

Efficient Fair Queuing using Deficit Round Robin

M. Shreedhar and George Varghese *

October 16, 1995

Abstract

Fair queuing is a technique that allows each flow passing through a network device to have a fair share of network resources. Previous schemes for fair queuing that achieved nearly perfect fairness were expensive to implement: specifically, the work required to process a packet in these schemes was $O(\log(n))$, where n is the number of active flows. This is expensive at high speeds. On the other hand, cheaper approximations of fair queuing that have been reported in the literature exhibit unfair behavior. In this paper, we describe a new approximation of fair queuing, that we call *Deficit Round Robin*. Our scheme achieves nearly perfect fairness in terms of throughput, requires only $O(1)$ work to process a packet, and is simple enough to implement in hardware. Deficit Round Robin is also applicable to other scheduling problems where servicing cannot be broken up into smaller units (such as load balancing) and to distributed queues.

1 Introduction

When there is contention for resources, it is important for resources to be allocated or scheduled fairly. We need firewalls between contending users, so that the “fair” allocation is followed strictly. For example, in an operating system, CPU scheduling of user processes controls the use of CPU resources by processes, and insulates well-behaved users from ill-behaved users. Unfortunately, in most computer networks there are no such firewalls; most networks are susceptible to sources that behave badly. A rogue source that sends at an uncontrolled rate can seize a large fraction of the buffers at an intermediate router; this can result in dropped packets for other sources sending at more moderate rates! A solution to this problem is needed to *isolate* the effects of bad behavior to users that are behaving badly.

An isolation mechanism called Fair Queuing [DKS89] has been proposed, and has been proved [GM90] to have nearly perfect isolation and fairness. Unfortunately, Fair Queuing (FQ) appears to be expensive to implement. Specifically, FQ requires $O(\log(n))$ work per packet, where n is the number of packet streams that are concurrently active at the gateway or router. With a large number of active packet streams, FQ is hard to implement¹ at high speeds. Some attempts have been made to improve the efficiency of FQ; however such attempts either do not avoid the $O(\log(n))$ bottleneck or are unfair. We will use the capitalized “Fair Queuing (FQ)” to refer to the implementation in [DKS89], and the uncapitalized “fair queuing” to refer to the generic idea.

*Washington University in St. Louis.

¹Alternately, while hardware architectures could be devised to implement FQ, this will probably drive up the cost of the router.

In this paper we shall define an isolation mechanism that achieves nearly perfect fairness (in terms of throughput), and which takes $O(1)$ processing work per packet. Our scheme is simple (and therefore inexpensive) to implement at high speeds at a router or gateway. Further we provide analytical results that do not depend on assumptions about traffic distributions.

Flows: Our intent is to provide firewalls between different packet streams. We formalize the intuitive notion of a packet stream using the more precise notion of a *flow* [Zha91]. A flow has two properties:

- A flow is a stream of packets which traverse the same route from the source to the destination and that require the same grade of service at each router or gateway in the path.
- In addition, every packet can be uniquely assigned to a flow using prespecified fields in the packet header.

The notion of a flow is quite general and applies to datagram networks (e.g., IP, OSI) and Virtual Circuit networks (e.g., X.25, ATM). For example, in a virtual circuit network a flow could be identified by a Virtual Circuit Identifier (VCI). On the other hand, in a datagram network, a flow could be identified by packets with the same source-destination addresses.² While source and destination addresses are used for routing, we could discriminate flows at a finer granularity by also using port numbers (which identify the transport layer session) to determine the flow of a packet. For example, this level of discrimination allows a file transfer connection between source A and destination B to receive a larger share of the bandwidth than a virtual terminal connection between A and B.

As in all fair queuing variants, our solution can be used to provide fair service to the various flows that thread a router, regardless of the way a flow is defined.

Organization: The rest of the paper is organized as follows. In the next section, we review the relevant previous work. A new technique for avoiding the unfairness of round-robin scheduling called *deficit round-robin* is described in Section 3. Round-robin scheduling [Nag87] can be unfair if different flows use different packet sizes; our scheme avoids this problem by keeping state, per flow, that measures the “deficit” or past unfairness. We analyze the behavior of our scheme using both analysis and simulation in Sections 4-5. Basic deficit round-robin provides throughput in terms of fairness but provides no latency bounds. In Section 6, we describe how to augment our scheme to provide latency bounds.

2 Previous Work

Existing Routers: Most routers use first-come-first-serve (FCFS) service on output links. In FCFS, the order of arrival completely determines the allocation of packets to output buffers. The presumption is that congestion control is implemented by the source. In feedback schemes for congestion control, connections are supposed to reduce their sending rate when they sense congestion. However, a rogue flow can keep increasing its share of the bandwidth and cause other (well-behaved) flows to reduce

²Note that a flow might not always traverse the same path in datagram networks, since the routing tables can change during the lifetime of a connection. Since the probability of such an event is low we shall assume that it traverses the same path during a session.

their share. With FCFS queuing, if a rogue connection sends packets at a high rate, it can capture an arbitrary fraction of the outgoing bandwidth. This is what we want to prevent by building firewalls between flows.

Typically routers try to enforce some amount of fairness by giving fair access to traffic coming on different input links. However, this crude form of resource allocation can produce exponentially bad fairness properties as shown below.

In Figure 1 for example, assume that all four flows $F1 - F4$ wish to flow through link L to the right of node D , and that all flows always have data to send. If node D does not discriminate flows, node D can only provide fair treatment by alternately serving traffic arriving on its input links. Thus flow $F4$ gets half the bandwidth of link L and all other flows combined get the remaining half. A similar analysis at C shows that $F3$ gets half the bandwidth on the link from C to D . Thus without discriminating flows, $F4$ gets $1/2$ the bandwidth of link L , $F3$ gets $1/4$ of the bandwidth, $F2$ gets $1/8$ of the bandwidth, and $F1$ gets $1/8$ of the bandwidth. In other words, the portion allocated to a flow can drop exponentially with the number of hops that the flow must traverse. This is sometimes called the *parking lot* problem because of its similarity to a crowded parking lot with one exit.

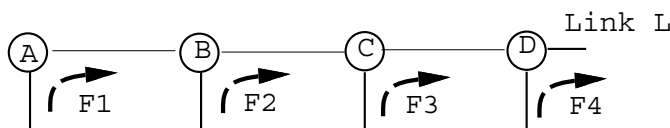


Figure 1: The parking lot problem.

Nagle’s solution: In Figure 1, the problem arose because the router allocated bandwidth based on input links. Thus at router D , $F4$ is offered the same bandwidth as flows $F1$, $F2$ and $F3$ combined. It is unfair to allocate bandwidth based on topology. A better idea is to distinguish flows at a router and treat them separately.

Nagle [Nag87] proposed an approximate solution to this problem for datagram networks by having routers discriminate flows, and then providing round-robin service to flows for every output link. Nagle proposed identifying flows using source-destination addresses, and using separate output queues for each flow; the queues are serviced in round-robin fashion. This prevents a source from arbitrarily increasing its share of the bandwidth. When a source sends packets too quickly, it merely increases the length of its own queue. An ill-behaved source’s packets will get dropped repeatedly.

Despite its merits, there is a flaw in this scheme. It ignores packet lengths. The hope is that the average packet size over the duration of a flow is the same for all flows; in this case each flow gets an equal share of the output link bandwidth. However, in the worst case, a flow can get $\frac{Max}{Min}$ times the bandwidth of another flow, where Max is the maximum packet size and Min is the minimum packet size.

Fair Queuing: Demers, Keshav and Shenker devised an ideal algorithm called *bit-by-bit round robin (BR)* which solves the flaw in Nagle’s solution. In the BR scheme, each flow sends one bit at a time in round robin fashion. Since it is impossible to implement such a system, they suggest approximately *simulating* BR. To do so, they calculate the time when a packet would have left the router using the BR

algorithm. The packet is then inserted into a queue of packets sorted on departure times. Unfortunately, it is expensive to insert into a sorted queue. The best known algorithms for inserting into a sorted queue require $O(\log(n))$ time, where n is the number of flows. While the BR guarantees fairness [GM90], the packet processing cost makes it hard to implement cheaply at high speeds.

A naive FQ server would require $O(\log(m))$, where m is the number of packets in the router. However Keshav [Kes91] shows that only one entry per flow need be inserted into a sorted queue. This still results in $O(\log(n))$ overhead. Keshav's other implementation ideas [Kes91] take at least $O(\log(n))$ time in the worst case.

Stochastic Fair Queuing (SFQ): SFQ was proposed by McKenney [McK91] to address the inefficiencies of Nagle's algorithm. McKenney uses hashing to map packets to corresponding queues. Normally, one would use hashing with chaining to map the flow ID in a packet to the corresponding queue. One would also require one queue for every possible flow through the router. McKenney, however, suggests that the number of queues be considerably less than the number of possible flows. All flows that happen to hash into the same bucket are treated equivalently. This simplifies the hash computation (hash computation is now guaranteed to take $O(1)$ time), and allows the use of a smaller number of queues. The disadvantage is that flows that collide with other flows will be treated unfairly. The fairness guarantees are probabilistic; hence the name *Stochastic Fair Queuing*. However, if the size of the hash index is sufficiently larger than the number of active flows through the router, the probability of unfairness will be small. Notice that the number of queues need only be a small multiple of the number of *active flows* (as opposed to the number of *possible* flows, as required by Nagle's scheme).

Queues are serviced in round robin fashion, without considering packet lengths. When there are no free buffers to store a packet, the packet at the end of the longest queue is dropped. McKenney shows how to implement this buffer stealing scheme in $O(1)$ time using bucket sorting techniques. Notice that buffer stealing allows better buffer utilization as buffers are essentially shared by all flows. The major contributions of McKenney's scheme are the buffer stealing algorithm, and the idea of using hashing and ignoring collisions. However, his scheme does nothing about the inherent unfairness of Nagle's round-robin scheme.

Other Relevant Work: Golestani introduced [Gol94] a fair queuing scheme, called Self-clocked Fair Queuing. This scheme uses a virtual time function which makes computation of the departure times simpler than in ordinary Fair Queuing. However, his approach retains the $O(\log(n))$ sorting bottleneck.

Together with weighted fair queuing, a pioneering approach to queue management is the Virtual Clock approach of Zhang [Zha91]. Delay bounds based on this queuing discipline have recently been discovered [FP95]. However, the approach still has the computational cost associated with sorting.

Van Jacobson and Sally Floyd have proposed a resource allocation scheme called Class-based queuing that has been implemented. In the context of that scheme, and independent of our work, Sally Floyd has proposed a queuing algorithm [Flo93a, Flo93b, FJ95] that is similar to our Deficit Round Robin scheme described below. Her work does not have our theorems about throughput properties of various flows; however, it does have interesting results on delay bounds and also considers the more general case of multiple priority classes.

A recent paper [SA94] has (independently) proposed a similar idea to our scheme: in the context of a specific LAN protocol (DQDB) they propose keeping track of remainders across rounds. Their

algorithm is, however, mixed in with a number of other features needed for DQDB. We believe that we have cleanly abstracted the problem; thus our results are simpler and applicable to a variety of contexts.

A paper by Parekh and Gallagher [PG93] showed that Fair Queuing could be used together with a leaky bucket admission policy to provide delay guarantees. This showed that FQ provides more than isolation; it also provides end-to-end latency bounds. While it increased the attractiveness of FQ, it provided no solution for the high overhead of FQ.

3 Deficit Round Robin

Ordinary round-robin servicing of queues can be done in constant time. The major problem, however, is the unfairness caused by possibly different packet sizes used by different flows. We now show how this flaw can be removed, while still requiring only constant time. Since our scheme is a simple modification of round-robin servicing, we call our scheme *deficit round-robin*.

We use Stochastic Fair Queuing to assign flows to queues. To service the queues, we use round-robin servicing with a quantum of service assigned to each queue; the only difference from traditional round-robin is that *if a queue was not able to send a packet in the previous round because its packet size was too large, the remainder from the previous quantum is added to the quantum for the next round*. Thus deficits are kept track off; queues that were shortchanged in a round are compensated in the next round.

In the next few sections, we will describe and precisely prove the properties of deficit round-robin schemes. We start by defining the figures of merit used to evaluate different schemes.

Figures of Merit: Currently, there is no uniform figure of merit defined for fair queuing algorithms. We define two measures: *FairnessMeasure* (which measures the fairness of the queuing discipline) and *Work* (which measures the time complexity of the queuing algorithm). Similar fairness measures have been defined before, but no definition of work has been proposed. It is important to have measures that are not specific to deficit round robin, so that they can be applied to other forms of fair queuing.

To define the work measure, we assume the following model of a router. We assume that packets sent by flows arrive to an Enqueue Process that queues a packet to an output link for a router. We assume there is a Dequeue Process at each output link that is active whenever there are packets queued for the output link; whenever a packet is transmitted, this process picks the next packet (if any) and begins to transmit it. Thus the work to process a packet involves two parts: enqueueing and dequeueing.

Definition 3.1 *Work is defined as the maximum of the time complexities to enqueue or dequeue a packet.*

For example, if a fair queuing algorithm takes $O(\log(n))$ time to enqueue a packet and $O(1)$ time to dequeue a packet, we say that the *Work* of the algorithm is $O(\log(n))$.

We will use a throughput fairness measure *FairnessMeasure* due to Golestani [Gol94], which measures the worst case difference between the normalized service received by different flows that are backlogged during any time interval. Clearly it makes no sense to compare a flow that is not backlogged with

one that is, because the former does not receive any service when it is not backlogged. If the fairness measure is very small, this amounts to saying that the service discipline closely emulates a bit-by-bit round robin server [DKS89], which is considered an ideal fair queueing system. Note that if the service discipline is idealized in a fluid-flow model to offer arbitrarily small increments of service, then *FairnessMeasure* becomes 0.

Definition 3.2 A flow is **backlogged** during an interval I of an execution if the queue for flow i is never empty during interval I .

We assume there is some quantity f_i , settable by a manager, which expresses the ideal share to be obtained by flow i . Let $sent_i(t_1, t_2)$ be the total number of bytes sent on the output line by flow i in the interval (t_1, t_2) . Fix an execution of the DRR scheme. We can now express the fairness measure of an interval (t_1, t_2) as follows. We define it to be the worst case (across all pairs of flows i and j that are backlogged during (t_1, t_2)), of the difference in the normalized bytes sent for flows i and j during (t_1, t_2) . More precisely:

Definition 3.3 Let $FairnessMeasure(t_1, t_2)$ be the maximum, over all pairs of flows i, j that are backlogged in the interval (t_1, t_2) , of $(sent_i(t_1, t_2)/f_i - sent_j(t_1, t_2)/f_j)$. Define *FairnessMeasure* to be the maximum value of $FairnessMeasure(t_1, t_2)$ over all possible executions of the fair queueing scheme and all possible intervals (t_1, t_2) in an execution.

Finally, we can define a service discipline to be fair if *FairnessMeasure* is a small constant. In particular, $FairnessMeasure(t_1, t_2)$ should not depend on the size of the interval [Gol94].

Algorithm: We propose an algorithm called *Deficit Round Robin* (Figure 2, Figure 3) for servicing queues in a router (or a gateway). We will assume that the quantities f_i , which indicate the share given to flow i ,³ are specified as follows. We assume that each flow i is allocated $Quantum_i$ worth of bits in each round. Define $Quantum = Min_i(Quantum_i)$. The share f_i allocated to flow i is simply $Quantum_i/Quantum$. Finally, since the algorithm works in rounds, we can measure time in terms of rounds. A round is one round-robin iteration over the queues that are backlogged.

Packets coming in on different flows are stored in different queues. Let the number of bytes sent out for queue i in round k be $bytes_i(k)$. Each queue i is allowed to send out packets in the first round subject to the restriction that $bytes_i(1) \leq Quantum_i$. If there are no more packets in queue i after the queue has been serviced, a state variable called *DeficitCounter_i* is reset to 0. Otherwise, the remaining amount $(Quantum_i - bytes_i(k))$ is stored in the state variable *DeficitCounter_i*. In subsequent rounds, the amount of bandwidth usable by this flow is the sum of *DeficitCounter_i* of the previous round added to $Quantum_i$. Pseudocode for this algorithm is shown in Figure 4.

To avoid examining empty queues, we keep an auxiliary list *ActiveList* which is a list of indices of queues that contain at least one packet. Whenever a packet arrives to a previously empty queue i , i is added to the end of *ActiveList*. Whenever index i is at the head of *ActiveList*, the algorithm services up to $Quantum_i + DeficitCounter_i$ worth of bytes from queue i ; if at the end of this service opportunity, queue i still has packets to send, the index i is moved to the end of *ActiveList*; otherwise, *DeficitCounter_i* is set to zero and index i is removed from *ActiveList*.

³more precisely, this is the share given to queue i and to all flows that hash into this queue. However, we will ignore this distinction until we incorporate the effects of hashing.

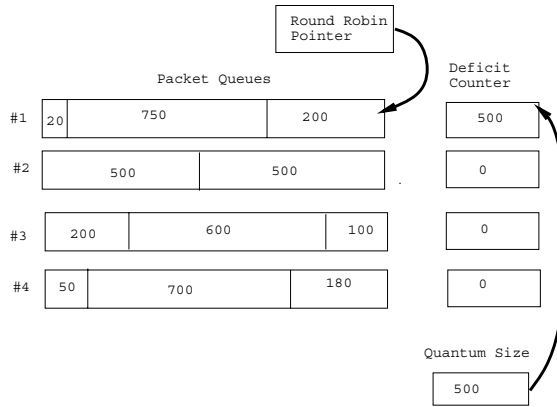


Figure 2: Deficit Round Robin: At the start, all the *DeficitCounter* variables are initialized to zero. The round robin pointer points to the top of the active list. When the first queue is serviced the *Quantum* value of 500 is added to the *DeficitCounter* value. The remainder after servicing the queue is left in the *DeficitCounter* variable.

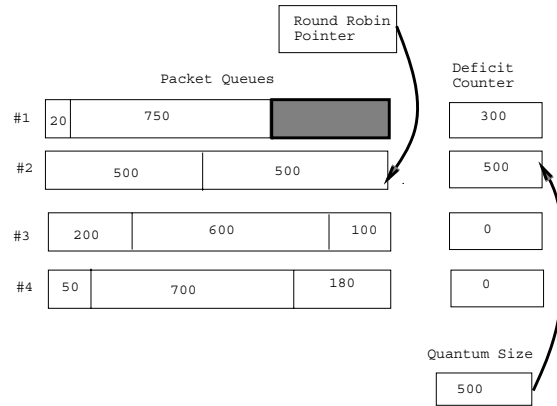


Figure 3: Deficit Round Robin (2): After sending out a packet of size 200, the queue had 300 bytes of its quantum left. It could not use it the current round, since the next packet in the queue is 750 bytes. Therefore, the amount 300 will carry over to the next round when it can send packets of size totaling 300 (deficit from previous round) + 500 (quantum).

In the simplest case $Quantum_i = Quantum_j$ for all flows i, j . Exactly as in Weighted Fair Queuing [DKS89], however, each flow i can ask for a larger relative bandwidth allocation and the system manager can convert it into an equivalent value of $Quantum_i$. Clearly if $Quantum_i = 2Quantum_j$, the manager intends that flow i get twice the bandwidth of flow j when both i and j are active.

Comparing DRR with BR The reader may be tempted to believe that DRR is just a crude approximation of BR. For instance, the reader may suspect that when the quantum size is 1 bit, the two schemes are identical. This plausible conjecture is incorrect.

Consider an example. Suppose $n - 1$ flows have large packets of size Max and the n -th flow is empty. Even with a quantum size of 1 bit, the deficit counter will eventually count up to $Max - 1$ (after $Max - 1$ scans of the $n - 1$ queues). Now assume that a small 1 bit packet arrives to the n -th flow queue. On the next scan, DRR will allow all other flows to send a maximum sized packet before sending the 1 bit

Consider any output link for a given router.
 $Queue_i$ is the i th queue, which stores packets with flow id i . Queues are numbered 0 to $(n - 1)$, n is the maximum number of output link queues.

$Enqueue()$, $Dequeue()$ are standard *Queue* operators.
 We use a list of active flows, *ActiveList*, with standard operations like *InsertActiveList*, which adds a flow index to the *end* of the active list.
 $FreeBuffer()$ frees a buffer from the flow with the longest queue using using McKenney's buffer stealing.
 $Quantum_i$ is the quantum allocated to $Queue_i$.
 $DeficitCounter_i$ contains the bytes that $Queue_i$ did not use in the previous round.

Initialization:

```
For ( $i = 0; i < n; i = i + 1$ )
   $DeficitCounter_i = 0;$ 
```

Enqueuing module: on arrival of packet p

```
 $i = ExtractFlow(p)$ 
If ( $ExistsInActiveList(i) == FALSE$ ) then
   $InsertActiveList(i);$  (*add  $i$  to active list*)
   $DeficitCounter_i = 0;$ 
If no free buffers left then
   $FreeBuffer();$  (* using buffer stealing *)
   $Enqueue(i, p);$  (* enqueue packet  $p$  to queue  $i$ *)
```

Dequeuing module:

```
While(TRUE) do
  If ActiveList is not empty then
    Remove head of ActiveList, say flow  $i$ 
     $DeficitCounter_i = Quantum_i + DeficitCounter_i;$ 
    while ( $(DeficitCounter_i > 0)$  and
      ( $Queue_i$  not empty)) do
       $PacketSize = Size(Head(Queue_i));$ 
      If ( $PacketSize \leq DeficitCounter_i$ ) then
         $Send(Dequeue(Queue_i));$ 
         $DeficitCounter_i = DeficitCounter_i$ 
           $-PacketSize;$ 
      Else break; (*skip while loop *)
    If ( $Empty(Queue_i)$ ) then
       $DeficitCounter_i = 0;$ 
    Else  $InsertActiveList(i);$ 
```

Figure 4: Code for Deficit Round Robin

packet of the n -th flow. Thus even with 1 bit quanta, the maximum delay suffered by a 1 bit packet (once it comes to the head of the queue) can be as bad as $(n - 1) * Max$, while in bit-by-bit, it can never be worse than $n - 1$ bit delays. Thus DRR is off by a multiplicative factor in delays and the two schemes are not identical.

4 Analytical Results

We begin with an invariant that is true for all executions of the DRR algorithm (not just for the backlogged intervals that are used to evaluate fairness). Recall that an invariant is meant to be true after every program action; it is not required to be true in the middle of a program action.

Lemma 4.1 *For all i , the following invariant holds for every execution of the DRR algorithm: $0 \leq DeficitCounter_i < Max$.*

Proof: Initially, $DeficitCounter_i = 0 \implies DeficitCounter_i < Quantum_i$. Notice that $DeficitCounter_i$ only changes value when queue i is serviced. During a round, when the servicing of queue i completes there are two possibilities:

- If a packet is left in the queue for flow i , then it must be of size strictly greater than $DeficitCounter_i$. Also, by definition, the size of any packet is no more than Max ; thus $DeficitCounter_i$ is strictly less than Max . Also, the code guarantees that $DeficitCounter_i \geq 0$.
- If no packets are left in the queue, the algorithm resets $DeficitCounter_i$ to zero.

□

The router services the queues in a round robin manner according to the DRR algorithm defined earlier. A round is one round-robin iteration over the queues that are backlogged.

We first show that during any period in which a flow i is backlogged, the number of bytes sent on behalf of flow i is roughly equal to $m \cdot Quantum_i$, where m is the number of round robin service opportunities received by flow i during this interval.

Lemma 4.2 *Consider any execution of the DRR scheme and any interval (t_1, t_2) of any execution such that flow i is backlogged during (t_1, t_2) . Let m be the number of round-robin service opportunities received by flow i during the interval (t_1, t_2) . Then:*

$$m \cdot Quantum_i - Max \leq sent_i(t_1, t_2) \leq m \cdot Quantum_i + Max.$$

Proof: We start with some definitions. Let us use the term *round* to denote service opportunities received by flow i within the interval (t_1, t_2) . Number these rounds sequentially starting from 1 and ending with round m . For notational convenience, we regard t_1 , the start of the interval, as the end of a hypothetical Round 0.

Let $DeficitCounter_i(k)$ be the value of $DeficitCounter_i$ for flow i at the end of round k . Let $bytes_i(k)$ be the bytes sent by flow i in round k . Let $sent_i(k)$ be the bytes sent by flow i in rounds 1 through k . Thus $sent_i(m) = \sum_{k=1}^m bytes_i(k)$.

The main observation (which follows immediately from the protocol) is: $bytes_i(k) + DeficitCounter_i(k) = Quantum_i + DeficitCounter_i(k - 1)$. We use the assumption that flow i always has a backlog in the above equation. Thus in round k , the total allocation to flow i is $Quantum_i + DeficitCounter_i(k - 1)$. Thus if flow i sends $bytes_i(k)$, then the remainder will be stored in $DeficitCounter_i(k)$, because queue i never empties during the interval (t_1, t_2) . This equation reduces to:

$$bytes_i(k) = Quantum_i + DeficitCounter_i(k - 1) - DeficitCounter_i(k).$$

Summing the last equation over m rounds of servicing of flow i we get a telescoping series. Since $sent_i(m) = \sum_{k=1}^m bytes_i(k)$ we get:

$$sent_i(m) = m \cdot Quantum_i + DeficitCounter_i(0) - DeficitCounter_i(m).$$

The lemma follows because the value of $DeficitCounter_i$ is always non-negative and $\leq Max$ (using Lemma 4.1). \square

The following theorem establishes the fact that the fairness measure for any interval is bounded by a small constant.

Theorem 4.3 *For an interval (t_1, t_2) in any execution of the DRR service discipline:*

$$FairnessMeasure(t_1, t_2) \leq 2Max + Quantum, \text{ where } Quantum = Min_i(Quantum_i)$$

Proof: Consider any interval (t_1, t_2) in any execution of DRR and any two flows i and j that are backlogged in this interval.

A basic invariant of the DRR algorithm (Figure 4) is that during any interval in which two flows i and j are backlogged, between any two round robin opportunities given to flow i , flow j must have had a round robin opportunity. This is easy to see because at the end of a flow i opportunity, the index i is put at the rear of the active queue. Since flow j is backlogged, index j is in the active queue, and thus flow j will be served before flow i is served again. Thus if we let m be the number of round robin opportunities given to flow i in interval (t_1, t_2) and if we let m' be the number of round robin opportunities given to flow j in the same interval, then $|m - m'| \leq 1$.

Thus from Lemma 4.2, we get:

$$sent_i(t_1, t_2) \leq m \cdot Quantum_i + Max.$$

Thus:

$$sent_i(t_1, t_2) \leq (m - 1) \cdot Quantum_i + Quantum_i + Max.$$

From the definition, f_i , the share given to any flow i , is equal to $Quantum_i/Quantum$. Thus we can calculate the normalized service received by i as:

$$sent_i(t_1, t_2)/f_i \leq (m - 1) \cdot Quantum + Quantum + Max/f_i,$$

since $Quantum_i = f_i \cdot Quantum$. Recall that $Quantum$ is the smallest value of $Quantum_i$ over all flows i . Similarly we can show for flow j (using Lemma 4.2) that:

$$sent_j(t_1, t_2) \geq m' \cdot Quantum_j - Max.$$

and so:

$$sent_j(t_1, t_2)/f_j \geq m' \cdot Quantum - Max/f_j,$$

Subtracting the equations for the normalized service for flows i and j , and using the fact that $m' \geq m - 1$, we get:

$$sent_i(t_1, t_2)/f_i - sent_j(t_1, t_2)/f_j \leq Quantum + Max/f_j + Max/f_i, \quad (1)$$

The theorem follows because both f_i and f_j are ≥ 1 for DRR. \square

Having dealt with the fairness properties of DRR, we analyze the worst-case packet processing work. It is easy to see that the size of the various $Quantum$ variables in the algorithm determines the number of packets that can be serviced from a queue in a round. This means that the latency for a packet (at low loads) and the throughput of the router (at high loads) is dependent on the value of the $Quantum$ variables.

Theorem 4.4 *The Work for Deficit Round Robin is $O(1)$, if for all i , $Quantum_i \geq Max$.*

Proof: Enqueuing a packet requires finding the queue used by the flow ($O(1)$ time complexity using hashing since we ignore collisions), appending the packet to the tail of the queue, and possibly stealing a buffer ($O(1)$ time using the technique in [McK91]). Dequeuing a packet requires determining the next queue to service by examining the head of *ActiveList*, and then doing a constant number of operations (per packet sent from the queue) in order to update the deficit counter and *ActiveList*. If $Quantum \geq Max$, we are guaranteed to send at least one packet every time we visit a queue, and thus the worst-case time complexity is $O(1)$. \square

Note that if we use hashing and we do not ignore collisions, then *Work* for DRR becomes $O(1)$ expected (as opposed to $O(1)$ worst case) because of possible collisions.

4.1 Comparison with other fair queuing Schemes

Golestani [Gol94] states the following result for the fairness measure of Self-clocked Fair Queueing (for any time interval (t_1, t_2)):

$$sent_i(t_1, t_2)/f_i - sent_j(t_1, t_2)/f_j \leq Max/f_j + Max/f_i,$$

On the other hand, Equation 1 shows that for DRR:

$$sent_i(t_1, t_2)/f_i - sent_j(t_1, t_2)/f_j \leq Quantum + Max_j/f_j + Max_i/f_i,$$

Thus the only difference in the fairness measure is the additive term of *Quantum* caused by DRR. Since we need $Quantum \geq Max$ to make the work complexity $O(1)$, this translates into an additive term of Max , assuming $Quantum = Max$.

Let us refer to the Demers-Keshav-Shenker scheme [DKS89] as DKS Fair Queuing. In summary: DKS Fair Queuing has a maximum value of *FairnessMeasure* of Max ; Self-clocked Fair Queuing has a maximum value of *FairnessMeasure* of $2Max$; and DRR fair queuing has a maximum value of *FairnessMeasure* of $3Max$. In all three cases, the fairness measure is a small constant that does not depend on the interval size, and becomes negligible for large intervals. Thus the small extra discrepancy caused by DRR in throughput fairness seems insignificant.

We compare the *FairnessMeasure* and *Work* of the major fair queuing algorithms that have been proposed until now in Figure 1. For this comparison only, assume that DRR does hashing but does not incorporate collisions. This is the only reasonable way to compare the algorithms because the trick of using hashing and ignoring collisions (as pioneered by [McK91]) can be applied to *all* fair queuing algorithms. At the same time, DRR can easily be modified to treat each flow separately (as opposed to treating all flows that hash into the same bucket equivalently). We will analyze the effect of ignoring collisions in the next subsection. We have also taken the fairness measure for round robin schemes as infinity. This is because if we consider two flows, one that uses large packets only, and a second that uses small packets only, then over any infinitely large interval, the first flow will get infinitely more service than the second. We have also taken the *Work* of Round Robin to be $O(1)$ expected, because even in ordinary round robin schemes, we need to look up the state for a flow, using say hashing.

From the table, Deficit Round Robin is the only algorithm that provides a fairness measure equal to a small constant and a *Work* of expected $O(1)$.

Table 1: Performance of fair queuing Algorithms

Algorithm	Fairness Measure	Work Complexity
Round Robin ([Nag87])	∞	$O(1)$ expected
Fair Queuing ([DKS89])	Max	$O(\log(n))$
Self-clocked Fair Queuing	$2Max$	$O(\log(n))$
Deficit Round Robin	$3Max$	$O(1)$ expected

4.2 Incorporating Hashing

In the previous analysis, we showed that if we did not lump together flows that hashed into the same bucket, then Deficit Round Robin (DRR) achieves a *FairnessMeasure* equal to $3Max$ and a *Work* = $O(1)$, expected. This is a reasonable way to compare DRR with other schemes (except Stochastic Fair Queuing) that do the same thing.

On the other hand, we have argued that implementations are likely to profit from the use of McKenney's idea of using hashing and ignoring collisions. The consequence of this implementation idea is that there is now some probability that two or more flows will collide; the colliding flows will then share the bandwidth allocated to that bucket.

The average number of other flows that collide with a flow can be shown [CLR90] to be $\frac{n}{Q}$, where n is the number of flows and Q is the number of queues. For example, if we have 1000 concurrent flows and 10,000 queues (a factor of 10, which is achievable with modest amounts of memory) the average number of collisions is 0.1. If B is the bandwidth allocated to a flow, the effective bandwidth in such a situation becomes $\frac{B}{1+\frac{n}{Q}}$. For instance, with 10,000 queues and 1000 concurrent flows, this means that two backlogged flows with identical quanta can differ in terms of throughput by 10 percent on the average, in addition to the additive difference of $3Max$ guaranteed by the fairness measure. Thus assuming a reasonably large number of queues, the unfairness caused by ignoring collisions should be small.

5 Simulation Results

We wish to answer the following questions about the performance of DRR.

- We would like to confirm experimentally that DRR provides isolation superior to FCFS as the theory indicates, especially in the backlogged case.
- The theoretical analysis of DRR is for a single router (i.e., 1 hop). How are the results affected in a multiple hop network?
- We want to confirm that the fairness provided by DRR is still good when the flows arrive at different (not necessarily backlogged) rates and with different distributions. Is the fairness sensitive to packet size distributions and arrival distributions?

Since we have multiple effects, we have devised experiments to isolate each of these effects. However, there are other parameters (such as number of packet buffers and Flow Index size) that also impact performance. We first did parametric experiments to determine these values before investigating our main questions. For lack of space we only present a few experiments and refer the reader to [S93] for more details.

5.1 Default Simulation Settings

Unless otherwise specified, the default for all the later experiments is as specified here. We measure the throughput in terms of delivered bits in a simulation interval, typically 2000 seconds.⁴

In the single router case (see Figure 5), there are one or more hosts. Each host has twenty flows, each of which generates packets at a Poisson average rate of 10 packets/ second. The packet sizes are randomly selected between 0 and Max (which is 4500 bits). Ill-behaved flows send packets at thrice the rate at which the other flows send packets (i.e., 30 packets/second). Each host in such an experiment is configured to have one ill-behaved flow.

In Figure 6, we show the typical settings in a multiple hop topology. There are hosts connected at each hop (a router) and each host behaves as as described in the previous section. In multiple hop

⁴Throughput is typically measured in bits/second. However, it is easy to convert our results into bits/second by dividing by the simulation time.

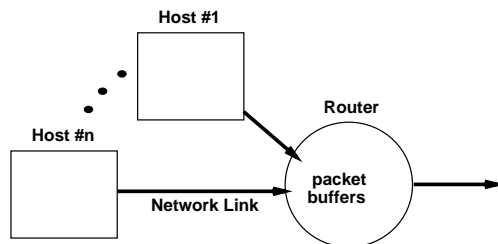


Figure 5: Single router configuration.

routes, where there are more than twenty flows through a router, we use large buffer sizes (around 500 packet buffers) to factor out any effects due to lack of buffers. In multiple hop topologies, all outgoing links are set at 1Mbps.

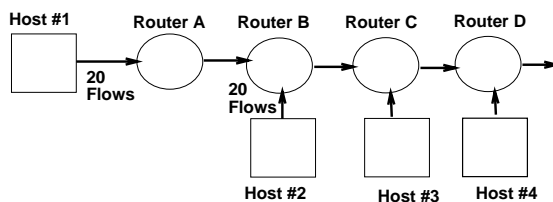


Figure 6: Multiple router configuration.

5.2 Comparison of DRR and FCFS

To show how DRR performs with respect to FCFS, we performed the following two experiments. In Figure 7, we use a single router configuration and one host with twenty flows sending packets at the default rate through the router. The only exception is that Flow 10 is a misbehaving flow. We use Poisson packet arrivals and random packet sizes (uniformly distributed between 0 and 4500 bits). All parameters used are the default settings. The figures shows the bandwidth offered to different flows using the FCFS and DRR queuing disciplines. We find that in FCFS the ill-behaved flow grabs an arbitrary share of bandwidth, while in DRR there is nearly perfect fairness.

We further contrast FCFS scheduling with DRR scheduling by examining the throughput offered to each flow at different stages in a multiple hop topology. The experimental setup is the default multiple hop topology described in Section 5.1. The throughput offered is measured at Router D (see Figure 8). This time we have a number of misbehaving flows. The figure shows that DRR behaves well in multihop topologies.

5.3 Independence with respect to Packet Sizes

We investigate the effect of different packet size on the fairness properties. The packet sizes in a train of packets can be modeled as random, constant or bimodal. We use a single router configuration with one host that has 20 flows. In the first experiment, we use random (uniformly distributed between 0 and 4500) packet sizes. In the next two experiments, instead of using the random packet sizes, we first use a constant packet size of 100 bits, and then a bimodal size that is either 100 or 4500 bits.

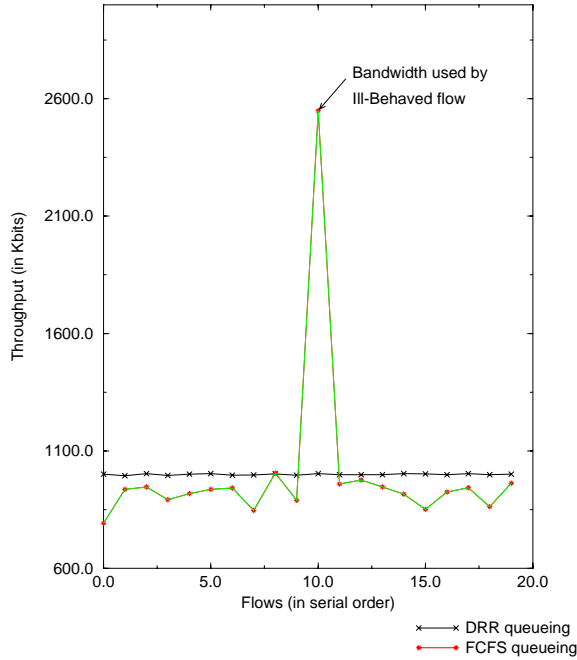


Figure 7: This is a plot of the bandwidth offered to flows using FCFS queuing and DRR. In FCFS, the ill-behaved flow (flow 10) obtains an arbitrary share of the bandwidth. The isolation property of DRR is clearly illustrated.

Figure 9 shows the lack of any particular pattern in response to the usage of different packet sizes in the packet traffic into the router. The difference in bandwidth offered to a flow while using the three different packet size distributions is negligible. The maximum deviation from this figure while using constant, random and bi-modal cases turned out to be 0.3%, 0.4699% and 0.32% respectively. Thus DRR seems fairly insensitive to packet size distributions.

This property becomes clearer when the DRR algorithm is contrasted with the performance of the SFQ Algorithm. While using the SFQ algorithm, flows sending larger packets consistently get higher throughput than the flows sending random sized packets; while all flows get equal bandwidth while using DRR.

It is clear from Figure 10 that DRR offers equal throughput to all kinds of sources. SFQ offers higher throughput to flows sending larger packets (for example, Flow Indices 44, 53, 863, 875 etc. which are flows originating from the “clever” host); this affects the throughput offered to normal flows (e.g. Flow Indices 60, 107, 190, etc. which are flows originating in the host sending random sized packets), which get substantially lower throughput.

5.4 Independence with respect to Traffic Patterns

We show that DRR’s performance is independent of the traffic distribution of the packets entering the router. We used two models of traffic generators: exponential and constant interarrival times respectively, and collected data on the number of bytes sent by each of the flows. We then examined the bandwidth obtained. The other parameters (e.g. number of buffers in the router) are kept constant at sufficiently high values in this simulation.

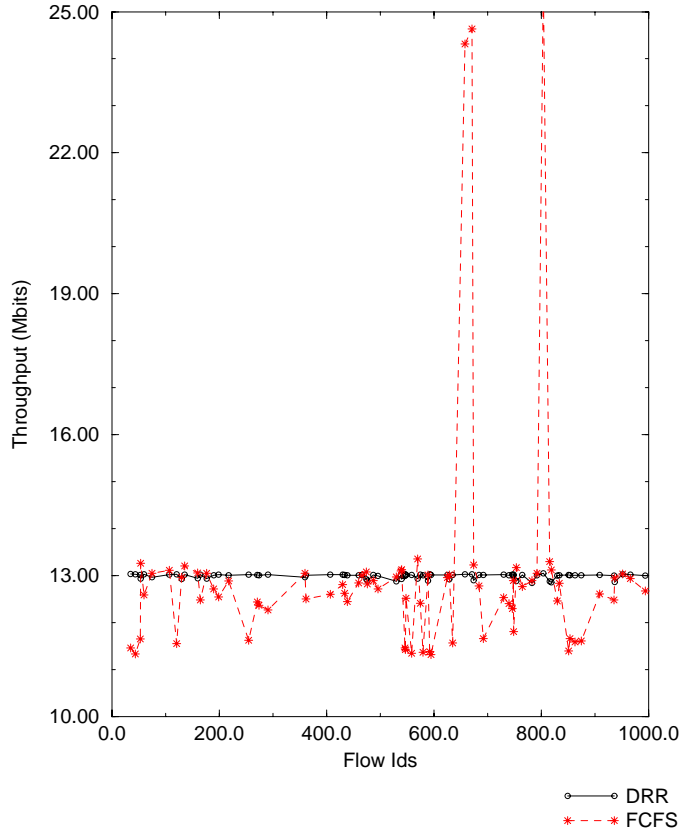


Figure 8: The bandwidth allocated to different flows at Router D using FCFS and DRR to schedule the departure of packets.

The experiment used a single router configuration with default settings. The outgoing link bandwidth was set to 10Kbps. Therefore, if there are 20 input flows each sending at rates higher than 0.5Kbps, there is contention for the outgoing link bandwidth. We found almost equal bandwidth allocation for all flows. The maximum deviation from the average throughput offered to all flows in the Constant traffic sources case is 0.3869%. In the Poisson case it is bounded by 0.3391% from the average throughput. Thus, DRR appears to work well regardless of the input traffic distributions.

6 Latency Requirements

Consider a packet p for flow i that arrives at a router. Assume that the packet is queued for an output link instantly and there are no other packets for flow i at the router. Let s be the size of packet p in bits. If we use bit-by-bit round-robin then the packet will be delayed by s round-robin rounds, Assuming that there are no more than n active flows at any time, this leads to a latency bound of $n * s / B$, where B is the bandwidth of the output line in bits/sec. In other words, a small packet can only be delayed by an amount proportional to its own size by every other flow. The DKS (Demers-Keshav-Shenker) approximation only adds a small error factor to this latency bound.

The original motivation in both the work of Nagle and DKS was the notion of isolation. Isolation is essentially a throughput issue: we wish to give each flow a fair share of the overall throughput. In terms

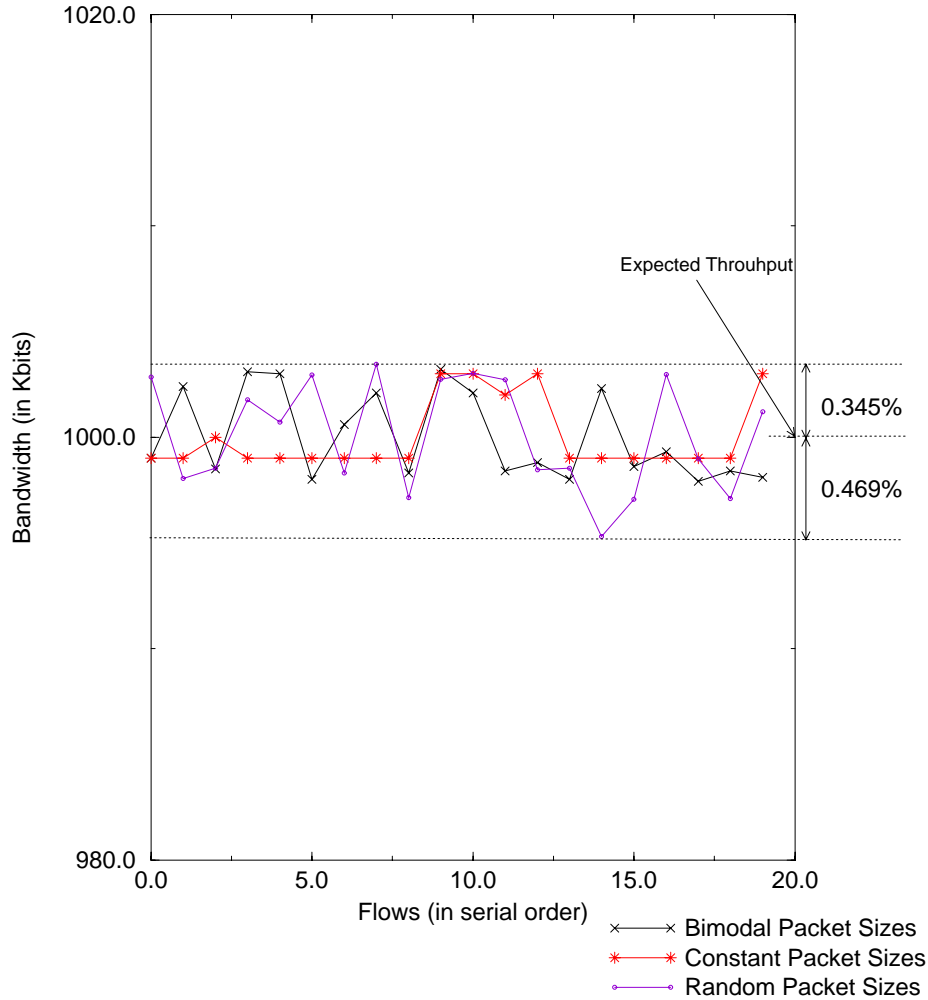


Figure 9: The bandwidth offered to different flows with exponential interpacket times and constant, bimodal and random packet sizes.

of isolation, the proofs given in the previous sections indicate that Deficit Round Robin is competitive with DKS Fair Queuing. However, the additional latency properties of BR and DKS have attracted considerable interest. In particular, Parekh and Gallager [PG93] have calculated bounds for end-to-end delay, assuming the use of DKS Fair Queuing at routers and token bucket traffic shaping at sources.

At first glance, Deficit Round Robin (DRR) fails to provide strong latency bounds. In the example of the arrival of packet p given above, the latency bound provided by DRR is $\frac{\sum_i^n \text{Quantum}_i}{B}$. In other words, a small packet can be delayed by a quantum's worth by every other flow. Thus in the case where all the quanta are equal to Max (which is needed to make the work $O(1)$), the ratio of the delay bounds for DRR and BR is Max/Min .

We note that the counterexample given in Section 3 indicates that *reducing the quantum size does not help improve the worst-case latency bounds for DRR*.

However, the real motivation behind providing latency bounds is to allow real-time traffic to have predictable and dependable performance. Since most traffic will consist of a mix of best-effort and real-

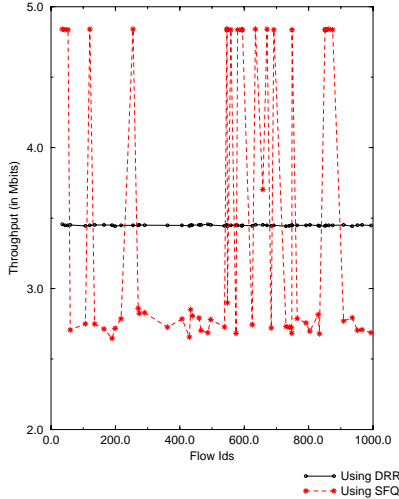


Figure 10: Comparison of DRR and SFQ: SFQ offers higher bandwidth to flows sending large packets, while DRR offers equal bandwidth to all flows.

time traffic, the simplest solution is to reserve a portion of the bandwidth for real-time traffic and use a separate Fair Queuing algorithm for the real-time traffic while continuing to use DRR for the best-effort traffic. This allows efficient packet processing for best-effort traffic; at the same time it allows the use of other fair queuing schemes that provide delay bounds for real-time traffic at reasonable cost.

As a simple example of combining fair queuing schemes, consider the following modification of Deficit Round Robin called DRR+. In DRR+ there are two classes of flows: *latency critical* and *best-effort*. A latency critical flow must contract to send no more than x bytes in some period T . If a latency critical flow f meets its contract, whenever a packet for flow f arrives to an empty flow f queue, *the flow f is placed at the head of the round-robin list*.

Suppose for instance that each latency critical flow guarantees to send at most a single packet (of size at most s) every T seconds. Assume that T is large enough to service one packet of every latency critical flow as well as one quantum's worth for every other flow. Then if all latency critical flows meet their contract, it appears that each latency critical flow is at most delayed by $((n' * s) + Max)/B$, where n' is the number of latency critical flows. In other words, a latency critical flow is delayed by one small packet from every other latency critical flow, as well as an error term of one maximum size packet (the error term is inevitable in all schemes unless the router preempts large packets). In this simple case, the final bound appears better than the DKS bound because a latency critical flow is only delayed by other latency critical flows.

In the simple case, it is easy to police the contract for latency critical flows. A single bit that is part of the state of such a flow is cleared whenever a timer expires and is set whenever a packet arrives; the timer is reset for T time units when a packet arrives. Finally, if a packet arrives and the bit is set, the flow has violated its contract; an effective (but user-friendly) countermeasure is to place the flow ID of a deviant flow at the end of the round-robin list. This effectively moves the flow from the class of latency critical flows to the class of best effort flows.

Figure 11 shows the results of a simulation experiment using DRR+ instead of DRR. The simulation parameters are as follows. We use a single router configuration as described earlier. The structure of

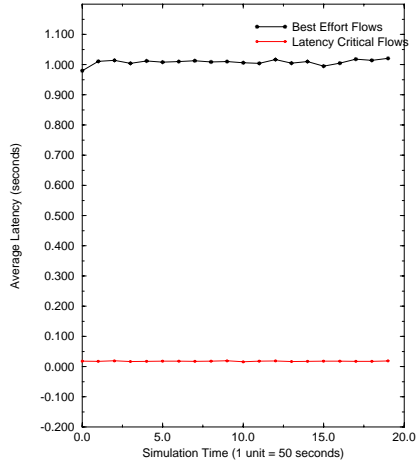


Figure 11: The latency experienced by latency critical flows and best-effort flows when there are three latency critical flows through the router. The algorithm used for servicing the queues was DRR+.

a BONEs host for these experiments is as follows. Each host has 20 flows, out of which one is marked as latency critical. Therefore, to have three *LatencyCritical* flows, we use three hosts. The default parameters are set in the DRR+ router and the hosts (The *LatencyCritical* and *BestEffort* flows send packets with exponential interarrival times at an average of 10 packets/sec and they have random packet sizes). The delay experienced by the *LatencyCritical* flows is averaged into 20 batch intervals spread over the simulation time. The rest of the flows (i.e., the *BestEffort* flows) are also batched together and averaged.

It is evident from the figure that the *LatencyCritical* flows experience bounded latency in a DRR+ environment. However, note that this experiment only considered a single router. DRR+ is a simple example of combining DRR with other fair queuing algorithms to handle latency critical flows. By using other schemes for the latency critical flows, we can provide better bounds while allowing more general traffic shaping rules.

7 Other Applications of DRR

Deficit round robin schemes can be applied to other scheduling contexts in which jobs must be serviced as whole units. In other words, jobs cannot be served in several “time slices” as in a typical operating system. This is true for packet scheduling because packets cannot be interleaved on the output lines, but it is true for other contexts as well. Note also that DRR is applicable to distributed queues because it needs only local information to implement. We describe two other specific applications: token ring protocols and load balancing schemes.

The current 802.5 token ring uses token holding timers that limit the number of bits a station can send at each token opportunity. If a station’s packets do not fit exactly into the allowed number of bits, the remainder is not kept track off; this allows the possibility of unfairness. The unfairness can easily be removed by keeping a deficit counter at each ring node and by a small modification to the token ring protocol.

Another application is load balancing or, as it is sometimes termed, striping. Consider a router that

has traffic arriving on a high speed line that needs to be sent out to a destination over multiple slower speed lines. If the router sends packets from the fast link in a round robin fashion across the slower links, then the load may not balance perfectly if the packet sizes are highly variable. For example, if packets alternate between large and small packets, then round robin across two lines can cause the second line to be underutilized. But load balancing is almost the inverse of fair queuing. It is not hard to see that deficit round robin solves the problem; we send up to a quantum limit per output line but we keep track of deficits. This should produce nearly perfect load balancing; as usual it can be extended to weighted load balancing. In [APV95], we show how to obtain perfect load balancing and yet guarantee FIFO delivery. Our load balancing scheme [APV95] appears to be a novel solution to a very old problem.

8 Conclusions

We have described a new scheme, Deficit Round Robin (DRR), that provides near-perfect isolation at very low implementation cost. As far as we know, this is the first fair queuing solution that provides near-perfect throughput fairness with $O(1)$ packet processing. DRR should be attractive to use while implementing fair queuing at gateways and routers.

A number of readers have conjectured that DRR should reduce to bit-by-bit round robin (BR) when the quantum size is 1 bit. This is not true! The two schemes have radically different behaviors. A counterexample is described in Section 3. The counterexample also indicates that reducing the quantum size to be less than Max does not improve the worst-case latency at all. Thus DRR is not just a crude approximation of BR; it has more interesting behaviors of its own.

We have described theorems that describe the behavior of DRR in ensuring throughput fairness. In particular, we show that DRR satisfies Golestani's [Gol94] definition of throughput fairness — i.e., the normalized bandwidth allocated to any two backlogged flows in any interval is roughly equal. Earlier versions of this paper [SV95, S93] had only shown throughput fairness for the case when all flows were backlogged. Our use of Golestani's definition makes the results in this paper more general, and indicates why DRR works well in non-backlogged and backlogged scenarios. Our simulations indicate that DRR works well in all scenarios.

The *Quantum* size is required to be at least Max for the work complexity to be $O(1)$. We feel that while Fair Queuing using DRR is general enough for any kind of network, it is best suited for datagram networks. In ATM networks, packets are fixed size cells; therefore Nagle's solution (simple round robin) will work as well as DRR. However, if connections in an ATM network require weighted fair queuing with arbitrary weights, DRR will be useful.

DRR can be combined with other fair queuing algorithms such that DRR is used to service only the best-effort traffic. We described a trivial combination algorithm called DRR+ that offers good latency bounds to Latency Critical flows as long as they meet their contracts. However, even if the source meets the contract, the contract may be violated due to “bunching” effects at intermediate routers. Thus other combinations need to be investigated. Recall that DRR requires having the quantum size be at least a maximum packet size in order for the packet processing work to be low; this does affect delay bounds.

We believe that DRR should be easy to implement using existing technology. It only requires a few instructions beyond the simplest queuing algorithm (FCFS), and this addition should be a small

percentage of the instructions needed for routing packets. The memory needs are also modest; 6K size memory should give a small number of collisions for about 100 concurrent flows. This is a small amount of extra memory compared to the buffer memory used in many routers. Note that the buffer size requirements should be identical to the buffering for FCFS because in DRR buffers are shared between queues using McKenney's buffer stealing algorithm.

Deficit round robin schemes can be applied to other scheduling contexts in which jobs must be serviced as whole units. We have described two other applications: token rings with holding timers and load balancing. For these reasons, we believe that deficit round robin scheduling is a general and useful tool. We hope our readers will use it in other ways.

Acknowledgments: A number of people listened patiently to our ideas and provided us with valuable feedback. Prominent among them are: Dave Clark, Sally Floyd, Andy Fingerhut, Srinivasan Keshav, Paul McKenney, and Lixia Zhang. We thank Adam Costello, Andy Fingerhut and S. Keshav for their careful reading of the paper. We thank S. Keshav and the anonymous SIGCOMM referees who prodded us into using Golestani's fairness definition. We thank R. Gopalakrishnan and Apostolos Dalianis for valuable discussions.

References

- [APV95] H. Adishesu, G. Parulkar, and G. Varghese. *Reliable FIFO Load Balancing over Multiple FIFO Channels*. Washington University Technical Report 95-11, available by FTP.
- [CLR90] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. MIT Press/McGraw-Hill, 1990.
- [DKS89] A. Demers, S. Keshav, and S. Shenker. Analysis and simulation of a fair queueing algorithm. *Proc. Sigcomm '89*, 19(4):1–12, September 1989.
- [FP95] N. Figuera and J. Pasquale Leave-in-time: A new service discipline for real-time communication in a packet-switching data network. *Proc. Sigcomm '95*, September 1995.
- [Flo93a] S. Floyd. Notes on guaranteed service in resource management. Unpublished note. 1993.
- [Flo93b] S. Floyd. Personal communication. 1993.
- [FJ95] S. Floyd and V. Jacobson Link-Sharing and Resource Management Models for Packet Networks. In *IEEE Trans. Networking*, August 95.
- [GM90] A. Greenberg and N. Madras. How fair is fair queueing. In *Proc. Performance '90*, 1990.
- [Gol94] S. Golestani. A self clocked fair queueing scheme for broadband applications. In *Proc. IEEE Infocomm '94*, 1994.
- [JR86] R. Jain and S. Routhier. Packet trains measurement and a new model for computer network traffic. *IEEE Journal on Selected Areas in Communications*, Sept. 1986.
- [Kes91] S. Keshav. On the efficient implementation of fair queueing. In *Internetworking: Research and Experience Vol.2, 157-173*, September 1991.

- [McK91] P. McKenney. Stochastic fairness queueing. In *Internetworking: Research and Experience Vol.2, 113-131*, January 1991.
- [Nag87] John Nagle. On packet switches with infinite storage. *IEEE Trans. on Comm.*, COM-35(4), April 1987.
- [PG93] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks. In *Proc. IEEE Infocomm '93*.
- [SA94] D. Saha and M. Saksena and S. Mukherjee and S. Tripathi. On Guaranteed Delivery of Time-Critical Messages in DQDB. In *Proc. IEEE Infocomm '94*.
- [S93] M. Shreedhar. *Efficient Fair Queuing Using Deficit Round Robin*. M.S. Thesis, Dept. of Computer Science, Washington University. December 1994.
- [SV95] M. Shreedhar and G. Varghese. *Efficient Fair Queuing Using Deficit Round Robin*. Older version of this paper, to appear in Proceedings SIGCOMM 95, Boston, August 1995.
- [Zha91] L. Zhang. Virtual clock: A new traffic control algorithm for packet switched networks. *ACM Trans. on Computer Systems*, 9(2):101–125, May 1991.