

Time-Shift Scheduling: Fair Scheduling of Flows in High Speed Networks

Jorge A. Cobb, Mohamed G. Gouda, Amal El Nahas

Abstract— We present a scheduling protocol, called Time-Shift Scheduling, to forward packets from multiple input flows to a single output channel. Each input flow is guaranteed a predetermined packet rate and an upper bound on packet delay. The protocol is an improvement over existing protocols because it satisfies the properties of rate-proportional delay, fairness, and efficiency, while existing protocols fail to satisfy at least one of these properties. In Time-Shift Scheduling, each flow is assigned an increasing timestamp, and the packet chosen for transmission is taken from the flow with the least timestamp. The protocol features the novel technique of time shifting, in which the scheduler's real-time clock is adjusted to prevent flow timestamps from increasing faster than the real-time clock. This bounds the difference between any pair of flow timestamps, thus ensuring the fair scheduling of flows.

Keywords—real-time network protocols, real-time scheduling, quality of service

1 Introduction

Consider a computer network with point-to-point communication channels. Assume a source computer wishes to transfer a sequence of packets to a destination computer. We call such a sequence a flow, that is, a flow is a sequence of packets generated by the same source and addressed to the same destination. Assume also that the source computer requires a lower bound on the rate at which its packets are forwarded through the network, and an upper bound on the packet delay from source to destination.

To solve this problem, a particular type of scheduling protocols, which we call rate-reservation protocols, were developed to forward packets from each flow at a designated rate. Examples of these protocols can be found in [26] [27]. In these protocols, the source of a flow finds a network path that leads to its desired destination. Then, it notifies each computer in the path about its desired packet rate. Each computer determines if it has enough available bandwidth in its output channel to forward the packets from the new flow. The new flow is accepted if and only if all computers in the path accept the new flow.

Due to the reservation of bandwidth, the network can provide service guarantees to each flow, such as end-to-end packet delays, provided the rate of the flow does not exceed the reserved rate. These service guarantees are of particular importance to real-time applications, such as interactive

audio and video [11].

The desirable properties of rate-reservation protocols are the following.

1. Rate-proportional delay: Let f be an input flow of a computer, with a reserved rate of $R.f$ bits/sec. Assume f is also the sole input to a constant rate server that forwards the bits of each packet of f at a precise rate of $R.f$ bits/sec. The delay of each packet of f through the computer should be at most the delay of the same packet through the constant rate server (plus a small constant).
2. Efficiency: The time to enqueue a received packet or to dequeue a packet for transmission is $O(\log(N))$, where N is the number of flows sharing the output channel.
3. Fairness: A flow should not be "punished" if it temporarily exceeds its reserved rate to take advantage of unused bandwidth in the channel. In addition, unused bandwidth should be shared among the flows in proportion to their reserved rates.

The rate-proportional delay property guarantees to each flow that the upper bound on its packet delay depends solely on its reserved rate, and not on other factors, such as the number of flows sharing the output channel or the reserved rate of other flows. The efficiency property is desirable due to the high bandwidth requirements expected from future applications of rate-reservation protocols.

The fairness property is desirable, because it may be normal for some flows to violate their reserved packet rate. Examples of such flows are file transfers and multi-resolution video [19]. The sources of these flows may reserve from the network the smallest packet rate necessary to receive a minimum quality of service. If the source of a flow detects that additional bandwidth is available, then it generates packets at a rate higher than its reserved rate, in order to take advantage of the unused bandwidth. If the source detects that no additional bandwidth is available, it reduces its sending rate. (There are several techniques by which a source can detect if additional bandwidth is available, see for example [17] and [23]). Thus, because some flows may be of adjustable rate, the unused bandwidth of a channel should be shared in a fair manner among all flows traversing the channel.

Some rate-reservation protocols are inadequate for adjustable rate flows because they are not work-conserving [10] [13] [18] [22]. That is, they serve each flow at exactly the rate it reserved, and will not forward additional packets of the flow even if the outgoing channel is idle. These

J. Cobb is with the Department of Computer Science, University of Houston, Houston, Texas, 77204-3475. E-mail: cobb@cs.uh.edu. M. Gouda is with the Department of Computer Sciences, The University of Texas at Austin, Austin, Texas 78712-1188. Email: gouda@cs.utexas.edu. Amal El-Nahas is with the Department of Computer Science, Alexandria University, Alexandria, Egypt.

protocols do not allow adjustable rate flows to take advantage of any additional bandwidth, and thus do not satisfy the fairness property.

Other rate-reservation protocols assign a timestamp to each packet, and packets are forwarded in increasing timestamp order. These protocols are work conserving, that is, the output channel is never idle as long as its packet queue is non-empty. However, some of these protocols are unfair [25], in the sense that they “punish” a flow if it sends packets at a rate higher than its reserved rate. Other protocols are fair [14] [21], but they are either inefficient or do not have rate-proportional delay.

In this paper, we introduce a new rate-reservation protocol, called Time-Shift Scheduling [4]. The protocol is based on flow timestamps (i.e., assigning a timestamp per flow rather than per packet), is work-conserving, and satisfies all the desirable properties mentioned above. Time-Shift Scheduling is based on the novel technique of time shifting, in which the real-time clock is periodically adjusted to prevent flow timestamps from increasing faster than the real-time clock. This ensures fairness by placing an upper bound on the difference between the timestamps of any pair of flows.

The paper is organized as follows. In Section 2, we discuss the use of flow timestamps in rate-reservation protocols. In Section 3, we present the Time-Shift Scheduler, and its formal definition is given in Section 4. We show that the Time-Shift Scheduler has rate-proportional delay in Section 5, and its fairness is shown in Section 6. In Section 7, we derive the end-to-end delay bounds for a path of Time-Shift Schedulers. Future work is given in Section 8.

Notation: Throughout the paper we use quantifications of the form:

$$(\otimes x : R(x) : B(x))$$

Above, \otimes is a commutative and associative operator, such as \max , \min , \sum (summation), \forall (conjunction), or \exists (disjunction). $R(x)$ is a boolean function defining the range of values for the dummy variable x , and $B(x)$ is a function defining the value given as an operand to \otimes . For example,

$$(\min x : 1 \leq x \leq 3 : x^2)$$

denotes the minimum of 1^2 , 2^2 , and 3^2 . If $R(x)$ is omitted, all values in the type of x are included.

2 Flow Timestamps

In this section, we discuss the strengths and weaknesses of the scheduling protocols of Virtual Clock [25], Weighted Fair Queuing [17] [20] [21], and Self-Clocking Fair Queuing [14]. Before doing so, we present some background on the scheduling technique used by these protocols.

A *computer network* consists of a set of computers interconnected via point-to-point bi-directional channels. A *flow* in a computer network is a potentially infinite sequence of packets generated by the same source and having the

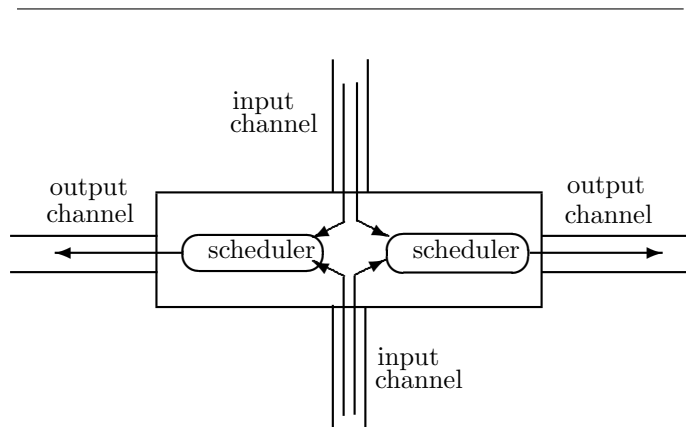


Figure 1: A computer with two input channels and two output channels.

same destination in the network. When a new flow wishes to join the network, the network finds a path from the source of the flow to the destination of the flow. Then, the network reserves a fraction of the packet rate of each channel along the path, and assigns this rate to the new flow. Finally, the source of the flow is given permission to generate packets at the rate reserved for it by the network. The chosen path and reserved rate of the flow remain fixed throughout the lifetime of the flow.

Each output channel of a computer is equipped with a *scheduler*, as shown in Figure 1. From the input channels, the scheduler receives packets from flows whose next hop to the destination is the output channel of the scheduler. Whenever its output channel becomes idle, the scheduler chooses a received packet and forwards the packet to the output channel. The rate at which the scheduler forwards the packets of a flow must be bounded from below by the reserved rate of the flow. To guarantee this minimum packet rate, the scheduler assigns a timestamp to each received packet. The timestamp is a function (among other things) of the flow’s reserved rate. The scheduler then forwards the packets in order of increasing timestamp.

The scheduler maintains a separate FIFO queue for the received packets from each flow. We say that a flow is *active* if its queue in the scheduler is non-empty.

We have shown in [3] that assigning a timestamp to a packet when it becomes the head of the queue of its flow is cleaner and more efficient than the original method of assigning a timestamp to the packet when it is received. We refer to the former technique as *flow timestamps*, since only one timestamp is maintained per flow. Thus, we base Time-Shift Scheduling on flow timestamps. To maintain a single timestamp paradigm throughout the paper, the flow timestamp versions of the Virtual Clock, Self-Clocking Fair Queuing, and Weighted Fair Queuing protocols will be discussed.

We adopt the following notation for a scheduler.

- N – number of flows in the scheduler.

- $queue.f$ – queue of received packets from flow f .
- $R.f$ – forwarding rate (in bits/sec.) reserved for flow f .
- $L.f$ – packet length (in bits) of the head of $queue.f$.
- $T.f$ – timestamp of flow f .
- $L.p$ – packet length (in bits) of packet p .
- $T.p$ – value of $T.f$ when packet p is at the head of $queue.f$, where f is the flow of packet p .
- L_{max} – upper bound on packet length for all flows.
- C – capacity in bits/sec. of the output channel.

We say that a packet is *forwarded to the output channel* by the scheduler when the first bit of the packet is being transmitted by the output channel. We say that a packet *exits the output channel* when the last bit of the packet is transmitted by the output channel. We say that a packet is *in the output channel* if it has been forwarded to the output channel but has not yet exited the output channel.

The goal of the scheduler is to forward the packets of each flow f at an average rate of at least $R.f$. Since all N flows share the output channel, the following constraint is necessary.

$$\sum_{f=0}^{N-1} R.f \leq C \quad (1)$$

Assume packet p from flow f is received. Before the packet is appended to $queue.f$, the scheduler checks if f is active. If f is not active, $T.f$ is updated as follows.

$$T.f := \max(v.p, T.f) + L.p/R.f$$

In this assignment, $v.p$ is some quantity related to packet p . The value chosen for $v.p$ varies from one scheduling protocol to another. On the other hand, if f is active, $T.f$ is not updated.

When the output channel becomes idle, the scheduler finds the active flow with the smallest timestamp. Let f be this flow. Then, the next packet from f is removed from its queue and forwarded to the output channel. If f remains active, its flow timestamp is updated as follows.

$$T.f := T.f + L.f/R.f$$

We next consider the case of Virtual Clock scheduling, which is defined by choosing $v.p$ as the time at which p is received by the scheduler. Virtual Clock has the property of rate-proportional delay, as proven in [3] [24]. In particular, the exit time of p is at most $T.p + L_{max}/C$. The scheduler is also efficient, requiring only $O(\log(N))$ operations to enqueue or dequeue a packet. However, the scheduler is unfair, as illustrated by the following well-known example.

Let packet sizes be constant, and the output channel have a rate of 1 packet/sec. The scheduler has two flows, f and g , each with a reserved rate of 1/2 packet/sec.. Consider the following sequence of events.

From time 0 up to time 100 sec., packets from flow f arrive at a rate higher than 1 packet/sec., and no packet is received from flow g . Thus, at time 100 sec., 100 packets

from f have been forwarded to the output channel, the queue of f is not yet empty, $T.f = 202$ sec., and $T.g = 0$.

At time 100 sec., packets from g arrive at a rate of at least 1 packet/sec., and packets from f continue to arrive. Note that when the next packet from flow g is received, $T.g = 102$ sec.. Since $T.f$, which equals 202 sec., is much larger than $T.g$, no packet from f is forwarded to the output channel until 50 packets from g are forwarded, i.e., until time 150 sec.. In effect, f is denied service for 50 seconds because it earlier took advantage of bandwidth unused by g .

This unfairness does not occur in the Weighted Fair Queuing protocol. Furthermore, the bound on packet delay is similar to that of Virtual Clock, and thus it also satisfies the rate-proportional delay property. In Weighted Fair Queuing, the timestamp of a packet is the time at which the packet would exit a virtual server. The input to the virtual server are the same input flows of the scheduler, and the virtual server shares its unused bandwidth among all active flows in proportion to their reserved rates. The value of $v.p$ is complicated and takes $O(N)$ time to compute.

The higher complexity of Weighted Fair Queuing led to the introduction of Self-Clocking Fair Queuing. In this scheduler, $v.p = T.q$, where q is the packet in the output channel at the time p is received. The time to enqueue or dequeue a packet is $O(\log(N))$, as it is in Virtual Clock.

This protocol is fair in the following sense. For any pair of active flows f and g ,

$$|T.f - T.g| \leq \max(L.f/R.f, L.g/R.g)$$

In this way, no flow can have a timestamp that is significantly greater than the timestamp of other flows. Thus, a flow that takes advantage of free bandwidth cannot be punished. On the other hand, in Virtual Clock scheduling, $|T.f - T.g|$ is unbounded.

Unfortunately, the packet delay increases as follows. Consider the same scheduler as above, with a fixed packet size and an output channel with rate 1 packet/sec.. Let the scheduler have 91 flows. Flows 1 through 90 have a rate of 1/100 packets/sec., and flow 0 has a rate of 1/10 packets/sec. Assume that one packet from each of flows 1 through 90 arrive at time 0. The timestamp of each of these packets is 100 sec.. Next, when the first of these packets is in the output channel, a packet from flow 0 arrives. The timestamp of this packet is 110 sec.. This packet must wait until all 90 packets with timestamp 100 sec. exit the output channel before itself may exit. Hence, the delay of this packet is 91 seconds, as opposed to a delay of at most 11 seconds that it would incur in Virtual Clock or Weighted Fair Queuing.

Note that the delay of flow 0 is related to the rate of the other flows. If the 90 flows with a rate of 1/100 packet/sec. are replaced by 900 flows with a rate of 1/1000 packet/sec., the delay of flow 0 increases by a factor of 10.

We conclude that each of the scheduling protocols discussed above satisfy only two out of the three desired properties. We next present a scheduling protocol that satisfies all three properties.




Figure 2: Drift of flow timestamps from the real-time clock.

3 Time-Shift Scheduling

We next present the intuition behind Time-Shift Scheduling. Its formal definition, its delay, and its fairness properties are given in later sections.

From its definition, $T.f$ can be viewed, intuitively, as the time at which the packet at the head of the queue of flow f should exit the output channel. That is, it should exit $L.f/R.f$ seconds later than the exit time of the previous packet from the same flow. Thus, we may define the “ideal arrival time” of the packet at the head of the queue of flow f as follows.

$$I.f = T.f - L.f/R.f$$

I.e., if the packet should exit at time $T.f$, and the flow is abiding by its reserved rate $R.f$, then, ideally, the packet should be received at time $I.f$.

Consider Figure 2, in which flows 0 through 2 are active, and the dashed line on the left indicates the current value of the real-time clock. Flow timestamp $T.f$ is an indicator of how much service has been given to flow f . If $T.f \gg \text{clock}$, then packets from f have been forwarded at a rate greater than $R.f$. This could occur if the flow is trying to take advantage of unused bandwidth. In the figure, all three flows have a flow timestamp significantly greater than the current value of the clock.

Assume that an additional flow, flow 3, has been inactive for some time, but becomes active once again. Assume we choose $v.p$ to be the arrival time of packet p . Then, when flow 3 becomes active, its flow timestamp is updated as follows.

$$T.3 := \max(\text{clock}, T.3) + L.3/R.3$$

Thus, $T.3 = \text{clock} + L.3/R.3$, which is significantly smaller than the flow timestamps of the other active flows. This implies that only packets from flow 3 are forwarded until $T.3$ reaches a value greater than the flow timestamps of the other flows.

To remedy this, $T.3$ should be given a value as close as possible to $T.0$, $T.1$, and $T.2$. Since $T.3$ is derived from the real-time clock, the clock should be close to $T.0$, $T.1$ and $T.2$. To do this, either we increase the value of the real-time

clock, or we reduce the value of each flow timestamp by an equal amount. We take the former approach, because the latter requires $O(N)$ time.

The remaining question is how much to advance the clock. If the clock is advanced beyond the minimum of the ideal arrival times, then the flow that becomes active has a timestamp larger than the timestamps of the active flows, and may be delayed excessively by these flows. If the clock is advanced to a value smaller than the minimum of the ideal arrival times, then the flow that becomes active has a timestamp smaller than the timestamps of the active flows, and it may unfairly delay the active flows. Thus, we choose to advance the clock to the minimum of the ideal arrival times of the active flows.

Because the scheduler advances the clock, i.e., it “shifts” the clock to the right, it is called a Time-Shift Scheduler. We refer to the adjustable real-time clock by the name ShiftClock¹

Whenever a flow becomes active, the scheduler performs a time shift by updating *ShiftClock* as follows:

$$\text{ShiftClock} := \max(\text{ShiftClock}, I_{\min})$$

where I_{\min} is the minimum ideal arrival time of all active flows, that is,

$$I_{\min} = (\min f : \text{queue}.f \neq \mathbf{empty} : I.f)$$

The restriction on the reserved rates given by (1) above is sufficient, provided the rate assigned to each flow remains fixed. However, this is not always the case, because the application generating packets for the flow may terminate, and the flow may be reassigned to a new application that reserves a different rate. Thus, a restriction is needed to indicate when can the reserved rate of a flow be reassigned to another flow.

The restriction we choose is the following. We say that a flow f is *live* if either $\text{ShiftClock} \leq T.f$ or f is active. The rate of a flow f can be reassigned to another flow if flow f is no longer live. Thus, the following is required to be an invariant of the reserved rates.

$$\left(\sum f : \text{ShiftClock} \leq T.f \wedge \text{queue}.f \neq \mathbf{empty} : R.f \right) \leq C \quad (2)$$

Next, we revisit the scenarios for Virtual Clock and Self-Clocking Fair Queuing discussed in Section 2, and see how unfairness and long delays are avoided in the case of Time-Shift Scheduling. Recall that in both scenarios the packet size is constant, and the output channel has a capacity of 1 packet/sec..

In the first scenario, we have two flows f and g , each with a reserved rate of 1/2 packet/sec.. From time 0 up to time 100 sec., packets from flow f arrive at a rate higher than 1 packet/sec., and no packet is received from g . Thus, at time 100 sec., 100 packets from f have been forwarded

¹If the clock of the scheduler cannot be advanced, the scheduler may increment a variable, say *shift*, instead of incrementing the clock. Any reference to the clock is then substituted by the expression $\text{clock} + \text{shift}$.

to the output channel, the queue of f is not empty, $T.f = 202$ sec., and $T.g = 0$.

At time 100 sec., packets from g arrive at a rate of 1/2 packet/sec. (or any rate higher than this), and packets from f continue to arrive. When the first packet of g arrives at time 100 sec., a time shift is performed, because g becomes active. Thus, $ShiftClock = I_{min} = I.f = 200$ sec., and $T.g$ is assigned 202 sec., which equals $T.f$. Therefore, from this point onward, one packet of g is forwarded for every packet of f , and the unfairness experienced with Virtual Clock is avoided.

In the second scenario, we have 100 flows. Flow 0 has a rate of 1/10 packets/sec., and flows 1 through 99 have a rate of 1/100 packets/sec.. At time 0, a packet from each of flows 1 through 99 is received, and the timestamp of each of these is 100 sec.. After the first of these packets is forwarded to the output channel, a packet from flow 0 arrives. Since flow 0 is becoming active, a time shift is performed, and $ShiftClock := \max(ShiftClock, I_{min})$. Since $I_{min} = 0$, and the first packet is still in the channel, $0 \leq ShiftClock \leq 1$. Hence, the timestamp of flow 0 is at most 11 sec.. Since the other flows have a timestamp of 100 sec., the next packet to be forwarded is from flow 0. Thus, this packet experiences a delay of at most two seconds, rather than the delay of 91 seconds that it would experience in Self-Clocking Fair Queuing.

4 Protocol Specification

We next provide a more formal description of a Time-Shift Scheduler. We define the behavior of a scheduler process by a set of global constants, a set local inputs, a set of local variables, and a set of actions. Actions are separated from each other with the symbol \square , using the following syntax:

begin action \square action \square ... \square action **end**

Each action is of the form *guard* \rightarrow *command*. A guard is either a boolean expression involving the local variables of its process, or a receive statement of the form **receive** p **from any** f , that receives a packet from any input flow. A command is constructed from sequencing ($;$) and conditional (**if fi**) constructs that group together **skip**, assignment, and statements of the form **forward** p , where p is a packet. Similar notations for defining network protocols are discussed in [15] [16].

An action in a scheduler process is said to be *enabled* if its guard is either a boolean expression that evaluates to **true**, or a receive statement of the form **receive** p **from any** f , and there is a packet that may be received from some input flow.

An execution step of a protocol consists of choosing any enabled action from the process, and executing the action's command. If the guard of the chosen action is a receive statement **receive** p **from any** f , then, before the action's command is executed, the packet is stored in variable p , and its flow number is stored in variable f . If the statement to execute in the command is of the form **forward** p , then packet p is forwarded to the output channel.

Protocol execution is fair, that is, each action that remains continuously enabled is eventually executed.

The specification of the Time-Shift Scheduler is presented in Figure 3. The process has two inputs from its environment. The first is a boolean bit which indicates if the output channel is currently idle. It becomes false when the scheduler forwards a packet p to the output channel, and becomes **true** $L.p/C$ seconds later. The second input is the rate reserved for each flow.

We assume the following. Variable *ShiftClock* is an adjustable real-time clock. It increases automatically with the progression of time. Also, executing an action takes zero time, i.e., *ShiftClock* remains constant while an action is executed unless an assignment statement in the action changes its value. Finally, *ShiftClock* does not advance while the packet queue is non-empty and the output channel is idle, i.e., the next packet to forward is chosen immediately after the output channel becomes idle.

The Time-Shift Scheduler process may be specified as follows.

The process contains three actions. In the first action, a packet is received from a flow. If the flow becomes active, then the flow timestamp is updated.

In the second action, when the output channel is idle and there are still packets to forward, the active flow with the smallest timestamp is obtained from function $\text{least}(T)$, and a packet from this flow is forwarded. The flow's timestamp is updated if the flow remains active. If no active flows remain, *ShiftClock* is updated so that it is greater than the flow's timestamp. This is necessary to prove fairness in Section 6.

In the third action, a time shift is performed, provided there is at least one active flow. Note that the scheduler has a lot of freedom in choosing when to perform a time shift, i.e., it could be done often or seldom. Regardless of when a time shift is performed, the property of rate-proportional delay is satisfied. However, if fairness is desired, the scheduler must execute a time shift whenever a flow becomes active. That is, immediately before executing the first action for an inactive flow f , the third action must be executed. We will show the correctness of these statements in the sections that follow.

Implementing this protocol requires two ordered queues: one for the flow timestamps, and one for the ideal service times. Inserting or removing an element from either of these takes $O(\log(N))$ time. Thus, the desired efficiency is achieved.

5 Local Delay Bound

In this section, we show that the Time-Shift Scheduler has a bound on packet delay no greater than the bound on packet delay of a Virtual Clock Scheduler or a Weighted Fair Queuing Scheduler.

Henceforth, any reference to time refers to the value of *ShiftClock* and not to the true value of real-time. E.g., the expression "at time t " refers to the state of the system

process Time-Shift Scheduler**inputs**

idle : is the output channel idle?
 R.f : rate of flow f

variables

p : packet
 $L.p$: length of packet p
 f, g : $0 \dots N-1$
 queue.f : packet queue of flow f
 $T.f$: timestamp of flow f
 ShiftClock : adjustable real-time clock

begin

```

receive  $p$  from any  $f \rightarrow$ 
  if queue.f = empty  $\rightarrow$ 
     $T.f := \max(\text{ShiftClock}, T.f) + L.p/R.f$ 
   $\square$  queue.f  $\neq$  empty  $\rightarrow$ 
    skip
  fi;
  queue.f := append(queue.f,  $p$ )
 $\square$ 
idle  $\wedge$  ( $\exists g ::$  queue.g  $\neq$  empty)  $\rightarrow$ 
 $f := \text{least}(T)$ ;
 $p := \text{head}(\text{queue}.f)$ ;
forward  $p$ ;
queue.f := tail(queue.f);
if queue.f  $\neq$  empty  $\rightarrow$ 
   $T.f := T.f + L.p/R.f$ 
 $\square$  ( $\forall g ::$  queue.g = empty)  $\rightarrow$ 
  ShiftClock := max(ShiftClock,  $T.f$ )
 $\square$  ( $\exists g ::$  queue.g  $\neq$  empty)  $\wedge$ 
  queue.f = empty  $\rightarrow$ 
  skip
fi
 $\square$ 
( $\exists g ::$  queue.g  $\neq$  empty)  $\wedge$  ShiftClock  $< I_{\min}$   $\rightarrow$ 
  ShiftClock :=  $I_{\min}$ 
end

```

Figure 3: Specification of the Time-Shift Scheduler.

when $\text{ShiftClock} = t$.

The bound on packet delay is based on the following theorem.

Theorem 1 In a Time-Shift Scheduler, for every active flow f ,

$$\text{ShiftClock} \leq T.f + L_{\max}/C - L.f/C$$

Proof

Consider an active flow f , and let p be the packet currently at the head of the queue of flow f . Let t be the time when p became the head of the queue of flow f , and let s be the latest time, no later than t ($s \leq t$), such that one of the following action executions occurred.

1. A time shift occurred, i.e., *ShiftClock* was increased, and $s = \text{ShiftClock}$ after the increase.
2. A packet q from a flow g , $f \neq g$, was forwarded, where $T.q > T.p$.
3. A packet was received when all the queues were empty.

In all three cases, at time s , for all active flows g ,

$$s \leq I.g \vee T.p < T.g \quad (3)$$

In case 1, this holds because after the time shift, $s = I_{\min} \leq I.g$. In case 2, since the packet chosen to be forwarded had $T.q > T.p$, then all active flows g have $T.g \geq T.q > T.p$. In case 3, the only active flow is the one whose packet is received at time s , and thus its ideal arrival time must be at least s .

Assume $T.p < s$. From (3), and $I.g < T.g$, for all active flows g at time s ,

$$T.p < T.g$$

Furthermore, when any flow g becomes active after s ,

$$T.p < s < T.g$$

Hence, $T.p < s$ is impossible, because flow f is active at time t , $s \leq t$, with $T.f = T.p$. We must have instead that $s \leq T.p$.

Let us first assume that no time shift occurs after s and before p is forwarded.

From the definition of s , only packets from flows g with $T.g \leq T.p$ are forwarded after s and until p is forwarded. Thus, from Lemma 1 below (replacing u by $T.p$) and from Relation (3), these packets, which include p , can be at most $(T.p - s) \cdot C$ bits. Since no time shift occurs after s , the last bit of packet p exits the output channel no later than time

$$s + (T.p - s) + L_{\max}/C,$$

that is, no later than time $T.p + L_{\max}/C$. The term L_{\max}/C is needed because for cases 1 and 2 we did not count the packet currently being in the output channel at time s .

If a time shift occurs after s and before p is forwarded, let s' be the time of the last time shift before p is forwarded. From the definition of a time shift, $s' \leq I.g$ for all active flows g at time s' . Also, since $s < s'$, only packets whose timestamps are at most $T.p$ are forwarded to the output channel until p is forwarded. Thus, from Lemma 1 (replacing u by $T.p$), these packets, which include p , can be at most $(T.p - s') \cdot C$ bits. Since no time shift occurs after s' , the packet p exits the output channel no later than time

$$s' + (T.p - s') + L_{\max}/C,$$

That is, no later than time $T.p + L_{\max}/C$. Again, the term L_{\max}/C is needed because we did not count the packet currently in the channel at time s' .

Recall that p is currently the head of the queue of flow f . Thus, $T.p = T.f$ until p is forwarded. Packet p will be forwarded to the output channel no later than time $T.f + (L_{\max} - L.f)/C$, and at this time either f becomes inactive, or $T.f$ increases, which implies the theorem. \square

Lemma 1 If $s \leq u$, and at time s , for all active flows f ,

$$s \leq I.f \vee u < T.f$$

then, starting from time s , the size of the packets forwarded to the output channel with a timestamp at most u have a total of at most $(u - s) \cdot C$ bits.

□

The proof of Lemma 1 is found in the appendix.

Theorem 1 implies that each packet p will exit the output channel before *ShiftClock* reaches the value of its timestamp, $T.p$, plus L_{\max}/C . It has been shown that the exit time of a packet in Virtual Clock scheduling is at most the packet's timestamp plus L_{\max}/C [3] [24]. This is not directly comparable with the above bound for Time-Shift Scheduling, because the exit time in the Time-Shift Scheduler is measured with respect to *ShiftClock*, and not with respect to the real-time clock.

To show that the delay bound of the Time-Shift Scheduler is no greater than the delay bound of Virtual Clock, we consider an alternative protocol as follows. Instead of shifting the clock forward δ seconds during a time shift, the timestamps of all flows (active or inactive) are reduced by δ seconds. Because the relative values of the flow timestamps with respect to each other and with respect to the clock is the same as in the Time-Shift Scheduler, the order in which packets are forwarded, and hence the delay, remains the same. Furthermore, it is easy to show that Theorem 1 holds for this alternative protocol. Finally, note that in this alternative protocol, the timestamp of each packet is at most the timestamp of the same packet in a Virtual Clock scheduler, because the flow timestamps are being reduced. Thus, the delay bound for a Time-Shift Scheduler is at most the delay bound for a Virtual Clock scheduler.

The delay bound of Weighted Fair Queuing is similar to that of Virtual Clock [20]. Hence, the delay bound of Time-Shift Scheduling is also similar to that of Weighted Fair Queuing.

6 Fairness

We next examine the fairness of a Time-Shift Scheduler. To ensure fairness, the protocol given above must be changed slightly. To keep the difference between the timestamps of any pair of flows as small as possible, the scheduler should perform a time shift often, in particular, before a flow becomes active. This can be accomplished by having the scheduler execute its third action (if enabled) immediately before receiving a packet from an inactive flow.

We refer to the above scheduler as a Fair Time-Shift Scheduler.

Definition 1 An active flow f is a *minimum serviced flow* (MSF) if $I.f = I_{\min}$.

For a Fair Time-Shift Scheduler, since a time shift is executed always before receiving a packet from an inactive flow, it is easy to show that I_{\min} increases monotonically.

Thus, if flow f is a MSF, it remains an MSF until the next packet from f is forwarded.

The fairness of a Fair Time-Shift Scheduler is based on the following theorem.

Theorem 2 In a Fair Time-Shift Scheduler, for any pair f and g of active flows, there exist an MSF m , such that

$$|T.f - T.g| \leq L.m/R.m + |L.f/R.f - L.g/R.g| + L_{\max}/C$$

Proof

We prove the theorem by showing that the Fair Time-Shift Scheduler preserves the following invariant. For every flow f , one of the following conditions hold.

1. f is active, and $I.m \leq I.f \leq T.m + L_{\max}/C$, for some MSF m .
2. f is inactive, and

$$T.f \leq \text{ShiftClock} \vee T.f \leq I_{\min}$$

3. f is inactive, and $I.m \leq T.f \leq T.m$ for some MSF m .

The proof that the scheduler satisfies this invariant is found in the appendix. We next show that the invariant implies the theorem. For the theorem, we only need case 1 of the invariant. However, cases 2 and 3 are needed to show that 1 holds when an inactive flow becomes active.

Consider two active flows f and g . From case 1 above, there exists two MSF flows, m and m' , such that

$$I.m \leq I.f \leq T.m + L_{\max}/C$$

$$I.m' \leq I.g \leq T.m' + L_{\max}/C$$

Without loss of generality, assume that $T.m' \leq T.m$. Since m and m' are MSF, $I.m = I.m'$. Thus, from the definition of ideal arrival time, we obtain the following upper bounds for $T.f$ and $T.g$.

$$I.m + L.f/R.f \leq T.f \leq I.m + L.m/R.m + L.f/R.f + L_{\max}/C$$

$$I.m + L.g/R.g \leq T.g \leq I.m + L.m/R.m + L.g/R.g + L_{\max}/C$$

The upper and lower bounds on $T.f$ and $T.g$ imply the theorem.

□

The bound on the relative value of two flow timestamps given by Theorem 2 prevents flows that become active from "hogging" the output channel and denying service to flows that have exceeded their reserved rates. The Virtual Clock protocol has no similar bound.

The above bound is close to, but not as tight as, the bound provided by the Self-Clocking Fair Queuing protocol. However, the fairness bound above is achieved in conjunction with a delay bound that is significantly better than the delay bound of Self-Clocking Fair Queuing.

The fairness of the scheduler implies that when a packet p arrives at the head of the queue of its flow, it will exit the output channel within a time bound that is independent of the value of its timestamp or the timestamps of other flows, as shown next.

Theorem 3 In a Fair Time-Shift Scheduler, the packet at the head of the queue of flow f will exit the output channel in at most

$$L.f/R.f + L.m/R.m + L_{\max}/C$$

seconds, for some MSF m .

Proof

Let p be the packet at the head of the queue of f , and it became the head of the queue at time t . Let $T.f'$ be the timestamp of flow f before p becomes the head of the queue. We have two choices regarding how p became the head of the queue.

First, assume p was received when the queue of f was empty. If $I.p$ is assigned *ShiftClock*, then from Theorem 1, p will exit the output channel in $L.p/R.f + L_{\max}/C$ seconds. If $I.p$ is assigned $T.f'$, then from part 3 in the proof of Theorem 2 (parts 1 and 2 do not apply), we have

$$I.m \leq T.f' = I.p \leq T.m \quad (4)$$

for some MSF m . Since I_{\min} is nondecreasing, only packets whose timestamps are in the range $[I.m, T.p]$ are forwarded after t , which from the proof of Lemma 1, are at most $(T.p - I.m) \cdot C$ bits. Furthermore,

$$\begin{aligned} & (T.p - I.m) \cdot C \\ &= (I.p + L.p/R.f - I.m) \cdot C \\ &\leq (T.m + L.p/R.f - I.m) \cdot C \\ &= (L.m/R.m + L.p/R.f) \cdot C \end{aligned}$$

Thus, p will exit the output channel in $L.m/R.m + L.p/R.f + L_{\max}/C$ seconds, where L_{\max}/C seconds comes from the time required for the packet forwarded before p to exit the output channel.

Assume now that p became the head of its queue when another packet from f was forwarded. Thus, $T.f' = I.p$. Let m be an MSF when p becomes the head of its queue. Thus, $I.m \leq I.p$. If $m = f$, then Relation (4) holds. If $m \neq f$, then m was not chosen to be forwarded, and hence $I.p = T.f' \leq T.m$, and Relation (4) holds.

Thus, the same relation of the previous case holds, and p will exit the output channel within $L.m/R.m + L.p/R.f + L_{\max}/C$ seconds. \square

7 End-to-End Delay Bounds


In this section, we present the end-to-end delay bound for a flow traversing multiple Time-Shift Schedulers. To do so, we borrow some results from Flow Theory [1, 2], which we overview next. The theorems are presented without proofs. The proofs may be found in [5].

A flow f is an infinite sequence $f.0, f.1, f.2, \dots$, of non-negative real numbers.

Informally, we divide time into very small and fixed sized intervals which we call instants. Each $f.i$ represents the

number of bits that travel in flow f at instant i . We denote the sequence $f.0, f.1, \dots, f.i$ by $f.(0, i)$.

As the flow traverses the network, it experiences queuing delays, which we represent with flow operators. A *flow operator* has an input flow f and an output flow g . At the i th instant, the operator inputs $f.i$ on its input flow, outputs $g.i$ on its output flow, and stores the remainder in an internal buffer. The content of the internal buffer at the i th instant is denoted $b.i$. The infinite sequence $b.0, b.1, b.2, \dots$, is called the *buffer flow* of the flow operator.



Formally, a flow operator with an input flow f , an output flow g , and a buffer flow b is defined, for every $i = 0, 1, 2, \dots$, as follows.

$$g.i \text{ is a value in the interval } F.(f.(0, i), b.(i - 1), i) \quad (5)$$

$$b.i = b.(i - 1) + f.i - g.i \quad (6)$$

where $b.(-1) = 0$, and F is a function, called the *operation* of the flow operator, that returns an interval of real numbers.

Flow operators cannot output more than which they have received, and thus the operation of a flow operator is required to ensure the following flow *conservation property*:

$$\begin{aligned} & \text{(For every input flow } f, \text{ output flow } g, \text{ and buffer} \\ & \text{flow } b, \text{ satisfying (5) and (6),} \\ & \text{(For every } i, g.i \leq f.i + b.(i - 1)) \\ & \text{).} \end{aligned}$$

The buffer capacity B of a flow operator with input flow f is the smallest non-negative real number that satisfies the following condition:

$$\begin{aligned} & \text{(For every output flow } g \text{ and buffer flow } b, \\ & \text{satisfying (5) and (6),} \\ & \text{(For every } i, b.i \leq B) \\ & \text{).} \end{aligned}$$

The delay D of a flow operator with input flow f is the smallest non-negative integer that satisfies the following condition:

$$\begin{aligned} & \text{(For every output flow } g \text{ and buffer flow } b, \\ & \text{satisfying (5) and (6),} \\ & \text{(For every } i, b_i \leq g.(i + 1) + \dots + g.(i + D)) \\ & \text{).} \end{aligned}$$

To represent the changes that occur to a flow as it traverses its path in the computer network, we introduce linear networks of flow operators.

A finite sequence of flow operators $\langle Op_0, Op_1, \dots, Op_{n-1} \rangle$ is a *linear network* iff for each i , $0 \leq i < n - 1$, the output flow of operator Op_i is the input flow of operator Op_{i+1} . The input flow of operator Op_0 is the input flow of the network and the output flow of operator Op_{n-1} is the output flow of the network.

The buffer flow c of a linear network $\langle Op_0, Op_1, \dots, Op_{n-1} \rangle$ is an infinite sequence $c.(-1), c.0, c.1, \dots$, such that $c.(-1) =$

0 and for every $i = 0, 1, 2, \dots$,

$$c.i = c.(i - 1) + f.i - g.i$$

where f is the input flow of the linear network and g is the output flow of the linear network.

This definition of the buffer flow of a network coincides with the definition of the buffer flow of a flow operator where the input flow is f and the output flow is g . Hence, we define the buffer capacity and delay of a linear network in the same manner as was done previously for flow operators. It can be shown that the buffer capacity of a linear network is at most the sum of the buffer capacities of each of the flow operators in the linear network. A similar relation holds for the delay of the linear network [5].

The first flow operator we introduce is the R -limiter, where R is a positive real number. The operation of an R -limiter is defined as follows,

$$g_i = \begin{cases} R & \text{if } b_{i-1} + f_i > R \\ b_{i-1} + f_i & \text{if } b_{i-1} + f_i \leq R \end{cases}$$

where f is the input flow, g is the output flow, and b is the buffer flow of the R -limiter. The R -limiter ensures that each element of its output flow is at most R . Basically, the R -limiter may be viewed as a constant-rate server that forwards its input flow to its output flow at exactly the rate R .


In Virtual Clock scheduling, the timestamp of a packet of a flow with reserved rate R is the time at which the packet would exit a constant-rate server that serves the flow at a rate R . A packet in Virtual Clock exits the output channel no later than the time indicated in its timestamp plus L_{\max}/C , that is, no later than the time at which it would exit an R -limiter plus L_{\max}/C .

Since the delay in a Time-Shift Scheduler is at most that of Virtual Clock, a packet in a Time-Shift Scheduler also exits no later than the time at which it would exit an R -limiter plus L_{\max}/C . Note that a packet may exit at a time much earlier than that at which it would exit an R -limiter.

This behavior can be represented by a flow operator called R -filter, where R is a positive real number. An R -filter is a flow operator that outputs an arbitrary value at each instant, with the following restriction. If an R -limiter and an R -filter have the same input flow, then the sum of any prefix of the output flow of the R -filter is at least the sum of the same length prefix of the output flow of the R -limiter. I.e., the R -filter may forward data faster than the R -limiter, and it is never allowed to lag behind the R -limiter. The formal definition of an R -filter may be found in [1, 5]

From the definition of an R -filter, each packet in its input flow exits no later than the time it would exit an R -limiter. Thus, an R -filter can be used to represent the behavior of a single Time-Shift Scheduler. However, in a Time-Shift Scheduler, a packet may exit up to L_{\max}/C seconds after it would exit an R -limiter. Furthermore, in an R -filter, the first bit of a packet of size L may exit up to L/R instants earlier than the last bit of the packet. These bits

should exit together, since packets are indivisible units of data and must be transmitted as a whole. Therefore, we represent a single Time Shift Scheduler with the following linear network, where L is the maximum packet size of the flow, and R is the reserved rate of the flow.



A d -delayer, where d is positive integer, is a flow operator that delays its input flow in an arbitrary manner by at most d instants. Its formal definition may be found in [1, 5]. The R -filter and d -delayer operators have the following useful properties.

Theorem 4

1. The linear network $\langle R\text{-filter}, R\text{-filter} \rangle$ is equivalent to a single R -filter.
2. The linear network $\langle d\text{-delayer}, d'\text{-delayer} \rangle$ is equivalent to a single $(d + d')$ -delayer.

□

Note that, by induction on part 1, any linear network of R -filters is identical to a single R -filter.

Theorem 5

1. When their input flow is the same, any output flow of the linear network $\langle d\text{-delayer}, R\text{-filter} \rangle$ is also an output flow of the linear network $\langle R\text{-filter}, d\text{-delayer} \rangle$.
2. The buffer capacity of $\langle R\text{-filter}, d\text{-delayer} \rangle$ is at most the buffer capacity of a single R -filter plus $d \cdot R$.

□

Assume a flow with reserved rate R and maximum packet size L traverses a path of n Time-Shift Schedulers. Let $L.i_{\max}$ and $C.i$ be, respectively, the upper bound on packet size and the capacity of the output channel of the i th scheduler along the path. The i th scheduler in this path is represented by the pair of flow operators $\langle R\text{-filter}, (L/R + L.i_{\max}/C.i)\text{-delayer} \rangle$. The whole path consists of a linear network with n of these pairs.

Using Theorem 4, and part 1 of Theorem 5, it is easy to show that any output flow of the above linear network is also an output flow of the linear network $\langle R\text{-filter}, D\text{-delayer} \rangle$, where

$$D = \sum_{i=0}^{n-1} (L/R + L.i_{\max}/C.i).$$

Thus, the delay of a path of Time-Shift Schedulers is at most the delay of an $\langle R\text{-filter}, D\text{-delayer} \rangle$ network. Since the delay of an R -filter is at most the delay of an R -limiter [5], the delay of the $\langle R\text{-filter}, D\text{-delayer} \rangle$ network is at most the delay of an R -limiter plus D .

The delay of a flow through an R -limiter depends on the "burstiness" of the flow. That is, how much does the flow temporarily exceed its reserved rate R . There are several ways to characterize the burstiness of a flow [2, 6, 13]. One

way to characterize burstiness is with the (m, R) -uniform property [2], which is defined as follows.

Let m be a positive integer and R be a positive real number. A flow f is (m, R) -uniform iff, for every $j = 0, 1, 2, \dots$,

$$\sum_{i=j}^{j+m-1} f.i \leq m \cdot R.$$

In this definition, R can be regarded as an upper bound on the rate of the flow, and m can be regarded as the interval over which the rate is averaged. The burstiness of the flow is measured by m : the larger m becomes, the burstier the flow may become.

The delay of an (m, R) -uniform flow through an R -limiter is at most m [5]. Thus, the total end-to-end delay along the path of n Time-Shift Schedulers is at most $m + D$.

From Theorem 4 and part 1 of Theorem 5, the linear network consisting of the first j $\langle R$ -filter, $(L/R + L.i_{\max}/C.i)$ -delayer \rangle pairs, i.e., the first j Time-Shift Schedulers, has a buffer capacity of at most the buffer capacity of the pair $\langle R$ -filter, D' -delayer \rangle , where

$$D' = \sum_{i=0}^{j-1} (L/R + L.i_{\max}/C.i)$$

Since the buffer capacity of the j th Time-Shift Scheduler is at most the buffer capacity of the first j Time-Shift Schedulers, then the buffer capacity of the j th scheduler is at most the buffer capacity of the pair $\langle R$ -filter, D' -delayer \rangle .

The buffer capacity of an R -filter is at most the buffer capacity of an R -limiter [5], and thus, from part 1 of Theorem 5, the buffer capacity of the j th Time-Shift Scheduler is at most the buffer capacity of an R -limiter plus $D' \cdot R$. For an (m, R) -uniform flow, the buffer capacity of an R -limiter is $m \cdot R$, and thus, the buffer capacity of the j th Time-Shift Scheduler is at most $m \cdot R + D' \cdot R$.

The upper bounds on the end-to-end delay and buffer capacities of a series of Time-Shift Schedulers derived in this section are the same upper bounds on the end-to-end delay and buffer capacities reported for a series of Virtual Clock or Weighted Fair Queuing schedulers [8, 12, 20]. Thus, the Time-Shift Scheduler achieves the same end-to-end delay and buffer capacities, while at the same time being fair and efficient.

8 Related and Future Work

In Time-Shift Scheduling, the end-to-end delay increases by $L/R + L.i_{\max}/C$ with each hop in the path to the destination. It is possible to decrease this delay for some flows at the expense of either increasing the delay of other flows or leaving some bandwidth unreserved. This has already been done in other protocols that do not ensure fairness [9] [10] [7]. In a future paper, we will describe a schedulability test, similar to the one presented in [28], that allows Time-Shift Scheduling to provide a delay bound indepen-

dent of the reserved rate of the flow, while at the same time ensuring the fairness property presented in this paper.

Note that in the proof of the upper bound on packet delay (Theorem 1), we made no assumption about how often does the scheduler perform a time shift. Furthermore, the theorem's proof can be relaxed to show that, during a time shift, if the scheduler assigns a value smaller than I_{\min} to the clock, then the theorem remains valid. Thus, the scheduler has a lot of freedom in manipulating the clock's value.

In a future paper, we will describe a whole family of protocols that assign timestamps to packets and shift the clock periodically. The distinguishing factor of each member of the family is the extent to which the clock is shifted. All members of the protocol family will satisfy the rate-proportional delay property. We will argue that, contrary to popular belief, the clock does not always need to increase at least as fast as real-time in order to maintain rate-proportional delay.

Acknowledgment

The authors would like to thank the referees for their constructive criticism which improved the quality of this paper.

References

- [1] Cobb J., Gouda M., "Flow Theory", *IEEE/ACM Transactions on Communications*, Vol. 5., No. 5, 1997.
- [2] Cobb J., Gouda M., "Flow Theory: Verification of Rate-Reservation Protocols", *IEEE International Conference on Network Protocols*, 1993.
- [3] Cobb J., Gouda M., El-Nahas A., "Flow Timestamps", *Annual Joint Conference on Information Sciences*, 1995.
- [4] Cobb J., Gouda M., El-Nahas A., "Time-Shift Scheduling: Fair Scheduling of Flows in High-Speed Networks", *IEEE International Conference on Network Protocols*, 1996.
- [5] Cobb, J., "Flow Theory and The Analysis of Timed-Flow Protocols", Ph.D. Thesis, The University of Texas at Austin, May 1996.
- [6] Cruz R.L., "A Calculus for Network Delay, Part I: Network Elements in Isolation", *IEEE Transactions on Information Theory*, Vol. 37, No. 1, January 1991, p. 114-131.
- [7] Figueira N., Pasquale J., "Leave-in-Time: A New Service Discipline for Real-Time Communications in a Packet-Switching Data Network", *SIGCOMM Conference*, 1995 p. 207.
- [8] Figueira N. R., Pasquale J., "An Upper Bound on Delay for the Virtual Clock Service Discipline", *IEEE/ACM Transactions on Networking*, Vol. 3, No. 4, Aug. 1995.
- [9] Ferrari D., Verma D., "A Scheme for Real-Time Channel Establishment in Wide-Area Networks", *IEEE*

Journal of Selected Areas in Communication, Vol. 8, No. 4, pp. 368-379, April 1990.

- [10] Ferrari D., Zhang H., "Rate Controlled Static Priority Queueing", *IEEE INFOCOM*, 1993.
- [11] Gall D., "A Video Compression Standard for Multimedia Applications", *Communications of the ACM*, Vol. 34, No. 4, April 1991.
- [12] Goyal P., Lam S., Vin H., "Determining End-to-End Delay Bounds in Heterogeneous Networks", *NOSS-DAV Workshop*, 1995.
- [13] Golestani S. J., "A Stop-and-Go Queueing Framework for Congestion Management", *ACM SIGCOMM conference*, 1990.
- [14] Golestani S. J., "A Self-Clocking Fair-Queueing Scheme for Broadband Applications", *IEEE INFOCOM*, 1994.
- [15] Gouda M., "Protocol Verification Made Simple", *Computer Networks and ISDN Systems*, Vol. 25, 1993, pp. 969-980.
- [16] Gouda M., *The Elements of Network Protocols*, textbook in preparation.
- [17] Keshav S., "A Control Theoretic Approach to Flow Control", *ACM SIGCOMM Conference*, 1991.
- [18] Kalmanek C. R., Kanakia H., and Keshav S., "Rate Controlled Servers for Very High-Speed Networks", *IEEE GLOBECOM*, 1990.
- [19] Kanakia H., Mishra P., "A Hop-by-Hop rate Based Congestion Control Scheme", *ACM SIGCOMM Conference*, 1992.
- [20] Parekh A. K. J., "A generalized Processor Sharing Approach to Flow Control in Integrated Services Networks", Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Technical Report Number LIDS-TH-2089.
- [21] Parekh A. K. J., Gallager R., "A generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case", *IEEE/ACM Transactions on Networking*, Vol. 1, No. 3, pp. 344-357, June 1993.
- [22] Philp I.R., Liu J.W.S., "End-to-End Scheduling in Real-Time Packet Switched Networks", *IEEE International Conference on Network Protocols*, 1996.
- [23] Ramakrishnan K. K., Raj J., "A Binary Feedback Scheme for Congestion Avoidance in Computer Networks", *ACM Transactions on Computer Systems*, Vol. 8, No. 2, pp. 158-181.
- [24] Xie G., Lam S., "Delay Guarantee of Virtual Clock Server", *IEEE/ACM Transactions on Networking*, Vol. 3, No. 6, December 1995.
- [25] Zhang L., "Virtual Clock: A New Traffic Control Algorithm for Packet-Switched Networks", *ACM Transactions on Computer Systems*, Vol. 9, No. 2, May 1991.
- [26] Zhang H., Keshav S., "Comparison of Rate-Based Service Disciplines", *ACM SIGCOMM Conference*, 1991.
- [27] Zhang H., "Service Disciplines for Guaranteed Performance Service in Packet Switching Networks", *Proceedings of the IEEE*, Vol. 83, No. 10, Oct.




Figure 4: Timestamps contained by each packet.

1995.

- [28] Zheng Q., Shin K. G., "On the Ability of Establishing Real-Time Channels in Point-to-Point Packet Switched Networks", *IEEE Transactions on Communications*, March 1994.

Lemma 1

If $s \leq t$, and at time s , for all active flows f ,

$$s \leq I.f \vee t < T.f$$

then, starting from time s , the size of the packets forwarded to the output channel with a timestamp at most t have a total of at most $(t - s) \cdot C$ bits.

Proof

We associate a timestamp with each bit b of a packet p from flow f . We allow b to be a real number in the interval $[0, L.p]$. For bit b in packet p , its timestamp equals $I.p + b/R.f$. Thus, the first bit has a timestamp equal to $I.p$, and the last bit has a timestamp equal to the packet timestamp, $T.p$. A timestamp u is contained by a packet p if a bit in the packet has timestamp u , i.e., $I.p \leq u \leq T.p$. A flow f contains timestamp u if some packet of f contains u .

For example, consider Figure 4. The ideal arrival time and timestamp of each packet are denoted by a small vertical line. In flow 1, for its first packet, p , we have $I.p = 1$ and $T.p = 3$. Similarly, we have: $I.q = 3$, $T.q = 4$, $I.r = 5$ and $T.r = 8$. Thus, all timestamp values in the intervals $[1, 4]$ and $[5, 8]$ are contained by flow 1.

We begin by showing that if timestamp u is contained by a packet p of some flow f , then f is live at time u , and furthermore, the value of $R.f$ used to compute $T.p$ is the same value that $R.f$ has at time u . We have the following two cases.

- If p arrived at a time at most u , we have two cases. If p is in the queue at time u , then f is live at u . Since $R.f$ cannot change while f is live, the value of $R.f$ at time u is used to timestamp p . If p was forwarded before time u , then, since $T.p \geq u$, f is live at time u , and $R.f$ cannot change until $ShiftClock > u$.
- If p arrives after time u , then $queue.f$ cannot become empty at any time starting at time u and until p arrives, since this would imply $I.p > u$, contradicting our choice of p . Hence, f is live at time u , and $R.f$

cannot change until p is received and forwarded.

From the above observation, for packets p and q in Figure 4, we have $I.p < T.p = I.q < T.q$, then both packets were assigned a timestamp using the same value for $R.1$. This is because flow 1 never ceases to be live from time $I.p$ up to $T.q$. Thus, for any interval $[u, v]$ contained by a flow f , the number of bits whose timestamps are in this interval are $R.f \cdot (v - u)$, where $R.f$ is the rate used to timestamp the packets containing the interval.

We next examine how many bits have a timestamp in the interval $[s, t]$. Starting from time s , find the latest time $s.0$, $s \leq s.0 \leq t$, such that no timestamp between s and $s.0$ is contained by a packet. In the figure, $s.0 = 1$ if we assume $s = 0$ and $t = 10$.

Next, find the largest timestamp $s.1$, $s.0 < s.1 \leq t$, such that if a flow does not contain a timestamp x in the interval $(s.0, s.1)$, then the flow does not contain any timestamps in the interval $(x, s.1)$. In Figure 4, $s.1$ equals 5, because flow 1 does not contain any value in the interval $(4, 5)$.

From the observations above, for any flow f that contains timestamps in $[s.0, s.1]$, at most $(s.1 - s.0) \cdot R.f$ bits contain timestamps in the interval $[s.0, s.1]$, where $R.f$ is the rate of flow f at time $s.0$, and furthermore, f is live at time $s.0$. From Relation (2), the sum of the rates of the live flows at time $s.0$ is at most C , and hence at most $(s.1 - s.0) \cdot C$ bits contain timestamps in the interval $[s.0, s.1]$.

We next find a timestamp $s.2$, whose relationship to $s.1$ is similar to the relationship between $s.1$ and $s.0$. The same argument shows that at most $(s.2 - s.1) \cdot C$ bits contain timestamps in the interval $[s.1, s.2]$. By repeating the above steps until we reach a timestamp $s.n$ where $s.n = t$, we have that the total number of bits whose timestamp is in the interval $[s, t]$ are at most $(t - s) \cdot C$.

Since we are given that for all active flows f at time s , $s \leq I.f \vee t < T.f$, then each packet p forwarded starting from time s with timestamp at most t has $I.p \geq s$, and hence the lemma holds. \square

Lemma 2

The following is an invariant of the Fair Time-Shift Scheduler. For every flow f , one of the following holds.

1. f is active, and $I.m \leq I.f \leq T.m + L_{\max}/C$, for some MSF m .
2. f is inactive, and

$$T.f \leq \text{ShiftClock} \vee T.f \leq I_{\min}$$

3. f is inactive, and $I.m \leq T.f \leq T.m$ for some MSF m .

Proof

We show that if the protocol is in a state satisfying any of the three cases above, then it will continue to satisfy the same case or satisfy one of the remaining two cases.

We begin with case 1.

Notice that once a flow m has $I.m = I_{\min}$, this continues to hold until flow m forwards a packet. That is, when a flow k becomes active, a time shift is performed, and $I_{\min} \leq \text{ShiftClock} \leq I.k$. Hence, $I.m = I_{\min} \leq I.k$. Because I_{\min} does not decrease, case 1 continues to hold until a packet

from either m or f is forwarded.

Assume that a packet from flow f is forwarded, and flow f is the only active flow. Note that $m = f$ in this case. If $\text{queue}.f$ becomes empty, ShiftClock is increased, if necessary, so that 2 holds. If $\text{queue}.f$ does not become empty, then 1 continues to hold with $m = f$.

Assume f is not the only active flow, and a packet from f is forwarded. Consider first that $f \neq m$ in 1. Because the packet from m was not chosen to be forwarded, $T.f \leq T.m$ before updating $T.f$. Also, we are given that $I.m \leq I.f$, and thus $I.m \leq I.f < T.f \leq T.m$. After $T.f$ is updated from the next packet in the queue of f ($I.f$ is now the previous value of $T.f$), we have $I.m \leq I.f \leq T.m$. Thus, 1 continues to hold. If no more packets from f remain, then $T.f$ is not updated, and thus 3 holds.

Assume next that $f = m$ in 1. Let k be an MSF after the packet from f is forwarded and $T.f$ is updated. If f is an MSF in this case, then 1 continues to hold with $f = m$. Otherwise, $k \neq f$, and because k was not chosen to be forwarded, $T.f \leq T.k$ before forwarding the packet from f , and thus $T.k \geq I.f$ after $T.f$ is updated. Furthermore, because k is an MSF, $I.k \leq I.f$. Thus, 1 continues to hold, with k replacing m .

The above assumes that packets from f remain and $T.f$ is updated. If this is not the case, then, because k was not chosen to be forwarded, $T.f \leq T.k$. If $I.k \leq T.f$, then case 3 holds with k replacing m . Otherwise, $T.f < I.k = I_{\min}$, and case 2 holds.

Assume now that a packet from m is forwarded, and $m \neq f$. Let k be an MSF after the packet is forwarded. Let $T.m'$ be the timestamp of m just before the packet is forwarded. Because k is an MSF, then $I.k \leq I.f$ after the packet is forwarded. If $k \neq m$, then, because k was not chosen to be forwarded, $T.m' \leq T.k$. If $k = m$, then, after the packet is forwarded, $T.m' = I.k < T.k$. Thus, from 1, from $T.m' \leq T.k$, and from k being an MSF, we have $I.k \leq I.f \leq T.k + L_{\max}/C$ after the packet is forwarded. Hence, 1 holds with k replacing m .

Consider now case 2.

Recall that I_{\min} and ShiftClock increase monotonically. Hence, case 2 can only be falsified if f becomes active, which we consider next.

If no flow is active and f becomes active, then case 1 holds with $m = f$. If other flows were active when the packet from f is received, then a time shift is performed immediately before receiving the packet from f . Let m be an MSF after the time shift. From Theorem 1 and from the time shift,

$$I.m \leq \text{ShiftClock} \leq T.m + L_{\max}/C - L.m/C \quad (7)$$

Let $T.f'$ be the value of the timestamp of f before the packet is received. Recall that $I.f$ is assigned $\max(\text{ShiftClock}, T.f')$. If $I.f$ is assigned ShiftClock , then case 1 holds by relation (7). If $I.f$ is assigned $T.f'$ (i.e., $\text{ShiftClock} < T.f'$), then from case 2 we have $\text{ShiftClock} < T.f' \leq I_{\min} = I.m$, which contradicts Relation (7), and hence will not occur.

Consider now case 3.

This case is affected if flow f becomes active or if a packet from flow m is forwarded. If flow f becomes active, then a time shift is performed immediately before receiving the packet from f , and Relation (7) holds. If $I.f$ is assigned *ShiftClock*, then case 1 holds from Relation (7). If $I.f$ is assigned the previous value of $T.f$, then from case 3,

$$I.m \leq I.f \leq T.m$$

which implies that case 1 holds.

If a packet from flow m is forwarded, then let k be an MSF after the packet is forwarded, and let $T.m'$ be the timestamp of m before the packet is forwarded. If $k \neq m$, then, since k was not chosen to be forwarded, $T.m' \leq T.k$. If $k = m$, then, after the packet is forwarded, $T.m' = I.k < T.k$. Thus, from case 3, $T.f \leq T.k$. If after the packet of m is forwarded, $I.k < T.f$, then case 3 holds. Otherwise, $T.f \leq I.k = I_{\min}$, and case 2 will hold. If no active flow remains after forwarding the packet from m , then *ShiftClock* is advanced such that $\text{ShiftClock} \geq T.m' \geq T.f$, and hence case 2 holds.

All cases of the invariant are either preserved or another case holds. Thus, the invariant is true.

□