 Center Street, Suite 600, Berkeley, CA 94704.

1 Introduction

High speed networking introduces opportunities for new applications that have stringent performance requirements in terms of throughput, delay, delay jitter² and loss rate [Fer90]. Conventional packet switching data networks with window-based flow control and first-come-first-served service discipline are not able to provide services with 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.

A rate-based service discipline is one that provides clients with a minimum rate of service 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 switch³ in the network, manages the following resources: bandwidth, service priority and buffer space. Then, in conjunction with appropriate admission policies, clients can get performance guarantees in terms of throughput, delay, delay jitter and loss rate.

Several rate-based scheduling disciplines have been proposed recently: Delay-EDD [FV89][FV90b], Virtual Clock [Zha88], Fair Queueing⁴ [DKS90], the Stop-and-Go protocol [Gol90a][Gol90b], the Hierarchical Round Robin Discipline (HRR) [KKK90], and Jitter-EDD [VZF90]. We feel that it is instructive to compare and contrast these disciplines since they are closely related, yet they have some critical 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?

A similar comparison between Virtual Clock and Leaky Bucket [SLCG89] is presented in [WW90]. However we feel that comparing Leaky Bucket with Virtual Clock is inappropriate since Leaky Bucket is not a service discipline.

²Delay jitter is defined to be the maximum difference of delay experienced by any two packets [VZF90].

³In the literature, the notion of switch is used in the context of an ATM network, while 'gateway' or 'router' is used in the Internet environment. In this paper, we will uniformly call switching elements as 'switches'.

⁴Fair 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 [DKS90]

The guarantees discussed in this paper are deterministic guarantees as defined in [Fer90]: all the packets should meet the 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 overbooking resources in a controlled way, as proposed in [FV89][Gol91], we will not discuss those guarantees here.

2 Background

Guaranteed service requires that the network *protect* well-behaved users from malicious users and network load fluctuations. Malicious 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 channel⁵ 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 allow a channel to send packets at the higher rate so long as it won't affect the performance of other channels. Delay-EDD, Virtual Clock and Fair Queueing are in this category.
- Rate controlled service disciplines won't allow a channel to send packets at a higher rate under any circumstances. Stop-and-Go, HRR, Jitter-EDD are in this category.

A rate-based service discipline can also be classified as either work-conserving or non-work-conserving. In a work-conserving discipline, a server is never idle when there is a packet to send. In a non-work-conserving discipline, each packet is assigned an eligibility time (either explicitly or implicitly). Even when the server is idle, if no packets are eligible, none will be transmitted. As will be shown later in the paper, whether a service discipline is work-conserving or non-work-conserving affects both the buffer space requirements and the delay jitter characteristics.

Of the six disciplines considered in the paper, Delay-EDD, Virtual Clock and Fair Queueing are work-conserving, while Stop-and-Go, HRR and Jitter-EDD are non-work-conserving. Notice that all three rate-controlled service disciplines work are non-work-conserving. This

⁵There are different terms in the literature for the same or similar objects. A channel is also called a connection, a circuit, a conversation, a flow, etc..

is not just a coincidence: only non-work-conserving disciplines can achieve an upper bound on the service rate for a channel.

We now present intuitive descriptions of the six service disciplines. Technical details are presented later.

2.1 Fair Queueing

The aim of Fair Queueing [DKS90] is simple: if N channels share an output trunk, then each should get $1/N$ of the bandwidth, with the provision that if any channel uses less than its share, the slack is equally distributed among the rest. This can be achieved by doing a bit-by-bit round robin (BR) service among the channels. This is impractical, and so Fair Queueing tries to emulate BR. Each packet is given a finish number, which is the round number at which the packet would have received service, had the server been doing BR. By servicing packets in order of the finish numbers, it can be shown that Fair Queueing emulates BR. Channels can be given different fractions of the bandwidth by giving them weights; a weight corresponds to the number of bits of service the channel receives per round of BR service.

2.2 Virtual Clock

The Virtual Clock [Zha88] discipline aims to emulate the Time Division Multiplexing (TDM) service discipline in the same way as Fair Queueing emulates BR. Each packet is allocated a virtual transmission time, which is the time at which the packet would have been transmitted were the server actually doing TDM. By sending packets in virtual time order, Virtual Clock can be shown to emulate TDM.

2.3 Delay Earliest-Due-Date

In classic earliest-due-date (EDD) scheduling each packet is assigned a deadline, and the packets are sent in order on increasing deadlines. The Delay-EDD [FV89] service discipline is an extension where the server negotiates a service contract with each channel. The contract states that if a source obeys a peak and average sending rate, then the server will provide a delay bound. The key lies in the assignment of deadlines to packets. The server sets a packet's deadline to the time at which it should be sent had it been received according to the contract. This is just the expected arrival time added to the delay bound at the server. By reserving bandwidth at the peak rate, Delay-EDD can assure each channel a hard delay bound.

2.4 Jitter EDD

The Jitter-EDD discipline [VZF90] extends Delay-EDD to provide delay-jitter bounds (that is, a bound on the minimum as well on the maximum delay). After a packet has been served at each server, it is stamped with difference between its deadline and actual finishing time. A regulator at the entrance of the next switch holds the packet for this period before it is made eligible to be scheduled. This provides the required minimum and maximum delay guarantees.

2.5 Stop-and-Go

The Stop-and-Go service discipline [Gol90b] aims to preserve the 'smoothness' property of traffic as it traverses the network. Time is divided into frames. In each frame, only data that arrived at the server in the previous frame is sent. It can be shown that with this scheme, a packet receives both a minimum and a maximum delay as it goes through a channel. Since the delay and delay-jitter bounds are linked to the choice of the frame time, Stop-and-Go purposes for multiple frame sizes.

2.6 Hierarchical Round Robin

The Hierarchical Round Robin (HRR) [KKK90] server has several service levels, where each level provides round-robin service to a fixed number of slots. A channel is allocated some number of service slots at some level, and the server cycles through the slots at each level. The time a server takes to service all the slots at a level is called the frame time at that level. The key to HRR lies in its ability to give each level a fixed share of the bandwidth. Since a server completes one round through its slots once every frame time, it can provide a maximum delay bound to the channels allocated to that level.

3 Rate Based Service Disciplines

3.1 Traffic Specification

The traffic specifications in Virtual Clock, HRR and Stop-and-Go are essentially the same: a transmission rate(AR) averaged over an interval (AI). Fair Queueing was described in a datagram network context and no traffic specification was proposed. In this paper, we assume the same (AR, AI) specification applies to the Fair Queueing as well.

Delay-EDD and Jitter-EDD have a three parameter traffic specification: $Xmin$, which is

the minimal packet inter-arrival time, X_{ave} , which is an average packet inter-arrival time, and I , which is the interval over which X_{ave} is computed. The two parameters X_{min} and X_{ave} are used to characterize bursty traffic and to offer statistical guarantees[Fer90] (as opposed to deterministic guarantees).

In order to compare the three parameter EDD traffic models with the two parameter model for the other disciplines, we need to map the three parameter model to the two parameter model. How should this be done? In EDD disciplines a server uses X_{ave} to do bandwidth allocation, and X_{min} to ensure that delay bounds can be met. Thus, for the purpose of comparing the bandwidth allocations of the six disciplines, we can assume that X_{min} is equal to X_{ave} . Then, the following correspondence holds:

$$AR \iff \frac{1}{X_{min}}$$

3.2 Virtual Clock, Fair Queueing and Delay-EDD

In each of Virtual Clock, Fair Queueing and Delay-EDD, at each switch, there is a state variable associated with each channel to monitor and enforce the rate for that channel. In Virtual Clock, the variable is called auxiliary Virtual Clock ⁶ ($auxVC$); in Fair Queueing, it is called Finish Number(F); in Delay-EDD, it is called Expected Deadline(ExD)⁷.

In all three cases, $auxVC$, F and ExD are used as priority indices of packets; packets are served in the order of increasing $auxVC$, F or ExD values.

Virtual Clock	Fair Queueing	Delay-EDD
$auxVC_i^k \leftarrow \max(AT, auxVC_i^k)$	$F_i^k \leftarrow \max(R^k, F_i^k)$	$ExD_i^k \leftarrow ExD_i^k + X_{min_i}$
$auxVC_i \leftarrow auxVC_i + Vtick_i$	$F_i^k \leftarrow F_i^k + \frac{P_i}{n_i}$	$ExD_i^k \leftarrow \max(AT + d_i^k, ExD_i^k)$
Stamp packet with $auxVC_i^k$	Stamp packet with F_i^k	Stamp packet with ExD_i^k

Table 1

The computations of $auxVC$, F and ExD are shown in Table 1. The subscript i is the channel number, the superscript k is the switch number and d_i^k is the local delay bound⁸ assigned to the channel at channel establishment time. In Virtual Clock and delay-EDD, AT is the packet arrival time; in Fair Queueing, R^k is the number of rounds that have been completed for a hypothetical bit-by-bit round robin server at switch k , n_i is the weighting factor, and P_i is the packet length measured in number of bits.

⁶There is another variable *VirtualClock* associated with each connection that is periodically checked to measure the difference between the actual rate and the allocated rate. *VirtualClock* is irrelevant in this discussion.

⁷Notation used in the paper is summarized in Appendix

⁸The delay bound is the maximum difference between the arrival time and departure time of a packet [FV89].

The equivalence between Virtual Clock and Fair Queueing is complete, and can be established as follows:

$$\begin{aligned} AT &\Longleftrightarrow R^k \\ auxVC_i^k &\Longleftrightarrow F_i^k \\ Vtick_i &\Longleftrightarrow \frac{P_i}{n_i} \end{aligned}$$

Hence, we will not discuss Fair Queueing any further.

If we combine the two state equations for both Virtual Clock and Delay-EDD, we have

- In Virtual Clock,

$$auxVC_i^k \leftarrow \max(AT + Vtick_i, auxVC_i^k + Vtick_i)$$

- In Delay-EDD,

$$ExD \leftarrow \max(AT + d_i^k, ExD_i^k + Xmin_i)$$

We have the following mapping:

$$\begin{aligned} auxVC_i^k &\Longleftrightarrow ExD \\ Vtick_i &\Longleftrightarrow Xmin_i \text{ or } d_i^k \end{aligned}$$

Recall that in Delay-EDD, $Xmin_i$ is the minimum packet inter-arrival time and d_i^k is the local delay bound; in Virtual Clock, $Vtick_i$ is the average packet inter-arrival time. As can be seen, there are two differences here:

1. Delay-EDD imposes the restriction of minimum spacing between packets, while Virtual Clock does not.
2. Delay-EDD decouples the delay and bandwidth requirements by using both $Xmin_i$ and d_i^k , while Virtual Clock just has one counterpart $Vtick_i$.

The two differences are the reason why Delay-EDD, in conjunction with the establishment and admission control scheme described in [FV89], can provide both delay and bandwidth guarantees, while Virtual Clock can provide only bandwidth guarantees.

Virtual Clock does have a mechanism to handle priorities – the AT in the equation can be replaced by $AT - Priority$. In this case, packets from channels with higher *Priority* will get lower average delay, however, without imposing minimum spacing between packets and a priority allocation scheme similar to the delay bound allocation scheme in [FV89], Virtual Clock cannot provide deterministic delay bounds.

3.3 Delay-EDD and Jitter-EDD

The comparison between Delay-EDD and Jitter-EDD is shown in Table 2.

Delay-EDD	Jitter-EDD
$ExA_i^k \leftarrow \max(ExA_i^k + Xmin_i, AT)$	$ExA_i^k \leftarrow \max(ExA_i^k + Xmin_i, AT)$
$Ahead \leftarrow ExA_i^k - AT$	$Ahead \leftarrow \max(ExA_i^k - AT, PreAhead)$
$Deadline \leftarrow AT + Ahead + d_i^k$	Hold the packet for <i>Ahead</i> time
Stamp packet with <i>Deadline</i>	$Deadline \leftarrow AT + Ahead + d_i^k$
	Stamp packet with <i>Deadline</i>

Table 2

Notice that the state equations for Delay-EDD are written in a form different from that in Table 1. This is only for the sake of exposition. It can be shown that these two representations of the state equations are equivalent. The variable ExA_i^k is the expected arrival time of the next packet. It is related to ExD_i^k of Table 1 by:

$$ExD_i^k \iff ExA_i^k + Xmin_i$$

The variable *Ahead* is the amount of time the packet arrives ahead of schedule at the current switch.

As can be seen in Table 2, there are two differences between Delay-EDD and Jitter-EDD:

- There is one more term in calculation of *Ahead* in Jitter-EDD: *PreAhead*. *Preahead* is a value carried in from the previous switch; it tells the switch *how much the packet is ahead of schedule with respect to the previous switch*, i.e., the difference between the deadline and the actual finishing time of the packet in the previous switch.
- Jitter-EDD holds the packet for *Ahead* units of time, and then calculates its deadline and hands it to the scheduler, while Delay-EDD extends the deadline by *Ahead* time units and hands the packet immediately to the scheduler.

It is this holding time that makes Jitter-EDD a non-work-conserving discipline. It has been shown in [VZF90] that Jitter-EDD has the following properties:

1. The traffic pattern of a channel is preserved at each switch in spite of network load fluctuation; if the traffic from a channel obeys the $(Xmin, I)$ specification at the entrance of the network, then this traffic will also obey $(Xmin, I)$ at each switch.
2. Let switch 0 be the source, d_i^h be the local delay bound assigned to channel i at switch h and t_i^h be the service time of the packet at switch h . If a packet from channel i enters the network at time *EnterTime*, the earliest time it can arrive at switch k

is $(EnterTime + \sum_{h=0}^{k-2} d_i^h + t_i^{k-1})^9$; the earliest time and the latest time the packet can leave the switch are $(EnterTime + \sum_{h=0}^{k-1} d_i^h + t_i^k)$ and $(EnterTime + \sum_{h=0}^k d_i^h)$, respectively.

3. From property 2, it immediately follows that the maximum residence time of a packet on channel i at switch k is $d_i^{k-1} + d_i^k - t_i^{k-1}$.

Property 2 provides jitter bounds for channels[VZF90]; properties 1 and 3 give buffer bounds to prevent packet loss, which will be discussed in more detail later in Section 3.5. As will be shown in the next section, the other two non-work-conserving service disciplines, Stop-and-Go and HRR, have similar properties.

However, it should be noticed that the three work-conserving disciplines: Virtual Clock, Fair Queueing and Delay-EDD, do not have such properties. The following statements, that are in some sense parallel to properties 1 through 3 above, apply to Virtual Clock, Fair Queueing and Delay-EDD.

1. For all three service disciplines, even if the traffic pattern of a channel obeys an average rate bound at the entrance of the network, a switch may face a higher instantaneous input rate over that channel due to network load fluctuations.
2. For Delay-EDD, if a packet from channel i enters the network at time $EnterTime$, the earliest time it can arrive at switch k is $(EnterTime + \sum_{h=0}^{k-1} t_i^h)$; the earliest time and the latest time it can leave the switch are $(EnterTime + \sum_{h=0}^k t_i^h)$ and $(EnterTime + \sum_{h=0}^k d_i^h)$, respectively. Since Virtual Clock and Fair Queueing do not provide a worst-case delay bound, they do not have a similar property.
3. For Delay-EDD, the maximum residence time of a packet on channel i in switch k is $(\sum_{h=0}^k d_i^h - \sum_{h=0}^{k-1} t_i^h)$.

3.4 Stop-and-Go and HRR

Both Stop-and-Go and HRR use a multi-level framing strategy. In this section, one level of framing is examined before discussing multi-level framing

For Stop-and-Go and HRR, the time axis is divided into periods of some constant length T , each called a frame¹⁰. Bandwidth is allocated to each channel as a certain fraction of the frame time.

Stop-and-Go defines *departing* and *arriving* frames for each link. At each switch, the *arrivingframe* of each incoming link is mapped to the *departingframe* of the output

⁹For simplicity, the propagation delay is not included in this calculation

¹⁰The notation FT is used in HRR, while T is used in Stop-and-Go; we adopt T for both in this paper.

link by introducing a constant delay θ , where $0 \leq \theta < T$. All the packets from one *arriving frame* of an incoming link and going to output link l are put into the corresponding *departing frame* of l . According to the stop-and-go discipline: the transmission of a packet that has arrived on any link l during a frame f should always be postponed until the beginning of the next frame. From a non-work-conserving discipline's point of view, the packets arriving during a frame f of the output link are not eligible for transmission until the next frame. Within each frame, the service order of packets is arbitrary¹¹

One level HRR is equivalent to the non-work-conserving round robin service discipline. Each channel is assigned a fraction of the total available bandwidth, and will receive that amount of bandwidth in each frame if it has sufficient packets available for service. The server ensures that the assigned bandwidth is also the maximum service rate for that channel in each frame. This means that, in a frame, after providing the amount of bandwidth allocated to one channel, even if the server is available and there are more packets of that channel queueing for transmission, the packets will not be served until the next frame. From a non-work-conserving discipline's point of view, those extra packets are not eligible for transmission until the next frame. Within each frame, the service order of packets is arbitrary.

Statements similar to those we made above for Jitter-EDD are as follows for Stop-and-Go:

1. If the traffic over a channel obeys, at the entrance of the network, the average rate constraint, where the average interval is T , the traffic will obey the same constraint throughout the network¹².
2. If a packet from channel i enters the network at time $EnterTime$, the earliest time it can arrive at switch k is $EnterTime + \sum_{h=0}^{k-1} \theta_h + (k-2)T$; the earliest time and the latest time it can leave switch k are $(EnterTime + \sum_{h=0}^k \theta_h + (k-1)T)$ and $(EnterTime + \sum_{h=0}^k \theta_h + kT)$, respectively, where switch 0 is the source, and θ_h is the synchronization time between the input link and the outlink at switch h , $0 \leq \theta_h < T$.
3. From property 2, the maximum packet residence time in switch k is: $2T + \theta_k$.

Again, property 2 will give bounds on jitter, and properties 1 and 3 will give bounds on buffer space requirements.

For HRR we have:

1. If the traffic from a channel obeys, at the entrance of the network, the average rate constraint, the traffic will obey the same constraint at each switch. If the traffic rate

¹¹FIFO is used as the service discipline within each frame in [Gol90a][Gol90b]; however, as pointed out by the author, this is just for convenience and is not part of Stop-and-Go.

¹²This is called the (r, T) -smoothness property in [Gol90a][Gol90b], where r is the average rate and T is the frame time.

is higher at the entrance, the first switch will smooth out the traffic, and the traffic of the channel will obey the average rate constraint at downstream switches.

2. If a packet over channel i enters the network at time $EnterTime$, the earliest time it can arrive at switch k is $EnterTime + \sum_{h=0}^{k-1} t_i^h$; the earliest time and the latest time it can leave the switch are $(EnterTime + \sum_{h=0}^k t_i^h)$ and $(EnterTime + kT)$, respectively. Here, switch 0 is the source, t_i^h is the service time of a channel i packet at switch h .
3. The maximum packet residence time in a switch is $2T$.

Notice that property 2 here does not give a tight bound on the minimum delay that a packet can incur as in Jitter-EDD and Stop-and-Go. However, HRR does provide an upper bound on delay. Thus HRR provides delay bounds but not tight delay jitter bounds¹³. Also, in Jitter-EDD and Stop-and-Go, property 3 is derived from property 2, while in HRR property 3 holds for itself due to the non-work-conserving nature of the service discipline.

If the packet size is assumed to be constant, we have the following equation for both HRR and Stop-and-Go:

$$ServiceQuantum = \frac{PacketSize}{FrameTime}$$

where the service quantum is the minimum granularity of bandwidth allocation.

The delay bounds that can be provided at each switch by Stop-and-Go and HRR are $T + \theta (0 \leq \theta < T)$ and T , respectively. In the above equation, we can see that the service quantum and frame time cannot be decrease simultaneously, i.e., low delay bound and fine granularity of bandwidth allocation cannot be achieved simultaneously.

To get around this coupling between delay and service quantum, both HRR and Stop-and-Go propose the use of multiple frames with different frame times. In this case, it is possible to provide low delay bounds to some channels by putting them in frames with a smaller frame time, and to allocate bandwidth with fine granularity to other channels by putting them in levels with a larger frame time. However, the coupling between delay and service quantum still exists within each frame.

	Stop-and-Go	HRR
Delay Bound	$H * T < Q < 2 * H * T$	$2 * H * T$
Jitter Bound	$2 * T$	—

Table 3

In Multi-Frame-Time Stop-and-Go and Multi-Frame-Time HRR, it is possible that a channel is assigned to frames with different sizes in different switches. The three properties

¹³Delay jitter is defined as the maximum difference between delays experienced by any two packets in a channel. Jitter is defined differently in [KKK90]. Providing a delay jitter bound is sufficient to provide the jitter bound as defined in [KKK90], but the other direction does not hold.

discussed above for Stop-and-Go and HRR will hold if the channel is assigned to frames with same frame time at each switch. Under the same assumption, Table 3 summarizes the end-to-end delay characteristics of the channel for each of the two service disciplines. H and T are the number of hops of the channel and the frame time, respectively. The exact value of Q is $\sum_{h=0}^H \theta_h + HT$. When a channel is allocated different frame times at different switches, similar worst case delay bounds hold for both Stop-and-Go and HRR. However, the delay expressions are complicated, and so we will not discuss them here.

The most important difference between Stop-and-Go and HRR is that Stop-and-Go synchronizes the arriving frames of the input links and the departing frames of the output link at each switch. There are two implications:

1. By this synchronization, tight delay jitter bounds can be provided by Stop-and-Go.
2. The synchronization also means that in Multi-Frame-Time Stop-and-Go, the frame times of a channel along the path should be non-decreasing. HRR does not have this restriction. This gives HRR more flexibility in assigning channels with different frame times at different switches.

Another difference between Stop-and-Go and HRR is their response to ill-behaved gateways that allow a channel to send data at more than AR . A HRR server queues data per channel, so if a channel sends more data than it ought to, it can only hurt itself. On the other hand, a Stop-and-Go server has no way to prevent itself from being flooded, and malicious users (that evade flow regulation at the input to the network) could cause other users to drop packets.

3.5 Buffer Requirement

In this paper, we assume that the buffer space is statically allocated on a per channel basis to prevent packet loss even in the worst case.

The buffer space required for a channel can be calculated by multiplying the maximum rate packets can arrive by the longest residence time of a packet. This corresponds to the first and third properties discussed above for different service disciplines. Since the delay characteristics are unclear for Virtual Clock and Fair Queueing, only the other four disciplines are discussed here.

Delay-EDD	$\frac{\sum_{k=0}^K d_i^k}{X_{min,i}}$
Jitter-EDD	$\frac{d_i^{K-1} + d_i^K}{X_{min,i}}$
Stop-and-Go	$(2T_i + \theta_k) \times AR_i$
HRR	$2T_i \times AR_i$

Table 4

The amount of buffer space required for channel i at node k are given in table 4. As mentioned before, the following correspondence holds:

$$AR_i \iff \frac{1}{X_{min_i}}$$

If local delay is assumed to be the same for each switch in Delay-EDD and Jitter-EDD, the buffer space requirements for them are $\frac{(k+1)d_i}{X_{min_i}}$ and $\frac{2d_i}{X_{min_i}}$ respectively.

It can be seen that the buffer space requirements for the three non-work-conserving disciplines are almost constant for each node traversed by the channel, while the buffer space requirement for the work-conserving Delay-EDD increases linearly for each node along the path.

3.6 Admission Control Policies

In order to offer performance guarantees in terms of delay and bandwidth, resources need to be reserved at each node for each channel. Besides buffer space as described in the previous section, two other resources, schedulability (discussed below) and bandwidth, need to be managed. For the six service disciplines, if deterministic guarantees [Fer90] are to be offered, the overall bandwidth allocated to the channels should not exceed the bandwidth of the output links.

To provide delay guarantees, schedulability should also be considered. As pointed out in [FV89], scheduling saturation can occur even if bandwidth is not overbooked. For example, two packets with delay bounds of 4 time units and service time of 3 time units may arrive at a node at the same time. In this situation, it is impossible to meet the delay bounds for both packets. Stop-and-Go and HRR automatically handle this by multi-level framing, though this does not solve the problem of coupling between delay and bandwidth allocation granularity. In the case of EDD scheduling as in Delay-EDD, Jitter-EDD and possibly Virtual Clock with priority as discussed before, an explicit test has to be done to ensure scheduling saturation cannot happen[FV90b].

4 Implementation

Thus far, we have compared the service disciplines without regard to the implementation cost. This cost depends on whether the implementation is done in hardware or software, and the choice of data structure and algorithms. Thus, we will only sketch the steps necessary to implement the algorithms, and present some implementation issues.

4.1 Fair Queueing, Virtual Clock and Delay-EDD

The three algorithms have been shown to be quite similar, and their implementation is identical. In each case, when a packet arrives, per-channel state has to be retrieved and updated. Then, the packet is stamped with a priority index and placed in a priority queue. When the output trunk becomes free, it removes the head of the priority queue and sends it.

The efficient implementation of a priority queue for Fair Queueing has been investigated in [Kes91], and the results should hold for the other two disciplines as well.

An efficient hardware scheme to implement a priority queue is a dictionary machine ¹⁴. Here, a large number of processors (one per packet in the buffer) are placed at the leaves of a data distribution tree. Each processor stores a priority index and has a bidirectional data path to its two neighbors. A priority index is introduced at the root of the tree and asynchronously propagates to the leaves. Here, the processor compares the incoming value of the index with the one that it has stored, and shifts its stored value left or right. By coordinating these shifts, it is possible to accommodate the new index into one of the processors at the leaves. In this way, a single hardware step can update the priority queue. Deleting values from the tree is trivial.

4.2 Jitter-EDD

Jitter-EDD can be implemented in the same way as Delay-EDD with the addition of a regulator. A regulator has to delay a packet an additional amount of time in order to reconstruct the input traffic stream. It can do so by constructing lists of packets for each delay value and placing them in a calendar queue [Bro88]. This is described in [FV90a].

4.3 Stop-and-Go

Assuming the service order within each frame time is FIFO, Stop-and-Go can be implemented by a number of FIFO queues, a service controller, a transmission queue and a transmission server. Each queue corresponds to frame time and stores packets for channels that are allocated at that frame time. At the beginning of each frame time, the service controller is interrupted and it places the packets present in the corresponding queue (which have been collected in the previous frame time) at the *head* of the transmission queue. The server services packets from the head of the transmission queue. In this way, the service of the packets from a frame of larger frame time may be interrupted, on a non-preemptive basis, by the starting of a new frame of smaller frame time.

¹⁴Suggested by S. Rajagopal at UC Berkeley

The periodic interruption of the service controller requires a set of timers. No particular design has been suggested for the timers in the published literature. The major cost associated with the implementation is in the design of the timers and we conjecture that it can be implemented by a scheme resembling HRR. Another implementation issue is the synchronization of incoming and outgoing frames. This has to be achieved by synchronizing the output frames of all output links and introducing a fixed delay at each input link.

4.4 Hierarchical Round Robin

HRR is specifically designed to be implemented in hardware. HRR is implemented by three counters associated with each frame time. The counters are loaded at the beginning of each frame time. Transfer of control between levels is based on simple combinational logic operating on the values of the counters. The choice of the next packet to send involves only one step in the hardware, and can be done at very high speeds. Per-channel queueing of the data is done using custom VLSI (this has been built by C.R. Kalmanek and R.C. Restrict at AT&T Bell Laboratories, Murray Hill), and a prototype implementation of HRR has been designed to run at 1.3 Gbps.

5 Conclusions

In this paper, six rate-based service disciplines have been examined and compared in the context of connection-oriented packet switching networks. The emphasis has been on examining the mechanisms and the specific properties that can provide delay, delay jitter and bandwidth guarantees. Table 5 summarizes the results.

	VCL	FQ	D-EDD	SG	HRR	J-EDD
Work Conserving	Yes	Yes	Yes	No	No	No
Throughput Guarantee	Yes	Yes	Yes	Yes	Yes	Yes
Delay Guarantee	No	No	Yes	Yes	Yes	Yes
Jitter Guarantee	No	No	No	Yes	No	Yes
Constant Buffer	No	No	No	Yes	Yes	Yes

Table 5

We showed that the Virtual Clock and Fair Queueing disciplines are identical in their behavior, and that if the traffic is regular, they are both identical to Delay-EDD. While Fair Queueing and Virtual Clock cannot provide worst case delay bounds, Delay-EDD is able to do so because a) it restricts the peak arrival rate to be $1/X_{min}$, b) reserves bandwidth at this peak rate, and c) the admission control tests assure that neither bandwidth saturation nor scheduler saturation can occur.

We have shown that using a regulator at the input to a Delay-EDD server allows us to provide delay jitter bounds. The key idea in using a regulator is that it restores the traffic characterization to what it was when the traffic entered the network, which ensures that the network traffic load is *regular*. This makes the network easier to control. The synchronization of frames in Stop-and-Go achieves precisely the same purpose, which is why both the disciplines are able to provide delay jitter bounds. Since HRR does not restore traffic to its original form, delay jitter control is not possible, but bandwidth and delay guarantees can be made.

Both Stop-and-Go and HRR have the notion of frames. This simplifies the allocation of bandwidths, delays and jitters. However, a coupling is introduced between the frame size, the delay and the bandwidth allocation granularity. In both cases, multiple frame sizes are used to alleviate the coupling, but the solution is not complete. It would be interesting to see if it is possible to come up with a discipline that allows complete decoupling of these three things.

While work-conserving service disciplines are dominant in conventional networks, non-work-conserving service disciplines exhibit promising features that are suitable for offering guaranteed network performance. We feel that analysis and implementation of non-work-conserving disciplines is a promising area for future research.

This study limits its scope to smooth traffic specification, deterministic services and static buffer allocation scheme. Further work should study the behavior of different service disciplines when burstiness is introduced and statistical guarantees are offered.

6 Acknowledgements

We would like to thank Prof. D. Ferrari and the members of the Tenet group at UC Berkeley and International Computer Science Institute for their stimulating discussions and thoughtful criticism.

References

- [Bro88] Randy Brown. Calendar queues: A fast $O(1)$ priority queue implementation for the simulation event set problem. *Communications of the ACM*, 31(10):1220–1227, October 1988.
- [DKS90] Alan Demers, Srinivasan Keshav, and Scott Shenker. Analysis and simulation of a fair queueing algorithm. In *Journal of Internetworking Research and Experience*, pages 3–26, October 1990. Also in Proc. ACM SIGCOMM’89.

- [Fer90] Domenico Ferrari. Client requirements for real-time communication services, 1990. RFC 1193.
- [FV89] Domenico Ferrari and Dinesh Verma. A scheme for real-time channel establishment in wide-area networks. Technical Report TR-89-036, International Computer Science Institute, Berkeley, California, May 1989.
- [FV90a] Domenico Ferrari and Dinesh Verma. Quality of service and admission control in ATM networks. Technical Report TR-90-064, International Computer Science Institute, Berkeley, California, December 1990.
- [FV90b] Domenico Ferrari and Dinesh Verma. A scheme for real-time channel establishment in wide-area networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.
- [Gol90a] S. Jamaloddin Golestani. Congestion-free transmission of real-time traffic in packet networks. In *Proceedings of IEEE INFOCOMM'90*, pages 527–542, San Francisco, California, June 1990. IEEE Computer and Communication Societies.
- [Gol90b] S. Jamaloddin Golestani. A stop-and-go queueing framework for congestion management. In *SIGCOMM'90 Symposium, Communications Architecture & Protocols*, pages 8–18, Philadelphia Pennsylvania, September 1990. ACM SIGCOMM.
- [Gol91] S. Jamaloddin Golestani. Duration-limited statistical multiplexing of delay-sensitive traffic in packet networks, 1991. Preprint submitted to Infocom '91.
- [Kes91] Srinivasan Keshav. On the efficient implementation of fair queueing, 1991. Manuscript in revision for Journal of Internetworking Research and Experience.
- [KKK90] Charles R. Kalmanek, Hemant Kanakia, and Srinivasan Keshav. Rate controlled servers for very high-speed networks. In *IEEE Global Telecommunications Conference*, San Diego, California, December 1990.
- [SLCG89] M. Sidi, W.-Z. Liu, I Cidon, and I Gopal. Congestion control through input rate regulation. In *IEEE Global Telecommunications Conference*, December 1989.
- [VZF90] Dinesh Verma, Hui Zhang, and Domenico Ferrari. Guaranteeing delay jitter bounds in packet switching networks. In *First International Workshop on Network and Operating System Support for Digital Audio and Video*, Berkeley, California, November 1990. also in TriComm'91.
- [WW90] Abel Weinrib and L.T. Wu. Virtual clocks and leaky buckets: Flow control protocols for high speed networks. In *Second IFIP International Workshop on Protocols For High-Speed Networks*, Palo Alto, California, November 1990.
- [Zha88] Lixia Zhang. *A New Architecture for Packet Switched Network Protocols*. PhD dissertation, Massachusetts Institute of Technology, San Diego, California, July 1988.

A Notation

A.1 Virtual Clock

AR :	average rate.
AI :	average interval over which AR is computed.
$auxVC_i^k$:	auxiliary Virtual Clock of channel i at switch k .
$Vtick_i$:	the average packet inter-arrival time for channel i .
AT :	the arrival time of the packet

A.2 Fair Queueing

P_i :	the packet length measured in number of bits.
R^k :	the number of rounds that have been completed for a hypothetical bit-by-bit round robin server at switch k .
F_i^k :	the finish number of channel i at switch k .
n_i :	the relative weight of channel i .

A.3 Delay-EDD and Jitter-EDD

$Xmin_i$:	the minimum packet inter-arrival time of channel i .
$Xave_i$:	the average packet inter-arrival time of channel i .
t_i^k :	the maximum service time of a channel i packet at switch k .
I :	the interval over which $Xave$ is computed.
ExD_i^k :	the expected deadline of next packet of channel i at switch k .
ExA_i^k :	the expected arrival time of next packet of channel i at switch k .
d_i^k :	the local delay bound assigned to channel i at switch k .
$Ahead$:	the amount of time the packet arrives ahead of schedule at current switch.
$PreAhead$:	the amount of time the packet is transmitted ahead of schedule at the previous switch.

A.4 Stop-and-Go and HRR

T :	the frame time.
θ_k :	the amount of time between the end of the arriving frame of an incoming link and the beginning of the departing frame of an outgoing link.

